Local Structure for Vertex-Minors

by

Rose McCarty

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, 2021

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: Dan Král'
Professor, Faculty of Informatics, Masaryk University

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

Internal Member: Peter Nelson
Associate Professor, Dept. of Combinatorics and Optimization,
University of Waterloo

Internal Member: Bertrand Guenin
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

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the 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.

Statement of Contributions

This thesis is based on an ongoing project which is joint work with Jim Geelen and Paul Wollan. Chapter 2 is based on a published paper which is also joint with O-joung Kwon [64]. The explanation of the connection with quantum computing is based on an unpublished manuscript with James Davies.

Abstract

This thesis is about a conjecture of Geelen on the structure of graphs with a forbidden vertex-minor; the conjecture is like the Graph Minors Structure Theorem of Robertson and Seymour but for vertex-minors instead of minors. We take a step towards proving the conjecture by determining the "local structure".

Our first main theorem is a grid theorem for vertex-minors. We prove that any graph of sufficiently large rank-width has a big comparability grid as a vertex-minor. Equivalently, a class of graphs has unbounded rank-width if and only if it contains all circle graphs as vertex-minors, up to isomorphism.

Our second main theorem is more like the Flat Wall Theorem of Robertson and Seymour. Given a graph of large rank-width in a proper vertex-minor-closed class, we describe how the rest of the graph "attaches" onto a circle graph that it contains. Informally, this theorem says that the attachments are almost compatible with the circle graph, relative to a large comparability grid.

We believe that the results presented in this thesis provide a path towards proving the full conjecture. To make this area more accessible, we have organized the first chapter as a survey on "structure for vertex-minors".

Acknowledgements

I am grateful to so many people for their support that I can't do everyone justice. I would like to thank Troy Peace for inspiring my self-confidence, Dana Randall for showing me the path to a PhD, Anant Godbole for introducing me to the joy of research, and Xingxing Yu for teaching me how to plan an effective project. I cannot express my gratitute to Robin Thomas strongly enough; I would not be the researcher I am today without his encouragement to pursue my own ideas. I wish I could tell him that myself.

I'd like to thank Marthe Bonamy, O-joung Kwon, Sang-il Oum, Zdeněk Dvořák, Sergey Norin, Bartosz Walczak, and Tomasz Krawczyk for the fantastic visits. I hope we get the chance to work together again. Thanks to my local colleagues and conference collaborators for many exciting discussions; I'm excited for in-person conferences to resume. On a related note, thanks to Marcin and Michał Pilipczuk for continuing to try to work together despite the many cancellations. Shoutout to my local collaborator James Davies; I cannot wait to hear your next crazy brilliant idea.

Thanks to Louis Esperet for the feedback, and for pointing out the connection between Problem 4.1.1 and the work of Máčajová and Škoviera. Thanks to Patrice Ossona de Mendez for the conversations on sparsity, and for finding too many corrections to this thesis. Thanks to Ben Moore for introducing me to sparsity. Thanks to all of my committee members – Dan Král', Peter Nelson, Bertrand Guenin, and Anna Lubiw – for your feedback and mentorship.

Thanks to Jim for being the best supervisor I could have asked for. I remember being mildly concerned that your students only had good things to say (despite my promting) at grad visit day. Now I understand why.

I'd like to thank all of my friends and family for their support in this journey. Thanks to Brett Nasserden and Smile Tiger for keeping me sane, to Sina Baghal for helping me keep perspective, and to the Smathers for making me a part of the family.

Without my partner Weston none of this would have been possible; your support means the world to me.

Table of Contents

Chapter 1

Structure for vertex-minors

1.1 Introduction

Vertex-minors have been discovered several times independently, in biology [14, 48] and quantum computing [36] as well as in mathematics [4, 7, 99]. Brijder and Hoogeboom [14] showed how to model gene assembly in ciliates using transformations that relate to vertex-minors. Van den Nest, Dehaene, and De Moor [36] showed how to use vertex-minors to determine if two quantum graph states are equivalent (up to local Clifford operations). Originally, vertex-minors arose from the work of Bouchet [4, 7] on isotropic systems, although Oum [99] gave them their present, descriptive name. Very roughly, the vertex-minors of a graph G are those graphs that can be obtained from G by performing "local complementations at vertices" and by deleting vertices.

Vertex-minors typically yield dense graph classes; since we are only allowed to delete vertices but not edges, almost no graphs are vertex-minors of cliques. This is very different from most other notions of "minors", like graph minors [111], topological minors [44, 68], and immersion minors [38, 46, 124]. So in some ways vertex-minors behave more like induced subgraphs or induced subdivisions than "minors". Yet Oum conjectures that vertex-minors have the well-quasi-ordering property.

Well-Quasi-Ordering Conjecture (Oum [103, 102, Question 6])**.** *Every infinite set of graphs contains one graph that is isomorphic to a vertex-minor of another.*

Equivalently, the conjecture states that every proper vertex-minor-closed class can be characterized by a finite set of forbidden vertex-minors.

The Well-Quasi-Ordering Conjecture immediately brings to mind Wagner's conjecture on well-quasi-ordering for graph minors, which was famously proven by Robertson and Seymour [112] in the twentieth paper of their graph minors series. Remarkably, the two conjectures are in fact related; well-quasi-ordering for pivot-minors [102, Question 6] would imply them both (see Section 1.6).

On the algorithmic side, well-quasi-ordering is useful because it oftentimes leads to efficient algorithms for membership testing; we just need to test for a fixed graph as a vertex-minor. Oum conjectures the following.

Membership Testing Conjecture (Oum [103, 102, Questions 6 and 7])**.** *For any proper vertex-minor-closed class of graphs $\mathcal{F}$, there is a polynomial-time algorithm that determines if a given graph is in $\mathcal{F}$.*

It even seems that the degree of the polynomial should be independent of $\mathcal{F}$. This conjecture is again motivated by graph minors; Robertson and Seymour [110] proved that there is a cubic algorithm (in the number of vertices) for testing membership within any fixed minor-closed class. The running time has since been improved to quadratic [78], although these algorithms are still extraordinarily non-practical (which is interesting in its own right).

The cornerstone theorem of Robertson and Seymour's graph minors series is their Graph Minors Structure Theorem [111], which gives a constructive description of the graphs in any proper minor-closed class. Informally, this theorem says that such graphs "decompose" into parts that "almost embed" in a surface of bounded Euler genus. This theorem provides the starting point for proving both well-quasi-ordering and membership testing for graph minors.

So, motivated by the above conjectures, the topic of this thesis is an analogous structural conjecture for proper vertex-minor-closed classes. This conjecture was proposed by Geelen [59], and the thesis consists of joint work with Jim Geelen and Paul Wollan towards proving the conjecture. Chapter 2 is based on a paper that is also joint work with O-joung Kwon [64]. We postpone the formal statement of the conjecture until Section 1.5, but here is the idea.

Structural Conjecture (Informal statement – Geelen [59])**.** *For any proper vertex-minor-closed class of graphs $\mathcal{F}$, each graph in $\mathcal{F}$ "decomposes" into parts that are "almost" circle graphs.*

So, informally, the conjecture says that any proper vertex-minor-closed class of graphs $\mathcal{F}$ is contained in a constructively-defined class of graphs $\mathcal{F}'$ whose closure under vertex-minors remains proper. "Decomposes" and "almost" have very different meanings in this

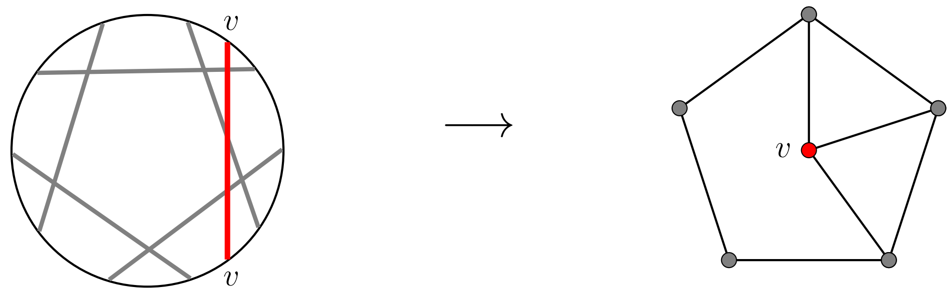

Figure 1.1: Chords of a circle (left, with a chord v labelled at its ends) and the corresponding circle graph (right).

conjecture than for graph minors; intuitively, the difference is that we care about the adjacency matrix rather than the incidence matrix. Moreover, the "basic graphs" have geometric rather than topological representations; a *circle graph* is the intersection graph of chords of a circle (so there is a vertex for each chord, and two chords are adjacent if they intersect; see Figure 1.1).

The rest of the introduction is loosely organized as a survey on "structure for vertex-minors". Throughout the thesis, we follow Diestel [39] for standard graph theoretic notation, and all graphs are finite, simple, and loopless unless otherwise noted.

1.2 Contributions of the thesis

Recall that our goal is to prove something like the Graph Minors Structure Theorem, but for vertex-minors instead of graph minors. The first step in Robertson and Seymour's proof of the Graph Minors Structure Theorem is their Grid Theorem [108], which says that for any planar graph H, every graph of sufficiently large tree-width has a minor isomorphic to H. In Chapter 2 we prove the following analogous theorem for vertex-minors.

The Grid Theorem for Vertex-Minors (Geelen, Kwon, McCarty, and Wollan [64])**.** *For any circle graph H, there exists an integer r_H so that every graph with rank-width at least r_H has a vertex-minor isomorphic to H.*

This theorem characterizes when a graph has large rank-width; rank-width does not increase when taking vertex-minors, and there are circle graphs of arbitrarily large rank-width. As in the Grid Theorem of Robertson and Seymour, we usually work with particular circle graphs of large rank-width called "comparability grids". So, in order to prove the Structural Conjecture, we just need to determine the structure of a graph relative to a large comparability grid. (It is more standard, although essentially equivalent [114], to work relative to a tangle as in [109]. However, we choose to work with a comparability grid for the sake of simplicity. The "global structure" can be recovered from the "relative structure" using the Tree of Tangles Theorem in [61, Theorem 9.1].)

The next main step is to determine the "local structure", as in the Flat Wall Theorem of Robertson and Seymour [110]. Consider a graph with a large grid minor in a proper minor-closed class; the Flat Wall Theorem roughly says that there is a planar subgraph containing much of the grid so that the rest of the graph "almost" attaches onto just the outer face. As in the Grid Theorem for Vertex-Minors, the key piece of the analogy is between planar graphs and circle graphs; we are interested in how the rest of the graph "attaches" onto a circle graph containing much of the comparability grid. Our notion of "attachment" is much more "local" than in the Flat Wall Theorem, however; it just refers to edges with an end in the circle graph.

Furthermore, instead of finding one particular circle graph where the "attachments" are well-behaved, we work with an arbitrary circle graph containing a comparability grid. This approach lets us grow the circle graph and adjust the comparability grid over time. For technical reasons having to do with unique representations of circle graphs (see [5] and [52]), it is convenient to assume that the circle graph is prime in the sense of Cunningham [29]. This is a minor connectivity-like condition that is analogous to 3-connectivity. The final main theorem of this thesis is the following, informally stated.

The Local Structure Theorem (Informal statement - Geelen, McCarty, Wollan)**.** *For any proper vertex-minor-closed class of graphs $\mathcal{F}$, and for any graph in $\mathcal{F}$ with an induced subgraph that is a prime circle graph containing a comparability grid, the rest of the graph "almost attaches" to the circle graph in a way that is "mostly compatible" with the comparability grid.*

We give a formal statement of the theorem in Section 3.3; it relies on a respresentation of circle graphs by connected 4-regular graphs due to Kotzig [83]. This representation shows that vertex-minors of circle graphs are roughly equivalent to immersion minors of connected 4-regular graphs. While the precise statement is somewhat technical, the main point is simple; for circle graphs we can reduce to a notion of "minor" that is already very well-understood. (In particular, Robertson and Seymour [113] proved well-quasi-ordering for weak immersion minors, verifying a conjecture of Nash-Williams [95]. There are also very short proofs regarding structure for both weak [38, 124] and strong [46] immersion minors.

For 4-regular graphs, weak and strong immersion minors are equivalent.) Our approach to the Structural Conjecture is based on this surprising connection with immersion minors.

The connection goes beyond just circle graphs; Bouchet [10] showed how to use a signature on the 4-regular graph to represent one additional vertex (that is, a vertex whose deletion results in a circle graph). This approach is very similar to Gerards' [67] proof of the forbidden minors for graphic matroids (which was originally proven by Tutte [117]). So, for each vertex v outside of the circle graph, we will have a signature on the 4-regular graph; that signature represents the neighbourhood of v within the circle graph.

Outline of the thesis

The rest of this chapter is a survey on "structure for vertex-minors". Some of the fundamental results discussed here will be used in later chapters.

In Chapter 2 we prove the Grid Theorem for Vertex-Minors. This theorem is not explicitly used in this thesis. However, in the future it will let us apply the Local Structure Theorem.

In Chapter 3 we give the formal statement of the Local Structure Theorem and outline its proof. The key idea is that we can represent a circle graph and its "attachments" by an associated graph called the "labelled tour graph". Bouchet [10] and Kotzig [83] discovered this representation and showed that it "efficiently captures" vertex-minors. We show how to use their theorem (Theorem 3.4.2) to work entirely in the labelled tour graph.

In Chapter 4 we prove a precise min-max theorem in the simplest case: when the labelled tour graph has just one signature. In fact we prove a somewhat more general result. We characterize the maximum number of non-zero circuits in a "rooted" circuit-decomposition of a signed Eulerian graph.

Finally, in Chapter 5 we use the results from the previous two chapters to prove the Local Structure Theorem. We then briefly return to the overall question of "structure for vertex-minors" in Section 5.6.

1.3 Vertex-minors and the main ingredients

For a vertex v of a graph G, *locally complementing at v* replaces the induced subgraph on the neighbourhood of v by its complement. We denote the new graph by $G * v$, as depicted in Figure 1.2. Then a graph H is a *vertex-minor* of a graph G if H can be obtained from

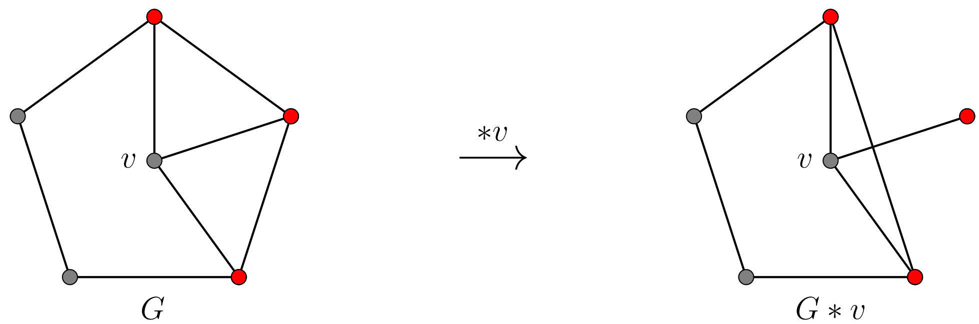

Figure 1.2: A graph before and after locally complementing at v, with the neighbours of v in red.

G by a sequence of vertex deletions and local complementations (in any order). If H can be obtained from G by local complementations alone, then H and G are *locally equivalent*. This is an equivalence relation because $G * v * v = G$. It is important to notice that all of the local complementations can be performed first; any vertex-minor of G is an induced subgraph of a graph that is locally equivalent to G.

So vertex-minors arise when we care not just about an individual graph, but about its whole local equivalence class. In Section 1.6 we will see that local equivalence classes can be further refined into "pivot-equivalence" classes. Given any graph G, we can construct a pivot-equivalence class that has one graph for each spanning tree of G. Deleting a vertex from the pivot-equivalence class corresponds to deleting or contracting an edge from G, depending on whether or not the edge is in the spanning tree. This hints at the remarkable connection with graph minors.

Next we will discuss two properties that are invariant under local complementation. The Structural Conjecture then says that, in some sense, these are the only two examples.

Circle graphs

Recall that a *circle graph* is the intersection graph of chords on a circle; we call a collection of chords of the unit circle a *chord diagram*. The class of circle graphs is closed under taking vertex-minors. To delete a vertex, we just delete its chord. To locally complement at a vertex v, we "flip" one of the two arcs of the circle that has the same ends as the chord v (see Figure 1.3, where we flip the arc on the right of v). In general we allow two chords to have a common end on the circle, but when applying this argument we first perturb the chords slightly so as to avoid this; it is always possible to do so.

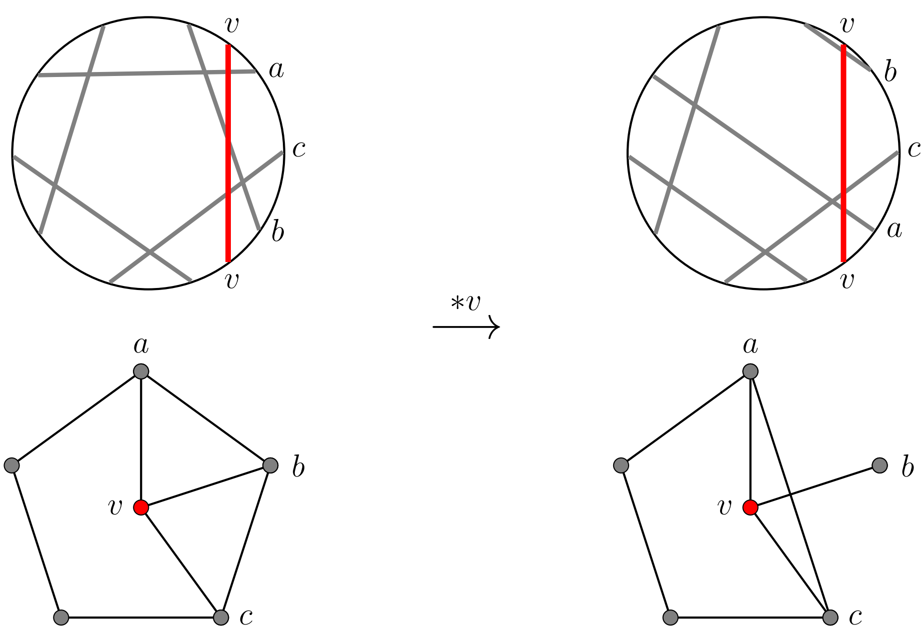

Figure 1.3: Chords on a circle (top) and the corresponding circle graphs (bottom), before and after locally complementing at v.

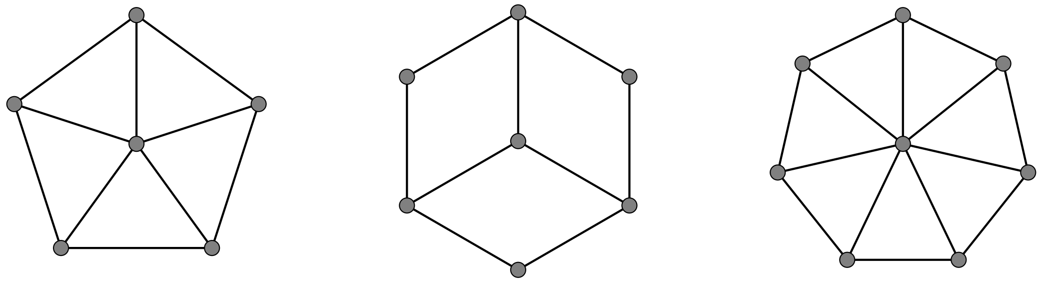

Figure 1.4: The three forbidden vertex-minors for circle graphs.

Since the class of circle graphs is closed under taking vertex-minors, we could hope to characterize circle graphs in a similar manner to Kuratowski's characterization of planar graphs. Bouchet [10] proved such a theorem.

Theorem 1.3.1 (Bouchet [10])**.** *A graph is a circle graph if and only if it has no vertex-minor isomorphic to any of the three graphs in Figure 1.4.*

In fact, Geelen and Oum [66] proved a common generalization of Bouchet's theorem and Kuratowski's theorem. Their theorem characterizes circle graphs by forbidden pivot-minors; the remarkable connection between planar graphs and circle graphs is due to de Fraysseix [34]. Unfortunately though, Geelen and Oum's theorem is difficult to state; there are 15 obstructions which were found by computer search.

There are other characterizations of circle graphs, such as Naji's [94] algebraic characterization (also see the proofs in [65] and [116]) and Brijder and Traldi's [17, 18] characterization via representations of a multimatroid. However, there is no known characterization of circle graphs by forbidden induced subgraphs (see the survey by Durán, Grippo, and Safe [41]).

Cut-rank

Cut-rank is a function that measures the "complexity" of each cut in a graph. The most standard such measure is the number of edges in a cut; cut-rank instead considers rank in the adjacency matrix. The advantage is that strictly more cuts have "low complexity" according to their cut-rank.

For this reason, cut-rank is particularly useful for algorithmic applications; it suggests a way of generalizing the "divide-and-conquer" approach (see [87]) beyond classes of graphs

$$\xrightarrow{*v}$$

$$\begin{array}{c} \\ v \\ a \\ b \\ c \\ d \\ e \end{array}\begin{array}{c} \begin{array}{cccccc} v & a & b & c & d & e \end{array} \\ \begin{bmatrix} 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \end{bmatrix} \end{array}\begin{array}{c} \\ \vec{r} \\ \\ \\ \\ \\ \\ \end{array} \qquad \begin{array}{c} \\ v \\ a \\ b \\ c \\ d \\ e \end{array}\begin{array}{c} \begin{array}{cccccc} v & a & b & c & d & e \end{array} \\ \begin{bmatrix} 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & \cdot & 0 & 1 & 0 & 1 \\ 1 & 0 & \cdot & 0 & 0 & 0 \\ 1 & 1 & 0 & \cdot & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \end{bmatrix} \end{array}\begin{array}{c} \\ \\ +\vec{r} \\ +\vec{r} \\ +\vec{r} \\ \\ \\ \end{array}$$

Figure 1.5: A graph (top) and its adjacency matrix (bottom), before and after locally complementing at v.

with small separators. This idea has been very effective; many problems that can be solved using dynamic programming on graphs of bounded tree-width [26] can likewise be solved on graphs of bounded rank-width [27, 104]. We will define rank-width later; for now let us just say that it is similar to branch-width but defined using the cut-rank function. This approach was introduced by Oum and Seymour [104].

Now here are the definitions. Let G be a graph, and recall that its *adjacency matrix*, denoted Adj_G, is the $V(G) \times V(G)$ matrix whose (u, v)-entry is 1 if $uv \in E(G)$ and 0 otherwise. In this thesis every matrix is over the binary field. This convention is important; it makes cut-rank behave appropriately under local complementations. The *cut-rank* of a set $X \subseteq V(G)$, denoted $\rho_G(X)$ (or just $\rho(X)$ if the graph is clear from context), is the rank of the submatrix of the adjacency matrix with rows X and columns $V(G) - X$. The cut-rank function is *symmetric*; that is, $\rho(X) = \rho(V(G) - X)$ for any $X \subseteq V(G)$.

Moreover, cut-rank is invariant under local complementation [99, Proposition 2.6]. To see this, consider what happens in the adjacency matrix when we locally complement at a

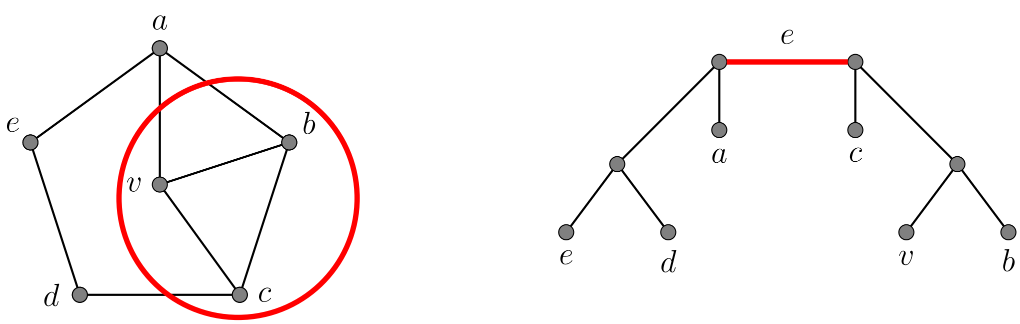

Figure 1.6: A graph (left, with the set of leaves of the rightmost component of $T-e$ circled in bold red) and a rank-decomposition T (right, with e in bold red).

vertex v. As depicted in Figure 1.5, the row of v is added to the rows of its neighbours, and then the diagonal is corrected. (It is alright to correct the diagonal because cut-rank only "sees" the off-diagonal entries.) So $\rho(X)$ does not change for any $X \subseteq V(G)$ that contains v; this is enough because cut-rank is symmetric.

Combining the examples

The Structural Conjecture says that every graph in a proper vertex-minor-closed class "decomposes" into parts that are "almost" circle graphs. Both "decomposes" and "almost" have to do with the cut-rank function. Since the decomposition step is fairly technical, for now we will just give the definition of rank-width instead. The informal intuition is that a graph has low rank-width if it "successively decomposes" (along cuts of low cut-rank) into the 1-vertex graph.

Here are the formal definitions; refer to Figure 1.6 for an example. First, a *rank-decomposition* of a graph G is a tree T, with $V(G)$ as its set of leaves, so that each vertex of T has degree either one or three. The *width* of an edge e of T is the cut-rank, in G, of the set of all leaves of one of the components of $T-e$. (The choice of component does not matter since cut-rank is symmetric.) Finally, the *rank-width* of G is the minimum, over all rank-decompositions T of G, of the maximum width of an edge of T. Graphs with at most one vertex do not technically admit rank-decompositions, so we define their rank-width to be zero.

Because cut-rank is invariant under local complementation, if H is a vertex-minor of G, then the rank-width of H is at most the rank-width of G. So for each $k \in \mathbb{N}$, the class of graphs with rank-width at most k is closed under taking vertex-minors. Oum [99] proved that each such class can be characterized by a finite number of forbidden vertex-minors. Furthermore, the Grid Theorem for Vertex-Minors gives a rough characterization.

Corollary 1.3.2. *A class of graphs has bounded rank-width if and only if it does not contain all circle graphs as vertex-minors, up to isomorphism.*

So we have two main examples of proper vertex-minor-closed classes: the class of circle graphs and, for each $k \in \mathbb{N}$, the class of graphs of rank-width at most k.

Now we formalize the notion of "almost"; this notion will let us create new proper vertex-minor-closed classes. So, a *rank-p perturbation* of a graph G is a graph whose adjacency matrix can be obtained from Adj_G by first adding (over the binary field) a symmetric matrix of rank at most p, and then changing all diagonal entries to be 0. Thus the cut-rank of each set changes by at most p. In Lemma 1.6.7, we will prove that for any $p \in \mathbb{N}$ and any proper vertex-minor-closed class $\mathcal{F}$, there exists a graph that is not isomorphic to a vertex-minor of any rank-p perturbation of a graph in $\mathcal{F}$. So this operation allows us to obtain a new proper vertex-minor-closed class. (Any graph in the new class is also a low-rank perturbation of a graph in $\mathcal{F}$, but the precise rank can go up.)

Nguyen and Oum [98] showed that a number of notions are approximately equivalent to low-rank perturbations. The most explicit definition comes from *complementing on a set* $X \subseteq V(G)$; this means to replace the induced subgraph on X by its complement. Every graph that can be obtained from G by complementing on p sets (that is, by repeating this operation p times) is a rank-p perturbation of G. In the other direction, every rank-p perturbation of G can be obtained from G by complementing on $\lfloor 3p/2 \rfloor$ sets. This follows from the next lemma; note that any rank-2 matrix can be written as the sum of three symmetric rank-1 matrices. (Furthermore, while we work over the binary field, this lemma actually holds for any field. It is related to Bunch–Kaufman decompositions for "nearly diagonalizing" symmetric matrices; see [19].)

Lemma 1.3.3. *For any $p > 0$, every symmetric rank-p matrix can be written as the sum of p_1 rank-1 matrices and p_2 rank-2 matrices, also symmetric, so that $p = p_1 + 2p_2$.*

Proof. We proceed by induction on p. We may assume that $p \geq 3$ as otherwise the lemma trivially holds. Let A denote the matrix in question, and suppose first that A has a non-zero entry on its diagonal. Then, where x denotes the corresponding column vector, the matrix $A - xx^\intercal$ has lower rank than A. So we may assume that the diagonal of A has no non-zero entries. Now consider an arbitrary non-zero entry A_{ij}; let x and y be the column vectors on columns i and j. Then $A - (xy^\intercal + yx^\intercal)$ has rank 2 less than the rank of A. This now completes the proof by induction on p. □

The examples of proper vertex-minor-closed classes discussed in this section give good intuition for the Structural Conjecture. The only piece still missing is the full "decomposition" step.

1.4 Motivating conjectures

In this section we discuss some conjectures that we hope can be approached using the Structural Conjecture. For a conjecture about proper vertex-minor-closed classes to seem reasonable and approachable using the Structural Conjecture, we would like for the following to hold.

(i) The conjecture is true for the class of circle graphs.

(ii) The conjecture is true for any class of graphs of bounded rank-width.

(iii) If the conjecture holds for a class of graphs $\mathcal{F}$, then it also holds for the class of all graphs that can be obtained from a graph in $\mathcal{F}$ by complementing on a set.

These three conditions are not quite enough overall because there is still the technical "decomposition" step in the Structural Conjecture. However, we believe that these three conditions provide good evidence that a conjecture is likely to be true. In fact, the third condition is most important to verify when $\mathcal{F}$ is the class of circle graphs.

Well-quasi-ordering and membership testing

Recall the Well-Quasi-Ordering Conjecture of Oum [102, 103] from Section 1; the conjecture says that every infinite set of graphs contains one graph that is isomorphic to a vertex-minor of another. Oum [100] proved this conjecture for classes of bounded rank-width; in fact he proved a stronger statement about pivot-minors. The conjecture also holds for the class of circle graphs. This fact is a corollary of well-quasi-ordering for immersion minors [113]; Kotzig [83] showed that vertex-minors of circle graphs are related to immersion minors of 4-regular graphs.

We need new techniques to deal with complementing on a set. The standard approach is to allow vertices (or sometimes edges) to be labelled by a fixed well-quasi-order. Then we would aim to use the fact that perturbations behave well with respect to taking vertex-minors (see Lemma 1.6.7). For now, however, this case is open.

Next recall the Membership Testing Conjecture of Oum [102, 103] from Section 1. The conjecture says that, for any fixed proper vertex-minor-closed class of graphs $\mathcal{F}$, there is a polynomial-time algorithm that determines if a given graph is in $\mathcal{F}$. This conjecture holds for classes of bounded rank-width by combining the well-quasi-ordering result of Oum [100] with the approach of Courcelle and Oum [28]. The conjecture also holds for the class of circle graphs [6, 52]. The general problem, however, likely requires a different approach.

Bounding the chromatic number

A graph class is *χ-bounded* if the chromatic number of each graph in the class is bounded above by some fixed function of its clique number. Such a function is called a *χ-bounding function*. In general it is very difficult to understand what else, besides a large clique, could possibly make the chromatic number of a graph large. Yet these graphs are surprisingly common; the classic random construction of Erdös [50] yields graphs of large chromatic number and large girth. So we study χ-bounded graph classes in order to understand what structures these graphs must contain; see the survey by Scott and Seymour [115].

For vertex-minors, Davies [31] recently proved the following theorem, which was conjectured by Geelen (see [43, Conjecture 1]).

Theorem 1.4.1 (Davies [31])**.** *Every proper vertex-minor-closed class of graphs is χ-bounded.*

Strikingly, Esperet [51, Conjecture 2.3.19] conjectures that every χ-bounded graph class which is closed under deleting vertices has a very efficient χ-bounding function: one that is a polynomial. Such a graph class is called *polynomially χ-bounded*. It is difficult to believe Esperet's conjecture, yet the conjecture remains open despite a great deal of effort aimed at disproving it. However, for vertex-minor-closed classes it seems likely to be true.

Polynomial χ-boundedness Conjecture (Esperet [51], Kim, Kwon, Oum, and Sivaraman [79])**.** *Every proper vertex-minor-closed class of graphs is polynomially χ-bounded.*

Polynomial χ-boundedness is particularly interesting because it implies the famous Erdős-Hajnal property (that there exists $\epsilon > 0$ so that every n-vertex graph in the class has a clique or stable set of size at least n^ϵ). It is unknown whether or not χ-boundedness alone implies the Erdős-Hajnal property (even though the Erdős-Hajnal Conjecture says that the property holds as long as an induced subgraph is forbidden); see the survey by Chudnovsky [20]. However, Chudnovsky and Oum [23] have already proven that every proper vertex-minor-closed class has the Erdős-Hajnal property.

We believe that the Structural Conjecture would imply the Polynomial χ-boundedness Conjecture. McCarty and Davies [33] proved that circle graphs are polynomially χ-bounded; in fact Davies [32] very recently found a χ-bounding function which is tight up to a constant factor. This improved on the singly exponential bound of Kostochka and Kratochvíl [82] and the original proof of Gyárfás [69]. Furthermore, Bonamy and Pilipczuk [1] proved that classes of bounded rank-width are polynomially χ-bounded (in fact they took care of the full "decomposition" step). This result improved the theorem of Dvořák and Král' [43] that such classes are χ-bounded.

For complementing on a set, first we consider complementing the entire graph. Kostochka and Kratochvíl [82] showed that the complements of circle graphs are polynomially χ-bounded. The complements of graphs of bounded rank-width still have bounded rank-width; so they are polynomially χ-bounded as well. Moreover, the following equivalent version of the conjecture is "stable" under complementing on a set; for any proper vertex-minor-closed class $\mathcal{F}$, both $\mathcal{F}$ and the class of all complements of graphs in $\mathcal{F}$ is polynomially χ-bounded.

Approximating the chromatic number and computing the clique number

We would like to make the Polynomial χ-boundedness Conjecture algorithmic; that is, we would like to efficiently approximate the chromatic number of graphs in a proper vertex-minor-closed class. It is natural to conjecture the following.

Chromatic Number Approximation Conjecture. *For any proper vertex-minor-closed class of graphs $\mathcal{F}$, there exists $\epsilon > 0$ so that the chromatic number of an n-vertex graph in $\mathcal{F}$ can be approximated to within a factor of $n^{1-\epsilon}$ in polynomial time.*

While this conjecture may not seem strong at first glance, it is typically very difficult to approximate the chromatic number; Zuckerman [125] showed that it is NP-hard to approximate the chromatic number of an n-vertex graph to within a factor of $n^{1-\epsilon}$ for any fixed $\epsilon > 0$. We also cannot expect to get an exact algorithm for proper vertex-minor-closed classes; it is NP-complete to compute the chromatic number of a circle graph [56].

In light of the Polynomial χ-boundedness Conjecture, there should be a rather simple approximation algorithm; return the size of a largest clique. Geelen conjectures that this can be done in polynomial time.

Max-Clique Conjecture (Geelen - see [115]). *For any proper vertex-minor-closed class of graphs $\mathcal{F}$, there is a polynomial time algorithm that computes the clique number of a graph in $\mathcal{F}$.*

The Polynomial χ-boundedness and Max-Clique Conjectures would together imply the Chromatic Number Approximation Conjecture.

The Max-Clique Conjecture is true when $\mathcal{F}$ is the class of circle graphs or their complements [57], and when $\mathcal{F}$ has bounded rank-width [27]. Geelen [59] believes that the conjecture also holds for the class of graphs obtained from circle graphs by complementing on a set, provided the set is given as input. (The proof would follow the methods of [57]).

For these conjectures it seems important to obtain an algorithmic version of the Structural Conjecture. For any fixed circle graph H, the Grid Theorem for Vertex-Minors yields a polynomial-time algorithm that finds either a vertex-minor isomorphic to H, or a rank-decomposition of bounded width as follows.

First, fix an integer r_H so that every graph with rank-width at least r_H has a vertex-minor isomorphic to H. Jeong, Kim, and Oum [73] provided an efficient algorithm that, for an input graph G, determines whether or not the rank-width of G is at most r_H, and, if it is, finds a rank-decomposition of width at most r_H. So we may assume that the rank-width is more than r_H. Since deleting a vertex decreases the rank-width by at most one, we can find an induced subgraph G' of G that has rank-width exactly r_H. Then, using a rank-decomposition for G' of width r_H, we can find a vertex-minor of G' that is isomorphic to H using dynamic programming [28]; for further details see the survey by Oum [102].

We believe that our proof of the Local Structure Theorem can also be made algorithmic, but this would require additional work which is not included in the thesis. On a related note, we should point out that local equivalence classes are not as complicated as they might seem; Bouchet [9] gave a polynomial time algorithm to determine if two graphs are locally equivalent.

First-order model-checking

We motivated vertex-minors by saying that they are like graph minors but for dense graph classes. We will eventually formalize this analogy in Section 1.6 by considering fundamental graphs; this will show that planar graphs are like circle graphs and branch-width is like rank-width. However, there is another nice way to formalize an analogy: by taking transductions. Sometimes the two approaches align; for example, a class of graphs has bounded rank-width if and only if it is contained in the image of a set of trees under a monadic second-order transduction (see [28, Proposition 5.4]). However this type of transduction is oftentimes too powerful; planar graphs can yield all graphs via a monadic second-order transduction. So we consider first-order transductions instead.

Informally, a *first-order sentence* is a logical statement that can quantify over vertices, express whether or not two vertices are adjacent or equal, and use logical connectives. A typical example is that, for any fixed graph H, there is a first-order sentence expressing that "there is an induced subgraph isomorphic to H". Likewise, for a fixed $t \in \mathbb{Z}^+$, there is a first-order sentence expressing that "there is a dominating set of size at most t". A sentence is either true or false for a given graph; if a sentence φ is true for a graph G, we say that G *models* φ. (In general, Ehrenfeucht–Fraïssé games give a nice combinatorial way to determine if two graphs model the same first-order sentences.)

We refer the reader to [53] for a short definition of first-order transductions. The definition is roughly as follows. First we non-deterministically specify p sets of vertices in a graph. Then we make m disjoint labelled copies. Finally we use first-order formulas (which are like sentences but with free-variables; so they may be true or false for a given vertex or pair of vertices) to specify which vertices and adjacencies to include. So a first-order transduction is specified by fixed integers p and m, and some first-order formulas. Because of the non-deterministic step, applying a transduction to a graph G yields a collection of graphs. For example, for each $p \in \mathbb{N}$, there is a first-order transduction that yields, when applied to a graph G, the collection of all graphs that can be obtained from G by complementing on p sets. Likewise, a first-order transduction can be used to obtain all induced subgraphs of a graph (but not necessarily all subgraphs because their edge-sets may not be "definable").

So we have another potential way to generalize theorems about proper minor-closed classes. In fact, this approach applies more generally to "sparse" graph classes in the sense of Nešetřil and Ossona de Mendez [96]. The new, more general graph classes obtained via first-order transductions are called "structurally sparse"; this approach was introduced in [54] and [97]. There are a number of algorithmic applications of sparsity (see the survey by Dvořák and Král' [42] and the book by Nešetřil and Ossona de Mendez [96]), but one of the most important is first-order model-checking.

We say that *first-order model-checking is fixed-parameter tractable* on a class of graphs $\mathcal{F}$ if there is an algorithm that takes as input an n-vertex graph $G \in \mathcal{F}$ and a first-order sentence φ, and determines if G models φ in time $f(|\varphi|)n^c$, for some fixed function f and constant c. (We write $|\varphi|$ for the length of the sentence φ, so this is the parameter.) We conjecture the following.

First-Order Model-Checking Conjecture. *For any proper vertex-minor-closed class of graphs $\mathcal{F}$ and any integer ω, first-order model-checking is fixed-parameter tractable on the class of all graphs in $\mathcal{F}$ that have clique number at most ω.*

In fact it seems possible that such classes have bounded twin-width [3]; if the appropriate linear order could be found in polynomial time, then this would imply the First-Order Model-Checking Conjecture. We believe that this stronger conjecture holds for circle graphs, graphs obtained from a circle graph by complementing on a set (when the set is provided as part of the input), and graphs of bounded rank-width. The key fact is that twin-width is preserved under first-order transductions, even when the linear order is part of the structure [3]. So we can "obtain a circle graph from its chord diagram", which can be found in polynomial time [5, 52].

This conjecture about twin-width would imply Davies' theorem [31] that every proper vertex-minor-closed class is χ-bounded, because classes of bounded twin-width are χ-bounded [2]. In fact, Gajarský, Pilipczuk, and Toruńczyk [55] very recently showed that any monadically stable class of bounded twin-width has a linear χ-bounding function. (Informally, a class is monadically stable if there is no first-order transduction that yields all half graphs; half graphs will come up again in Section 2.4.) Circle graphs, however, do not have a linear χ-bounding function as shown by Kostochka [80, 81]. We believe that their twin-width must somehow depend on their clique number; there is no obvious relationship with the Polynomial χ-boundedness Conjecture.

The same difficulty occurs in the First-Order Model-Checking Conjecture. While it seems likely that ω could be made a parameter, the dependence on ω cannot be removed; under the Exponential Time Hypothesis, first-order model-checking is not fixed-parameter tractable for the class of circle graphs [72]. In that paper, Hliněný, Pokrývka, and Roy proposed the First-Order Model-Checking Conjecture for the class of circle graphs; this was our original motivation for looking into the case of vertex-minors.

The First-Order Model-Checking Conjecture should be seen as one concrete conjecture in part of a larger project to relate vertex-minors to structural sparsity (and beyond). It is tempting to think there should be a connection because graph minors are very closely related to sparsity (see [96]).

Simulating measurement-based quantum computation

Measurement-based quantum computation (MBQC) is an alternate model of quantum computation that was introduced by Raussendorf and Briegel [106]. MBQC is "polynomially equivalent" to the standard quantum gate model; see the survey by Jozsa [75]. The advantage of MBQC is that the "level of entanglement" can only decrease throughout the course of computation; so it is useful for studying what resources can potentially give quantum computing a speed-up over classical computing.

Here is the idea of how MBQC works. First a quantum state is prepared, and then measurements are performed on individual qubits. Depending on the outcomes of prior measurements, it is decided what qubit to measure next, and in what basis (this classical step can use randomness). The output of the computation is then determined based on all of the measurement outcomes. When talking about running time, we require a uniformity condition on how the quantum state is prepared. That is, there is a separate classical algorithm which says, in polynomial time, which state to prepare. It is convenient to use n-qubit "graph states" (see [71]) because they are fully represented by an n-vertex graph. (The original model of Raussendorf and Briegel [106] uses graph states where the graphs are grids.)

So, in order to study what resources give quantum computing its power, we can restrict ourselves to preparing only graph states from a graph class $\mathcal{F}$. Informally, we write $\mathrm{BQP}_{\mathcal{F}}$ for the decision problems that can be solved in polynomial time (with bounded error) using measurement-based quantum computation where only graph states from $\mathcal{F}$ can be prepared. We said earlier that MBQC is "polynomially equivalent" to the standard quantum gate model; so in particular, if $\mathcal{G}$ is the class of all graphs, then $\mathrm{BQP}_{\mathcal{G}} = \mathrm{BQP}$ (bounded-error quantum polynomial time). On the other hand $\mathrm{BQP}_{\emptyset} = \mathrm{BPP}$ (bounded-error probabilistic polynomial time).

Geelen [59] conjectures that if we restrict ourselves to preparing graph states from any proper vertex-minor-closed class, then measurement-based quantum computation can be efficiently simulated classically. There are several ways to interpret this, but here is one.

Simulation Conjecture (Geelen [59])**.** *For any proper vertex-minor-closed class of graphs* $\mathcal{F}$, $BQP_{\mathcal{F}} = BPP$.

The conjecture is interesting because locally equivalent graphs yield graph states with the same "level of entanglement" (see [120]). We will see later (in Lemma 1.6.5) that there are really three main operations for vertex-minors; these three operations correspond to the three Pauli matrices (see [30] for a nice discussion).

There are two main pieces of evidence for the Simulation Conjecture. First, it is known that MBQC can be efficiently simulated when the prepared graph states correspond to graphs of bounded rank-width [121]. Furthermore, the same holds for "planar code states" under a minor connectivity-type condition on the measured qubits [12]; these states relate to the graph states of circle graphs.

1.5 Structural conjectures

This section is dedicated to formalizing the Structural Conjecture.

The weak structural conjecture

First we give a version of the conjecture that lets us avoid "decomposing" for "highly rank-connected" graphs. This notion of connectivity is rather restrictive, yet there exist circle graphs of arbitrarily large rank-connectivity (see Lemma 1.5.1). So, unlike the case of graph minors and topological minors (see Mader [89]), high "connectivity" does not force a vertex-minor. That is, it is not true that for each graph H, every graph of sufficiently large rank-connectivity has a graph isomorphic to H as a vertex-minor. However, Geelen [59] conjectures that the sufficiently rank-connected graphs in any proper vertex-minor-closed class have a simple structure: that they are low-rank perturbations of circle graphs.

For $k \in \mathbb{N}$, a graph G is *k-rank-connected* if it has at least $2k$ vertices and $\rho(X) \geq \min(|X|, |V(G) - X|, k)$ for each $X \subseteq V(G)$. Equivalently, there exists a set of cut-rank k, and every set X of cut-rank less than k has $\rho(X) = \min(|X|, |V(G) - X|)$ (since the right-hand side also upper bounds the cut-rank). Then the *rank-connectivity* of G is the minimum $k \in \mathbb{N}$ so that G is k-rank-connected. This definition is similar to "Tutte connectivity" for matroids [119]. In fact there is a direct connection for fundamental graphs; see Section 1.6 and [8, 99].

We have the following simpler version of the Structural Conjecture for highly rank-connected graphs.

Weak Structural Conjecture (Geelen [59])**.** *For any proper vertex-minor-closed class of graphs $\mathcal{F}$, there exist $k, p \in \mathbb{N}$ so that each k-rank-connected graph in $\mathcal{F}$ is a rank-p perturbation of a circle graph.*

This conjecture is very similar to the conjecture of Geelen, Gerards, and Whittle [62] about the structure of "highly vertically connected" binary matroids (vertical connectivity is less restrictive than Tutte connectivity). Geelen, Gerards, and Whittle [62] also conjecture that every binary matroid of sufficiently large vertical connectivity has either the graphic matroid of the t-vertex clique, or its dual, as a minor. This would generalize the result of Mader [89] that every graph of sufficiently large vertex-connectivity has the t-vertex clique as a minor. Furthermore, due to the connection with pivot-minors (see Section 1.6), this conjecture would imply that every bipartite graph of sufficiently large rank-connectivity has every t-vertex graph as a vertex-minor (up to isomorphism).

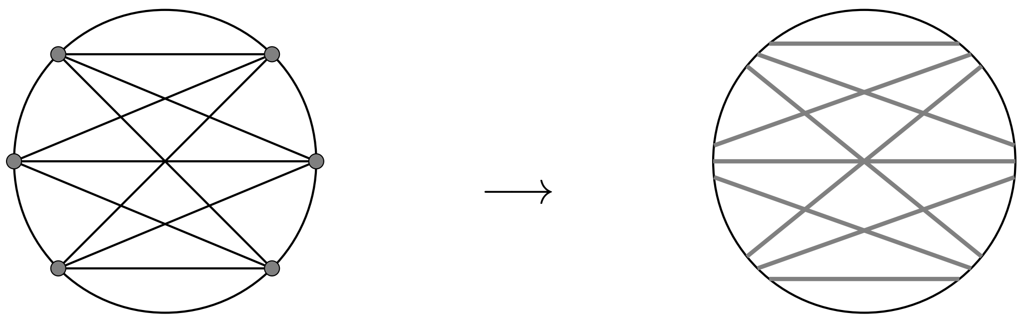

Figure 1.7: The graph G with its vertices placed according to ϕ on the unit circle (left) and the corresponding chord diagram (right).

So it is rather surprising that, as we will now show, there exists a proper vertex-minor-closed class with graphs of arbitrarily large rank-connectivity. In some informal sense, these graphs must be very far from bipartite.

Lemma 1.5.1. *For each $k \in \mathbb{Z}^+$, there exists a circle graph that is k-rank-connected.*

Proof. Fix $k \in \mathbb{Z}^+$, and set $d \coloneqq (4k)(2k)^{k-1}$. Let G be a graph of girth at least $4k$ and connectivity at least d; such a graph exists by, for instance, the classic random construction of Erdös [50]. Next fix an arbitrary injective function ϕ from $V(G)$ to the unit circle; this just fixes a cyclic order of $V(G)$. Then form a chord diagram by replacing each edge uv of G with a chord whose ends are very close to $\phi(u)$ and $\phi(v)$; arrange these chords so that for each vertex, the "incident" chords are non-intersecting (see Figure 1.7). We abuse notation by referring to these chords and the edges of G interchangeably.

We are not yet done adding chords. We will define a chord diagram $\mathcal{C}$ that is disjoint from $E(G)$, and then consider the circle graph of $\mathcal{C} \cup E(G)$. This chord diagram $\mathcal{C}$ is obtained by, for each vertex $v \in V(G)$, adding some non-intersecting chords "very close" to $\phi(v)$ as depicted in Figure 1.8. So, let $v \in V(G)$, let $D \geq d$ be the degree of v, and let $v_1, \ldots, v_D$ be the ends, in order and near $\phi(v)$, of the D chords incident to v. First add a chord that goes over all of $v_1, \ldots, v_D$. We say that this chord has *level 1*; there will be k levels of chords. Suppose that we have already defined the chords at levels $1, \ldots, i$ for some $i < k$. Then for each chord J of level i, there are $2k$ new chords of level $i+1$ "underneath" J; place the new chords so that they partition whichever of $v_1, \ldots, v_D$ are underneath J as evenly as possible. Thus, inductively, each chord of level i has at least $(4k)(2k)^{k-i}$ of $v_1, \ldots, v_D$ underneath it.

This completes the definition of $\mathcal{C}$. We say that two chords in $\mathcal{C}$ *nest* if they are at the same vertex and one is "underneath" the other. We will show that the circle graph of

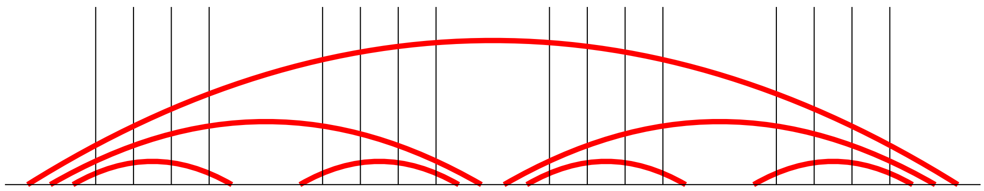

Figure 1.8: A depiction of the chords in $\mathcal{C}$ (bold red; they are drawn curved for convenience) and $E(G)$ (black) when we have "zoomed in" to be very close to $\phi(v)$. The exact number of chords does not align with the proof.

$\mathcal{C} \cup E(G)$ is k-rank-connected. So, let (X, Y) be a partition of $\mathcal{C} \cup E(G)$; we may assume that $|Y \cap E(G)| \geq |X \cap E(G)|$. We split into cases.

Case 1: $|X \cap E(G)| \geq 4k$

Since G is $4k$-connected and has girth at least $4k$, there are at least $4k$ vertices that are incident to both an edge in $X \cap E(G)$ and an edge in $Y \cap E(G)$. At least $2k$ of these vertices have level 1 chords in the same side of (X, Y). Then since each chord in $E(G)$ intersects precisely two chords of level 1, there is a perfect matching of size k induced between X and Y. So $\rho(X) \geq k$, which completes the case.

Case 2: $|X \cap E(G)| < 4k$ and $X \cap \mathcal{C}$ has $2k$ pairwise non-nested chords

Since $|X \cap E(G)| < 4k$, each chord in $X \cap \mathcal{C}$ intersects a chord in $Y \cap E(G)$. Furthermore, each chord intersects at most two of the $2k$ pairwise non-nested chords in $X \cap \mathcal{C}$. So there is a matching of size k induced between between X and Y, and $\rho(X) \geq k$, which completes the case.

Case 3: $|X \cap E(G)| < 4k$ and $X \cap \mathcal{C}$ has k pairwise nested chords

These k pairwise nested chords are at levels $1, 2, \ldots, k$ at some vertex v of G; denote them by $J_1, J_2, \ldots, J_k$, respectively. We claim that for each level $i < k$, there exists a chord in $Y \cap E(G)$ that intersects J_i but none of $J_{i+1}, \ldots, J_k$. This would imply that $\rho(X) \geq k$ (since any square matrix with 1's on the diagonal and 0's on the upper-right has full rank) and we would be done. The claim holds since $|X \cap E(G)| < 4k$ and, for each level $i < k$, there exists a chord other than J_{i+1} at level $i + 1$ underneath J_i.

Case 4: None of the above cases occurs.

In this case we will show that $\rho(X) \geq |X|$. It suffices to find an ordering $J_1, \ldots, J_{|X|}$ of the chords in X so that for each i, there is a chord in Y that intersects J_i but none of

$J_{i+1}, \ldots, J_{|X|}$. The chords in $X \cap E(G)$ will come first. Since G has girth at least $4k$, the set $X \cap E(G)$ is acyclic. We successively delete an edge in $X \cap E(G)$ that is incident to a leaf vertex v. There is a chord in $Y \cap \mathcal{C}$ at v (because $X \cap \mathcal{C}$ does not have k pairwise nested chords since *Case 3* does not occur); this chord does not intersect any of the "following" chords in X.

It just remains to consider the chords in $X \cap \mathcal{C}$. Among the chords in $X \cap \mathcal{C}$ that have not yet been ordered, choose a chord J whose level is minimum. Say that J is at a vertex $v \in V(G)$ and has level i. Since $X \cap \mathcal{C}$ does not contain $2k$ pairwise non-nested chords (as *Case 2* does not occur), there exists a chord underneath J at level $i + 1$ that is in Y. By again looking underneath that chord, and so on, and then applying the fact that $|X \cap E(G)| < 4k$ (since *Case 1* does not occur), we can find a chord in $Y \cap E(G)$ that intersects J but no other chord in $X \cap \mathcal{C}$. By adding the chord J next, and continuing in this fashion among all remaining chords in $X \cap \mathcal{C}$, we obtain the desired ordering $J_1, \ldots, J_{|X|}$. This completes the final case, and therefore also Lemma 1.5.1. □

The full structural conjecture

We are almost ready to formalize the Structural Conjecture. We just need to discuss "decompositions". We call the following structure a "tree-decomposition" even though it is different from the classic "tree-decompositions" that are used for tree-width.

So, a *tree-decomposition* of a graph G is a pair $\mathcal{T} = (T, \phi)$ where T is a tree and $\phi : V(G) \to V(T)$ is a function. So the vertices of T yield a partition of $V(G)$; for each $t \in V(T)$, there is a (possibly empty) part $\phi^{-1}(t)$ called the *bag* of t. The tree also shows how to decompose at each bag; for each $t \in V(T)$, there is a partition $\mathcal{A}_{\mathcal{T}}(t)$ of $V(G) - (\phi^{-1}(t))$ according to the connected components of $T - t$. More formally, for each $t \in V(T)$, we write $\mathcal{A}_{\mathcal{T}}(t)$ for the collection of all sets $\phi^{-1}(X)$ where X is the vertex set of a component of $T - t$.

The formal conjecture will say that there exists a tree-decomposition so that each bag "becomes" a circle graph after first performing a low-rank perturbation and then "decomposing on each set in $\mathcal{A}_{\mathcal{T}}(t)$". Unfortunately we do not have a direct analog of clique sums for vertex-minors. The Structural Conjecture is better compared with a version of the Graph Minors Structure Theorem where, informally, as many of the "clique sum vertices" as possible are turned into apex vertices (such statements are given in [35] and [37], for example). Roughly, for graph minors, we are just left with gluing onto apex vertices and facial triangles. So for vertex-minors, we first "lower the cut-rank as much as possible"

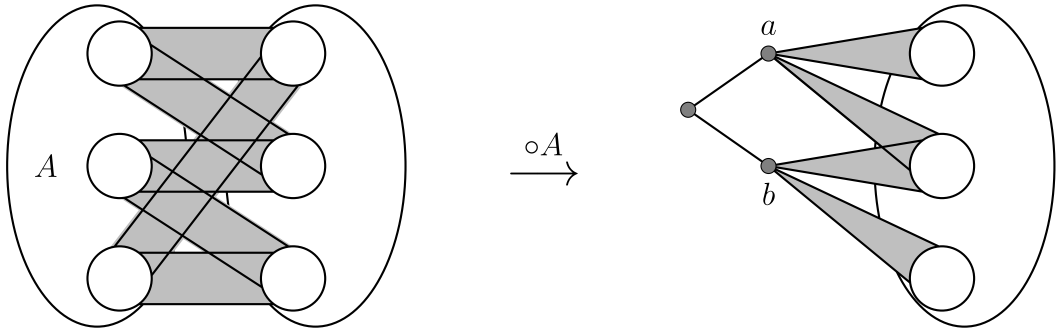

Figure 1.9: Decomposing on a set A with $\rho(A) = 2$.

using a low-rank perturbation. Then, afterwards, each set in $\mathcal{A}_{\mathcal{T}}(t)$ has cut-rank at most 2 (this is the analog of a 3-separation), and we can "decompose on each set in $\mathcal{A}_{\mathcal{T}}(t)$".

So, let G be a graph with $A \subseteq V(G)$ so that $\rho(A) \leq 2$. If $\rho(A) = 0$, then define $G \circ A$ to be the graph obtained from G by deleting A. If $\rho(A) = 1$, then define $G \circ A$ to be the graph obtained from G by deleting all vertices in A other than one vertex that has a neighbour in $V(G) - A$. Finally, if $\rho(A) = 2$, then there are vertices $a, b \in A$ with distinct neighbourhoods in $V(G) - A$ (see Figure 1.9); define $G \circ A$ to be the graph obtained from G by deleting all vertices in A other than a and b, and then adding a new vertex that is adjacent to just a and b. While there may be more than one choice for a and b, the graph $G \circ A$ is well-defined up to isomorphism and local equivalence; we omit the proof of this fact. We call $G \circ A$ the graph obtained from G by *decomposing on A*.

This decomposition operation can be successively performed on disjoint sets of vertices, as long as they each have cut-rank at most 2. So, if G is a graph and $\mathcal{A}$ is a collection of disjoint vertex sets each of cut-rank at most 2, then define $G \circ \mathcal{A}$ to be the graph $G \circ A_1 \circ \ldots \circ A_k$, where $A_1, \ldots, A_k$ is any enumeration of the sets in $\mathcal{A}$. Again this graph is well-defined up to isomorphism and local equivalence, though we omit the proof. We call $G \circ \mathcal{A}$ the graph obtained from G by *decomposing on* $\mathcal{A}$.

We are now ready to state the main conjecture.

Structural Conjecture (Geelen [59]). *For any proper vertex-minor-closed class of graphs $\mathcal{F}$, there exists $p \in \mathbb{N}$ so that each graph $G \in \mathcal{F}$ has a tree-decomposition $\mathcal{T} = (T, \phi)$ so that for each $t \in V(T)$, there is a rank-p perturbation $\widetilde{G}_t$ of G so that each set in $\mathcal{A}_{\mathcal{T}}(t)$ has cut-rank at most 2 in $\widetilde{G}_t$, and $\widetilde{G}_t \circ \mathcal{A}_{\mathcal{T}}(t)$ is a circle graph.*

The conjecture holds for classes of bounded-rank perturbations of circle graphs; take a tree-decomposition (T, ϕ) where T has one vertex. The conjecture also holds for classes

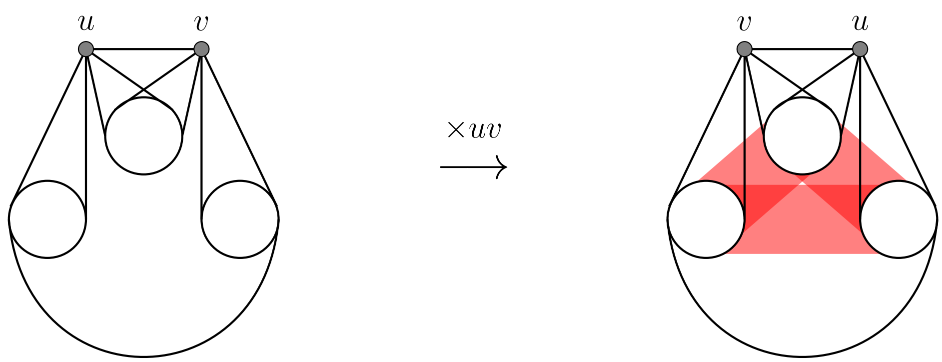

Figure 1.10: A graph before and after pivoting on uv, with the exchanged edges/non-edges in red.

of bounded rank-width; if T is a rank-decomposition of G so that the maximum width of an edge of T is minimum, then, where for a vertex $v \in V(G)$ we define $\phi(v) = v$, the tree-decomposition (T, ϕ) of G suffices. So these two examples are on opposite ends of the spectrum; that is why we view them as the most important examples of proper vertex-minor-closed classes.

The precise statement of the Structural Conjecture was motivated by the matroid minors project of Geelen, Gerards, and Whittle (see [58] for an overview and [62] for the case of "high connectivity"). We will discuss this motivation further in the next section.

1.6 Pivot-minors

Several times so far we have alluded to a common generalization of vertex-minors and graph minors. That common generalization is provided by pivot-minors, as shown by Bouchet [7]. Not only is this connection useful for motivating conjectures, but it also provides one of the most important proof techniques for vertex-minors.

For a graph G and an edge $uv \in E(G)$, *pivoting on* uv results in the graph $G{*}u{*}v{*}u$; this operation is well-defined because $G{*}u{*}v{*}u = G{*}v{*}u{*}v$ (see [99, Corollary 2.2]). We denote this new graph by $G \times uv$. Equivalently, $G \times uv$ is obtained from G by "complementing between" the sets of all vertices other than u and v which are 1) adjacent to u but not v, 2) adjacent to v but not u, and 3) adjacent to both u and v, and then switching the labels of u and v. See Figure 1.10 and [99, Proposition 2.1]. *Complementing between* pairwise disjoint sets of vertices means to exchange edges/non-edges between the sets.

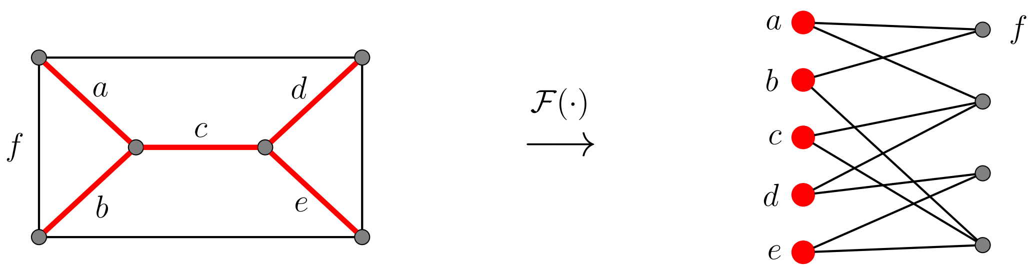

Figure 1.11: A graph (left) with a spanning tree T (in bold red) and its fundamental graph (right).

Now we can define *pivot-minors* and *pivot equivalence* in the same way that we defined vertex-minors and local equivalence. (So a *pivot-minor* of a graph G is any graph that can be obtained from G by pivoting on edges and deleting vertices. A graph is *pivot-equivalent* to G if it can be obtained from G by pivoting alone; this is an equivalence relation since $(G \times uv) \times uv = G$.) Thus every pivot-minor of a graph G is also a vertex-minor of G, but typically not vice-versa.

Remarkably, pivot-minors of bipartite graphs essentially generalize graph minors. (The class of bipartite graphs is closed under pivot-minors; in fact, a graph is bipartite if and only if it does not have a triangle as a pivot-minor.) We will explain this connection next.

Pivot-minors and graph minors

Let G be a graph; we will show how to obtain a pivot-equivalence class of bipartite graphs from G. First, a *fundamental graph* of G is obtained by selecting a maximal spanning forest T of G and creating a bipartite graph where one side is $E(T)$ and the other side is $E(G) - E(T)$; each $e \in E(G) - E(T)$ is adjacent to the edges in the unique cycle of $T + e$ (see Figure 1.11). Let $\mathcal{F}(G)$ denote the set of all fundamental graphs of G.

Bouchet [7] proved that $\mathcal{F}(G)$ is a pivot-equivalence class of bipartite graphs, and that each pivot-minor of a graph in $\mathcal{F}(G)$ is a fundamental graph of a minor of G (also see the nice explanation by Oum [99]). In the fundamental graph corresponding to a tree T of G, pivoting exchanges edges in T, deleting $e \in E(G) - E(T)$ deletes it from G, and deleting $e \in E(T)$ contracts it in G.

To go in the opposite direction (beginning with a pivot-equivalence class of bipartite graphs), we would need to work with binary matroids. We will discuss this fact in the next subsection, but for now we focus on graphs. So to state the following theorem in terms of graphs, we use Whitney's planarity criterion [122] and Whitney's theorem [123] about representations of graphic matroids.

Theorem 1.6.1 (Bouchet [11] and Oum [99])**.** *For any 3-connected graphs H and G, a graph in $\mathcal{F}(H)$ is isomorphic to a pivot-minor of a graph in $\mathcal{F}(G)$ if and only if H is isomorphic to a minor of G or, if G is planar, its planar dual.*

It is much more technically complicated to state a version of Theorem 1.6.1 for graphs that are not 3-connected, but it can be done.

Moreover, fundamental graphs let us formalize the relationship between planar graphs and circle graphs.

Theorem 1.6.2 (de Fraysseix [34])**.** *A bipartite graph is a circle graph if and only if it is a fundamental graph of a planar graph.*

This theorem characterizes circle graphs; every circle graph is a vertex-minor of one that is bipartite (folklore, see [16]).

We can formalize a relationship between branch-width and rank-width as well. Using a theorem of Mazoit and Thomassé [93] about the branch-width of a graphic matroid, we have the following.

Theorem 1.6.3 (Oum [99])**.** *The branch-width of a bridgeless graph is one more than the rank-width of each of its fundamental graphs.*

For graphs with bridges, the rank-width can be at most one off.

As was the case for circle graphs, every class of bounded rank-width comes from a bipartite class of bounded rank-width. Formally, it follows from a theorem of Kwon and Oum [84] that there exists a function f so that for each $r \in \mathbb{N}$, every graph of rank-width r is a vertex-minor of a bipartite graph of rank-width at most $f(r)$. We believe that this sort of statement holds for any proper vertex-minor-closed class.

Bipartite Generation Conjecture. *For any proper vertex-minor-closed class of graphs $\mathcal{F}$, there exists a proper vertex-minor-closed class of graphs $\mathcal{F}'$ so that each graph in $\mathcal{F}$ is a vertex-minor of a bipartite graph in $\mathcal{F}'$.*

Next we will see that this conjecture would relate vertex-minors to binary matroid minors.

Pivot-minors and binary matroid minors

We refer the reader to the book by Oxley [105] for an introduction to matroids. In this section we formalize the connection between binary matroid minors and pivot-minors.

Let M be a binary matroid. We define $\mathcal{F}(M)$ and the fundamental graphs of M in the same way as before. (So a *fundamental graph* of M is obtained by selecting a basis B and creating a bipartite graph where one side is B and the other side is $E(M) - B$; each element $e \in E(M) - B$ is adjacent to the other elements in the unique circuit of $B \cup \{e\}$. We let $\mathcal{F}(M)$ denote the set of all fundamental graphs of M.) Bouchet [7] really proved that $\mathcal{F}(M)$ is a pivot-equivalence class of bipartite graphs, and that each pivot-minor of a graph in $\mathcal{F}(M)$ is a fundamental graph of a minor of M.

Now we can go in the opposite direction as well; every pivot-equivalence class of bipartite graphs is the set of fundamental graphs of a binary matroid M. However, one side of a fundamental graph corresponds to a basis B, the other side corresponds to $E(M)-B$, and exchanging sides takes the dual of the matroid. More formally, a *component-wise dual* of a matroid M is a matroid that can be obtained from M by replacing zero or more of its components by their duals. We have the following theorem.

Theorem 1.6.4 (Bouchet [11] and Oum [99])**.** *For any binary matroids M and N, a graph in $\mathcal{F}(N)$ is isomorphic to a pivot-minor of a graph in $\mathcal{F}(M)$ if and only if N is isomorphic to a minor of a component-wise dual of M.*

This theorem can be used to make Theorem 1.6.1 precise for graphs that are not 3-connected. It is also why well-quasi-ordering for pivot-minors would imply well-quasi-ordering for graph minors (as mentioned by Oum [102]).

Now we can see why the Bipartite Generation Conjecture is interesting; it would let us obtain "structure for vertex-minors" from "structure for binary matroids". However, the only way we can imagine proving the conjecture in the first place is to use "structure for vertex-minors". (Brijder and Traldi [15, Corollary 35] gave an interesting way of constructing a pivot-equivalence class of bipartite graphs from a local equivalence class, but their construction does not yield proper vertex-minor-closed classes.) Still, this is why the (ongoing) matroid minors project of Geelen, Gerards, and Whittle [58] motivated the precise statement of the Structural Conjecture.

Pivot-minors and vertex-minors

The following definition is very useful for proving things about vertex-minors. For a graph G and a vertex $v \in V(G)$, define G/v to be the graph $G \times uv - v$ for some neighbour u

of v, or, if v has no neighbour, to be the graph $G - v$. We call G/v the graph obtained from G by *pivot-deleting* v; it is well-defined up to pivot-equivalence (and therefore also local equivalence). This is because, for any two neighbours u_1 and u_2 of v, we have $G \times vu_1 = (G \times vu_2) \times u_1u_2$; see [99, Proposition 2.5].

Surprisingly, there are only three ways to remove a vertex from a graph, up to local equivalence (see [66, Lemma 3.2] for a direct proof).

Lemma 1.6.5 (Bouchet [7])**.** *If H is a vertex-minor of a graph G and $v \in V(G) - V(H)$, then H is a vertex-minor of either $G - v$, $G * v - v$, or G/v.*

This fact gives us a very useful proof technique. Suppose that we are interested in two properties P_1 and P_2, and let G be a vertex-minor-minimal graph which satisfies P_1 and P_2. Then there is no vertex v so that for $i = 1, 2$, most of the three graphs $G - v$, $G * v - v$, and G/v satisfy P_i. Thus we can oftentimes deal with P_1 and P_2 separately. This approach is used in the proofs of both of our main theorems, which is a major obstacle to generalizing these theorems to pivot-minors (where there are only two ways to remove a vertex).

There is one particularly important example. For a graph G with disjoint sets of vertices S and T, the *connectivity between S and T*, denoted $\kappa_G(S, T)$ (or just $\kappa(S, T)$ when the graph is clear from context), is the minimum cut-rank of a set of vertices which contains S and is disjoint from T. Removing a vertex cannot increase the connectivity, and $\kappa(S, T) = \kappa(T, S)$. Moreover, at least two of the three ways to remove a vertex maintain the connectivity.

Theorem 1.6.6 (Oum [99, Lemma 4,4])**.** *If S and T are disjoint sets of vertices in a graph G, then for each $v \in V(G) - (S \cup T)$, at least two of $\kappa_{G-v}(S, T)$, $\kappa_{G*v-v}(S, T)$, and $\kappa_{G/v}(S, T)$ are equal to $\kappa_G(S, T)$.*

(Oum gave a slightly different statement of the result, but Theorem 1.6.6 follows; see [64, Theorem 4.1].) Oum's theorem directly implies Tutte's linking theorem [118] for binary matroids due to Theorem 1.6.4 on fundamental graphs of binary matroids. This connection makes Oum's Theorem an analog of Menger's theorem for vertex-minors.

Finally, we complete this chapter with another application of Lemma 1.6.5 (that there are three ways to remove a vertex). We use the lemma to prove that low-rank perturbations behave well with respect to taking vertex-minors.

Lemma 1.6.7. *For any $p \in \mathbb{N}$ and any proper vertex-minor-closed class $\mathcal{F}$, there exists a graph that is not isomorphic to a vertex-minor of any rank-p perturbation of a graph in $\mathcal{F}$.*

Proof. Let $\mathcal{F}'$ be the class of all vertex-minors of a rank-p perturbation of a graph in $\mathcal{F}$; we want to show that $\mathcal{F}'$ is a proper class. First we prove a claim.

Claim 1.6.7.1. *Each graph in $\mathcal{F}'$ can be obtained from a graph in $\mathcal{F}$ by complementing on $3p$ sets.*

Proof. Let $H' \in \mathcal{F}'$; so there is a graph $G' \in \mathcal{F}'$ that has H' as a vertex-minor and is a rank-p perturbation of a graph $G \in \mathcal{F}$. We claim that there is a graph $\hat{G}$ which has both G' and G as vertex-minors and has only p more vertices than G' and G. By Lemma 1.3.3, any symmetric rank-p matrix can be written as the sum of p_1 rank-1 matrices and p_2 rank-2 matrices, also symmetric, so that $p = p_1 + 2p_2$. Then, to obtain G' from G, we first add a stable set of size p_1 and a perfect matching of size p_2, then we locally complement at each vertex in the stable set and pivot on each edge in the matching. So such a graph $\hat{G}$ exists.

Now, by locally complementing in $\hat{G}$, we may assume that H' is an induced subgraph of $\hat{G}$. Since G is a vertex-minor of $\hat{G}$, by Lemma 1.6.5 (that there are three ways to remove a vertex), we know that for each $v \in V(\hat{G}) - V(G)$, the graph G is also a vertex-minor of either $G - v$, $G * v - v$, or G/v. Each of these three graphs can be obtained from $G - v$ by complementing on 3 sets (some of which may be empty). So, by repeating this argument for each of the p vertices in $V(\hat{G}) - V(G)$, we see that H' can be obtained from a vertex-minor of G by complementing on $3p$ sets. The claim follows. $\square$

Next we work with a particular type of graph that, informally, behaves like a complete graph does for graph minors. For each $t \in \mathbb{Z}^+$, let $K_t^{(1)}$ denote the graph that is obtained from a t-vertex clique by subdividing each edge once; subdividing an edge uv replaces it by a new vertex which is only adjacent to u and v. This graph $K_t^{(1)}$ has every t-vertex graph as a vertex-minor, up to isomorphism; we can choose to include an edge uv of the t-vertex clique by locally complementing at the vertex adjacent to u and v. We will show that $\mathcal{F}'$ forbids one of these graphs $K_t^{(1)}$.

First, let $r \in \mathbb{Z}^+$ be such that there exists an r-vertex graph that is not isomorphic to any graph in $\mathcal{F}$. Then fix $t \in \mathbb{Z}^+$ such that every t-vertex clique whose edges are coloured with 2^{3p} colours contains a monochromatic $((r+5)2^{3p})$-vertex clique; such an integer exists by the multicolour version of Ramsey's theorem. We have one more key claim.

Claim 1.6.7.2. *Any graph which is obtained from $K_t^{(1)}$ by complementing on $3p$ sets has every r-vertex graph as a vertex-minor, up to isomorphism.*

Proof. First we apply some Ramsey theory. The vertices of $K_t^{(1)}$ correspond to the vertices and edges of a t-vertex clique. So consider a t-vertex clique whose edges and vertices are coloured by binary vectors of length $3p$ indicating the $3p$ sets. By the multicolour version of Ramsey's theorem, this t-vertex clique contains an $((r+5)2^{3p})$-vertex clique whose edges all have the same colour. By the pigeonhole principle, this resulting clique contains an $(r+5)$-vertex clique whose vertices all have the same colour.

In the resulting induced subgraph $K_{r+5}^{(1)}$ of $K_t^{(1)}$, each of the $3p$ sets is either the entire vertex-set, or is one of the two sides of the bipartition. Let (X, Y) denote the bipartition of $K_{r+5}^{(1)}$, where Y is the set of vertices which correspond to subdivided edges. Let G_{r+5} denote the graph that is obtained from $K_{r+5}^{(1)}$ by complementing on $3p$ sets. Then G_{r+5} is obtained from $K_{r+5}^{(1)}$ by possibly complementing on X, on Y, and/or between X and Y.

We will use two of the five "extra" vertices to "undo" a complementation on Y. So, if we complemented on Y to obtain G_{r+5}, then locally complement on a vertex in Y in G_{r+5}. This vertex is either non-adjacent to all but two vertices in X, or adjacent to all but two vertices in X. In any case, there exists a vertex-minor G_{r+3} of G_{r+5} so that G_{r+3} is obtained from $K_{r+3}^{(1)}$ by possibly complementing on the remaining vertices in X, and/or between the two sides.

Now we will use the remaining three "extra" vertices to "undo" a complementation between the two sides. If we complemented between the sides to obtain G_{r+5}, then pivot on an edge going between the sides. The end in Y is adjacent to all but two of the remaining vertices in X and non-adjacent to the rest of Y. The end in X is adjacent to all of the "relevant" vertices in Y. In any case, we find a vertex-minor G_r of G_{r+3} so that G_r is obtained from $K_r^{(1)}$ by possibly complementing on the remaining vertices in X.

This graph G_r has every r-vertex graph as a vertex-minor, up to isomorphism; we can choose to remove an edge uv of the r-vertex clique by locally complementing at the vertex adjacent to u and v. This finishes the proof of Claim 1.6.7.2. □

Now, if the graph $K_t^{(1)}$ was in $\mathcal{F}'$, then by Claim 1.6.7.1 there would be a graph in $\mathcal{F}$ that is obtained from $K_t^{(1)}$ by complementing on $3p$ sets. Then by Claim 1.6.7.2, the class $\mathcal{F}$ would contain every r-vertex graph as a vertex-minor, up to isomorphism. This is a contradiction, which means that $\mathcal{F}'$ is a proper class and Lemma 1.6.7 holds. □

Chapter 2

The Grid Theorem for Vertex-Minors

2.1 Introduction

In this chapter we prove the Grid Theorem for Vertex-Minors, which is restated below for convenience.

The Grid Theorem for Vertex-Minors (Geelen, Kwon, McCarty, and Wollan [64]). *For any circle graph H, there exists an integer r_H so that every graph with rank-width at least r_H has a vertex-minor isomorphic to H.*

As in the Grid Theorem of Robertson and Seymour [108], we will first prove the theorem for particular graphs called "comparability grids". For a positive integer n, the $n \times n$ *comparability grid* is the graph with vertex set $\{(i,j) : i,j \in \{1,2,\dots,n\}\}$ where there is an edge between any two distinct comparable vertices; vertices (i,j) and (i',j') are *comparable* if either $i \leq i'$ and $j \leq j'$, or $i \geq i'$ and $j \geq j'$. We will prove that every circle graph is isomorphic to a vertex-minor of a comparability grid in Lemma 2.2.3. Thus the Grid Theorem for Vertex-Minors is equivalent to the following result.

Theorem 2.1.1. *There is a function $f : \mathbb{Z} \to \mathbb{Z}$ so that for any positive integer n, every graph of rank-width at least $f(n)$ has a vertex-minor isomorphic to the $n \times n$ comparability grid.*

Despite the resemblance, we see no way of directly proving the Grid Theorem of Robertson and Seymour from the Grid Theorem for Vertex-Minors or vice-versa. However, the following conjecture of Oum [101], if true, would imply both results.

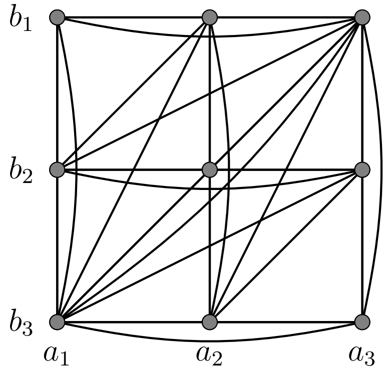

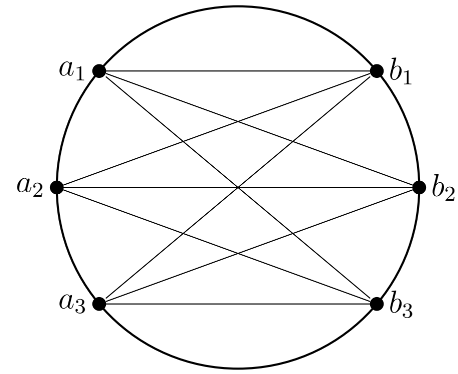

Figure 2.1: The $n \times n$ comparability grid with $n = 3$, where each vertex (i, j) is placed at position (a_i, b_{n+1-j}) (left), and the corresponding chord diagram (right).

Conjecture 2.1.2 (Oum [101])**.** *For any bipartite circle graph H, there exists an integer r'_H so that every graph with rank-width at least r'_H has a pivot-minor isomorphic to H.*

See Section 1.6 for an overview of the connection. The conjecture is known to hold for bipartite graphs, as that special case is equivalent to the grid theorem for binary matroids (see [60]). Oum [101] also proved Conjecture 2.1.2 for line graphs, and observed that, for circle graphs, it follows from Johnson's thesis [74]. It is natural to ask if something similar could hold for induced subgraphs, but this seems unlikely; see [13] and [25].

2.2 Circle graphs

In this section we prove that each circle graph is isomorphic to a vertex-minor of a comparability grid. To prove this result, we show that 1) every circle graph is a vertex-minor of a "permutation graph" and 2) every permutation graph is an induced subgraph of a comparability grid. For a permutation π of $\{1, \dots, n\}$, the *permutation graph* represented by π is the graph F_π with vertex set $\{1, \dots, n\}$ where vertices i and j, with $i < j$, are adjacent if $\pi_i > \pi_j$.

Lemma 2.2.1. *Every n-vertex permutation graph is isomorphic to an induced subgraph of the $n \times n$ comparability grid.*

Proof. To obtain a chord diagram for the $n \times n$ comparability grid, place distinct points $b_1, b_2, \dots, b_n, a_n, a_{n-1}, \dots, a_1$ in clockwise order around a circle and include every chord

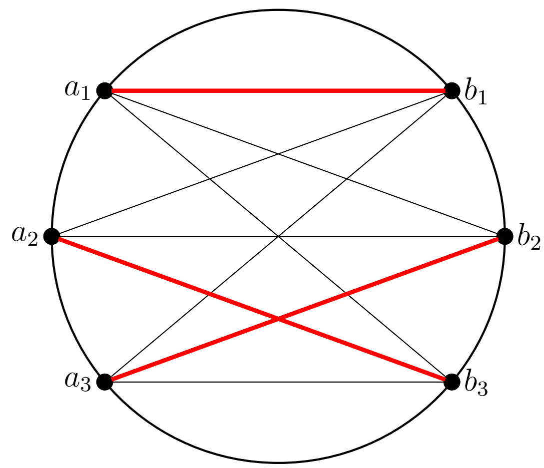

Figure 2.2: A chord diagram for the 3×3 comparability grid, where the chords in bold red represent the permutation graph F_π with $\pi = (1)(3, 2)$.

with one end in $\{a_1, \ldots, a_n\}$ and one end in $\{b_1, \ldots, b_n\}$ (see Figure 2.1). The circle graph of this chord diagram is isomorphic to the $n \times n$ comparability grid, where the vertex (i, j) of the comparability grid is associated with the chord $a_i b_{n+1-j}$. To obtain the permutation graph F_π of a permutation π of $\{1, \ldots, n\}$, represent each vertex $i \in \{1, \ldots, n\}$ by the chord connecting a_i to b_{π_i} (see Figure 2.2). In this manner, every n-vertex permutation graph is isomorphic to an induced subgraph of the $n \times n$ comparability grid.

$\square$

Now we find circle graphs as vertex-minors of permutation graphs.

Lemma 2.2.2. *Every circle graph on n vertices is a vertex-minor of a permutation graph on $3n$ vertices.*

Proof. Consider a chord diagram $\mathcal{C}$ for a circle graph G so that no two chords in $\mathcal{C}$ share an end (such a chord diagram exists). Let A be an arc of the unit circle whose ends are disjoint from $\mathcal{C}$. A chord is *crossing* if it has exactly one end in A. We may assume that there exists a non-crossing chord in $\mathcal{C}$ since otherwise G is itself a permutation graph and the result follows. We will construct a chord diagram $\mathcal{C}_2$ such that:

(*i*) $|\mathcal{C}_2| = |\mathcal{C}| + 2$,

(*ii*) $\mathcal{C}_2$ has fewer non-crossing chords than $\mathcal{C}$, and

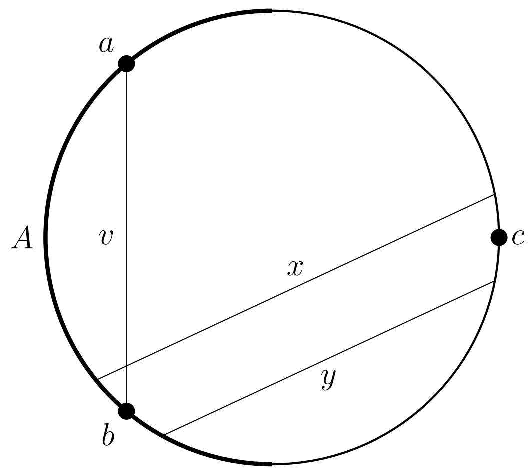

Figure 2.3: The chord diagram $\mathcal{C}_1$

(*iii*) the circle graph of $\mathcal{C}_2$ contains G as a vertex-minor.

The result follows by iterated applications of this construction.

Let $v \in \mathcal{C}$ be a non-crossing chord with ends a and b; we may assume that $a, b \in A$. Now select a point c on the unit circle disjoint from A and disjoint from $\mathcal{C}$. Let $\mathcal{C}_1$ be obtained from $\mathcal{C}$ by adding two parallel chords x and y immediately on either side of the chord $[b, c]$, and let $\mathcal{C}_2$ be obtained from $\mathcal{C}_1$ by replacing the chord v with the chord $[a, c]$. See Figure 2.3. Clearly $\mathcal{C}_2$ satisfies (*i*) and (*ii*). Let G_1 and G_2 denote the intersection graphs of $\mathcal{C}_1$ and $\mathcal{C}_2$ respectively. Then G_1 is isomorphic to $G_2 * x * y$ (recall that to locally complement at x, we "flip" one of the two arcs of the circle with the same ends as x) and G is an induced subgraph of G_1. Thus (*iii*) holds, as required. □

Finally, we combine the two lemmas to obtain the desired result.

Lemma 2.2.3. *Every circle graph on n vertices is isomorphic to a vertex-minor of the $3n \times 3n$ comparability grid.*

Proof. The lemma follows immediately from Lemmas 2.2.1 and 2.2.2. □

2.3 Main tools

The rest of this chapter is dedicated to proving Theorem 2.1.1, that every graph of sufficiently large rank-width has a vertex-minor isomorphic to the $n \times n$ comparability grid.

Our proof will use a result of Oum [99] to "obtain some connectivity", and a theorem of Kwon and Oum [85] to obtain a base case for induction. The main new tool is a Disentangling Lemma (Lemma 2.3.4) which "displays the connectivity" while "maintaining an induced subgraph". This lemma is particular to vertex-minors; it relies on the fact that there are three ways to remove a vertex (recall from Lemma 1.6.5 that a vertex can either be deleted, locally complemented and deleted, or pivot-deleted). In this section we will discuss these three results.

Obtaining connectivity

To prove the Grid Theorem for Vertex-Minors, it suffices to consider a vertex-minor-minimal graph of rank-width at least r_H. Oum [99] proved that in such a graph, any cut of low cut-rank has one side that is small. Formally, for a positive integer m and a function f, a graph G is (m, f)-*connected* if for every set $X \subseteq V(G)$ with $\rho(X) < m$, either $|X| \leq f(\rho(X))$ or $|V(G) - X| \leq f(\rho(X))$. Observe that if G is (m, f)-connected, t is an integer which is less than m, and S and T are disjoint subsets of $V(G)$ with cardinality greater than $f(t)$, then $\kappa(S, T) > t$ (recall from Section 1.6 that $\kappa(S, T)$ is the connectivity between S and T). So the following lemma lets us "obtain some connectivity".

Lemma 2.3.1 (Oum [99, Lemma 5.3])**.** *If $g_{2.3.1} : \mathbb{Z} \to \mathbb{Z}$ is the function defined by $g_{2.3.1}(n) = (6^n - 1)/5$, then, for any positive integer r, every vertex-minor-minimal graph of rank-width at least r is $(r, g_{2.3.1})$-connected.*

There is a partial converse to Lemma 2.3.1; any $(r, g_{2.3.1})$-connected graph with at least $3g_{2.3.1}(r-1)$ vertices has rank-width at least r. It follows that, with respect to proving the Grid Theorem for Vertex-Minors, it suffices to consider large $(r, g_{2.3.1})$-connected graphs.

Obtaining a base case

We will use the following theorem of Kwon and Oum [85] as a base case for an inductive proof (see Figure 2.4). In this result, a *star* is a tree with at most one non-leaf vertex.

Theorem 2.3.2 (Kwon and Oum [85, Theorem 1.6])**.** *There is a function $r_{2.3.2} : \mathbb{Z} \to \mathbb{Z}$ so that, for all positive integers m and k, every graph of rank-width at least $r_{2.3.2}(m, k)$ has a vertex-minor with m components, each of which is a star on $k + 1$ vertices.*

So, after locally complementing, we will have an induced subgraph with m components, each of which is a star on $k+1$ vertices. Informally, at the first inductive step, we will find

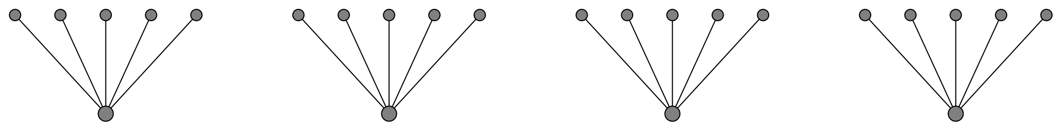

Figure 2.4: The base case: a graph with four components, each of which is a star on 6 vertices.

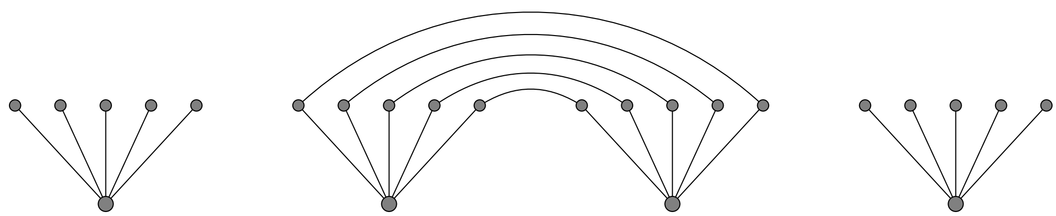

Figure 2.5: A depiction of a vertex-minor we might find after one inductive step.

the following graph as a vertex-minor (after "sacrificing" some of m and k). First, take the high-degree vertices x and y of two stars, and then change edges/non-edges between the neighbourhoods X of x and Y of x so that the submatrix of the adjacency matrix with rows X and columns Y has full rank. Then, in each inductive step, we "attach another star via a full-rank submatrix of the adjacency matrix". We do so in a "connected way" so that there is one "highly connected component" and all other components are stars. We will use Ramsey theory to control the "highly connected component"; the actual definitions will be given in Section 2.5.

Disentangling

The Disentangling Lemma will "display" the connectivity between disjoint sets of vertices S and T while "maintaining" the original subgraph induced by $S \cup T$. We will use the following definition to "display" the connectivity.

Let S and T be disjoint sets of vertices in a graph G; recall that Adj_G denotes the adjacency matrix of G, and all matrices are over the binary field. The *local connectivity of S and T*, denoted by $\sqcap_G(S,T)$ (or simply $\sqcap(S,T)$ if the graph is clear from context), is the rank of the submatrix of Adj_G with rows S and columns T. Notice that if $(S_1,\ldots,S_s)$

is a partition of S and $(T_1, \ldots, T_t)$ is a partition of T, then

$$\sqcap(S,T) \leq \sum_{i=1}^{s} \sum_{j=1}^{t} \sqcap(S_i, T_j);$$

we refer to this property as *sub-additivity.* Moreover, since a rank-k binary matrix has at most 2^k distinct columns, vertices in S have at most $2^{\sqcap(S,T)}$ distinct neighbourhoods within T.

Our proof will use Edmonds' Matroid Intersection Theorem (see [47]). So, for a set $T \subseteq V(G)$, let M_T denote the binary matroid represented by the submatrix of Adj_G with rows T and columns $V(G) - T$. Thus the ground set of M_T is $V(G) - T$, and each set $X \subseteq E(M_T)$ has rank $\sqcap(T, X)$. So a set $I \subseteq E(M_T)$ is independent in the matroid if $|I| = \sqcap(T, I)$; we refer to the independent sets of M_T as T-*independent* sets.

Recall that Theorem 1.6.6 of Oum [99] says that, if S and T are disjoint sets of vertices in a graph G and $v \in V(G) - (S \cup T)$, then at least two of $\kappa_{G-v}(S,T)$, $\kappa_{G*v-v}(S,T)$, and $\kappa_{G/v}(S,T)$ are equal to $\kappa_G(S,T)$. By repeatedly applying this theorem, we can find a pivot-minor $\widetilde{G}$ of G with $V(\widetilde{G}) = S \cup T$ so that $\sqcap_{\widetilde{G}}(S,T) = \kappa_G(S,T)$. Unfortunately, the graphs $\widetilde{G}[S \cup T]$ and $G[S \cup T]$ may be different. Informally, the Disentangling Lemma makes the connectivity between S and T "somewhat local" without changing the subgraph induced by $S \cup T$. Here is what we mean by "somewhat local".

Definition 2.3.3 (k-link)**.** *For a graph G with disjoint $S, T \subseteq V(G)$, a k-*link for (S,T) *is a pair (X_1, X_2) of k-element subsets of $V(G) - (S \cup T)$ such that X_1 is S-independent, X_2 is T-independent, and either*

(i) $X_1 = X_2$, or

(ii) X_1 and X_2 are disjoint, $\sqcap(X_1, X_2) = k$, all vertices in X_1 have the same set of neighbours in T, and all vertices in X_2 have the same set of neighbours in S.

We will not explicitly use this fact, but the motivation for k-links is that they certify high connectivity between S and T; that is, if there exists a k-link for (S,T), then $\kappa(S,T) \geq \frac{1}{3}k$. The Disentangling Lemma says that, if $\kappa(S,T) \gg k$ and $\sqcap(S,T) < k$, then we can find a k-link in a locally equivalent graph without changing the induced subgraph on $S \cup T$. We just assume that $\sqcap(S,T) < k$ for the sake of convenience. We believe that the lemma would not hold if "locally equivalent" were replaced by "equivalent up to pivoting". (We omit the proof, but the idea is to consider fundamental graphs of grids; see Section 1.6 for the definition of fundamental graphs.)

Lemma 2.3.4 (Disentangling Lemma)**.** *There is a function $L_{2.3.4} : \mathbb{Z} \to \mathbb{Z}$ so that, for every positive integer k, if G is a graph and $S, T \subseteq V(G)$ are disjoint sets with $\kappa(S,T) \geq L_{2.3.4}(k)$ and $\sqcap(S,T) < k$, then there exists a graph $\widetilde{G}$ that is locally equivalent to G such that $\widetilde{G}[S \cup T] = G[S \cup T]$ and $\widetilde{G}$ has a k-link for (S,T).*

Proof. Fix a positive integer k. Define $k_0 := 2^{k-1} + 1$ and

$$L_{2.3.4}(k) := 2^{k+k_0-2} + 2k - 1.$$

Suppose that the lemma fails for this function, and choose a counterexample (G, S, T) with $|V(G)|$ minimum. We begin with two claims.

Claim 2.3.4.1. *No two vertices in $V(G) - (S \cup T)$ have the same set of neighbours in $S \cup T$.*

Proof. Suppose that $u, v \in V(G) - (S \cup T)$ have the same set of neighbours in $S \cup T$. Consider the case that $uv \notin E(G)$. Then $G[S \cup T] = (G * v * u)[S \cup T]$, so $G[S \cup T]$ is a vertex-minor of both $G - v$ and $G * v - v$. However, by Theorem 1.6.6 (that at least two of the three ways to remove a vertex maintain connectivity), either $\kappa_{G-v}(S,T) = \kappa_G(S,T)$ or $\kappa_{G*v-v}(S,T) = \kappa_G(S,T)$, contradicting the minimality of G. In the case that $uv \in E(G)$, we see that $G[S \cup T]$ is a vertex-minor of both $G - v$ and $G/v = G \times uv - v$. Again we obtain a contradiction via Theorem 1.6.6. □

Claim 2.3.4.2. *There exist disjoint sets $Y_1, Y_2 \subseteq V(G) - (S \cup T)$ so that*

(i) $\sqcap(Y_1, Y_2) = |Y_1| = |Y_2| > 2^{k_0-1}$,

(ii) all vertices in Y_1 have the same set of neighbours in T, and

(iii) all vertices in Y_2 have the same set of neighbours in S.

Proof. Note that, if $X \subseteq V(G) - (S \cup T)$ is a common independent set of $M_S - T$ and $M_T - S$ with cardinality k, then (X, X) is a k-link for (S,T); however, there is no such k-link, and hence $M_S - T$ and $M_T - S$ do not have a common independent set of size k. So, by the Matroid Intersection Theorem, there is a partition (P, Q) of $V(G) - (S \cup T)$ so that $\sqcap(S, P) + \sqcap(T, Q) < k$.

Let $(P_1, \ldots, P_s)$ be the partition of P into equivalence classes of identical columns of $A[S, P]$ and let $(Q_1, \ldots, Q_t)$ be the partition of Q into equivalence classes of identical columns of $A[T, Q]$. Since $\sqcap(S,P) + \sqcap(T,Q) \leq k - 1$, we have $st \leq 2^{\sqcap(S,P)} 2^{\sqcap(T,Q)} \leq 2^{k-1}$.

Note that $\sqcap(S \cup Q, T \cup P) \geq \kappa(S,T) \geq L_{2.3.4}(k) \geq st2^{k_0-1} + 2k - 1$. Moreover, by sub-additivity and since $\sqcap(S,T) \leq k-1$,

$$\begin{aligned}\sqcap(Q,P) &\geq \sqcap(S \cup Q, T \cup P) - \sqcap(S,T) - \sqcap(S,P) - \sqcap(Q,T)\\ &\geq \sqcap(S \cup Q, T \cup P) - 2k + 2\\ &> st2^{k_0-1}.\end{aligned}$$

So, again using sub-additivity, there exist $i \in \{1, \ldots, t\}$ and $j \in \{1, \ldots, s\}$ such that $\sqcap(Q_i, P_j) > 2^{k_0-1}$. Now choose $Y_1 \subseteq Q_i$ and $Y_2 \subseteq P_j$ such that $\sqcap(Y_1, Y_2) = |Y_1| = |Y_2| > 2^{k_0-1}$. Now it is straightforward to see that (Y_1, Y_2) satisfies (i), (ii), and (iii), as required. □

By Claim 2.3.4.1 and part (ii) of Claim 2.3.4.2, no two vertices in Y_1 have the same set of neighbours in S. Then, since $|Y_1| > 2^{k_0-1}$, we have $\sqcap(Y_1, S) \geq k_0$. Let $Y_1' \subseteq Y_1$ be a k_0-element S-independent set. Since Y_1 is Y_2-independent, Y_1' is also Y_2-independent. So there is a k_0-element subset $Y_2' \subseteq Y_2$ that is Y_1'-independent. Now $|Y_2'| > 2^{k-1}$ so, by similar reasoning, there exist a k-element subset $X_2 \subseteq Y_2'$ that is T-independent and a k-element subset $X_1 \subseteq Y_1'$ that is X_2-independent. Then (X_1, X_2) is a k-link for (S,T), a contradiction. □

2.4 Ramsey theory

For the rest of this chapter, the vertex set of every graph G is an ordered set, and every set of vertices of G is considered as an ordered subset of $V(G)$; an *ordered set* is a sequence $X = (x_1, x_2, \ldots, x_k)$ with no repeated elements, and a *subset of* X is a subsequence of X. Oftentimes this assumption that the vertex-set of a graph is ordered will not matter, but it will matter when we discuss disjoint sets X and Y of $V(G)$ which are "coupled". This will mean that the bipartite subgraph of G which is "induced between" X and Y is one of a few specific graphs, like a perfect matching, where the vertices which are paired in the matching (for instance) are determined by the orderings of X and Y. So we use the ordering of $V(G)$ to induce fixed orderings on subsets which do not depend on the particular coupled pair under consideration. We sometimes remind the reader of these conventions by saying that the ordering of X is *induced by* the ordering of $V(G)$.

Furthermore, if H is a subgraph of G, we mean that $V(H)$ is a subset of $V(G)$ as ordered sets. Two graphs are *isomorphic* if they are isomorphic as graphs with unordered vertex sets. For each positive integer n, we fix a lexicographic ordering on the vertex set of

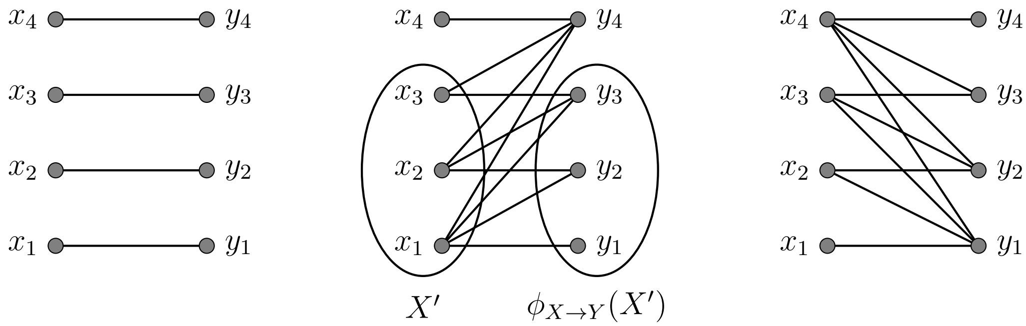

Figure 2.6: A coupled matching (left), up-coupled half graph (middle), and down-coupled half graph.

the $n \times n$ comparability grid. The rest this chapter is dedicated to proving Theorem 2.1.1 with these conventions; this is easily seen to be equivalent to the original statement.

Let X and Y be disjoint ordered sets of cardinality k, and let $X' \subseteq X$. We write $\phi_{X\to Y}(X')$ for the subset of Y induced by the ordering of X' with respect to X. That is, if $X = (x_1, x_2, \ldots, x_k)$, $Y = (y_1, y_2, \ldots, y_k)$, and $X' = (x_{i_1}, x_{i_2}, \ldots, x_{i_{k'}})$, then $\phi_{X\to Y}(X') = (y_{i_1}, y_{i_2}, \ldots, y_{i_{k'}})$. We also write $\phi_{X\to X}(X')$ for the set X' itself.

In the rest of this section we review some Ramsey theory for graphs with ordered vertex-sets. We give a number of definitions so that we can simply state the possible outcomes. First, for a graph G with disjoint sets $X, Y \subseteq V(G)$, we say that X and Y are *anticomplete* if G has no edges with one end in X and one end in Y, and *complete* if for all $x \in X$ and $y \in Y$, $xy \in E(G)$. We say that X and Y are *homogeneous* if they are either complete or anticomplete. So far these definitions do not care about the ordering, but the following definitions will.

Let $X = (x_1, x_2, \ldots, x_k)$ and $Y = (y_1, y_2, \ldots, y_k)$ be disjoint sets of vertices in a graph G with orderings induced by the ordering of $V(G)$. Refer to Figure 2.6 for the following definitions. We say that (X, Y) is

(1) a *coupled matching* if for each $i \in \{1, 2, \ldots, k\}$, the neighbourhood of x_i within Y is (y_i),

(2) an *up-coupled half graph* if for each $i \in \{1, 2, \ldots, k\}$, the neighbourhood of x_i within Y is $(y_i, y_{i+1}, \ldots, y_k)$, and

(3) a *down-coupled half graph* if for each $i \in \{1, 2, \ldots, k\}$, the neighbourhood of x_i within Y is $(y_1, y_2, \ldots, y_i)$.

We say that (X, Y) is the *complement of a coupled matching* if (X, Y) is a coupled matching in the complement of G. Similarly we will talk about the *complement of a down-coupled half graph* and the *complement of an up-coupled half graph*. If (X, Y) is either a down-coupled half graph, an up-coupled half graph, or one of their complements, then we say (X, Y) is a *coupled half graph*. If (X, Y) is either a coupled matching, the complement of a coupled matching, or a coupled half graph, we say X *and* Y *are coupled*. Notice that if X and Y are coupled and $X' \subseteq X$, then X' and $\phi_{X \to Y}(X')$ are coupled.

If X and Y are disjoint coupled sets in a graph G, then $\sqcap(X, Y) \geq |X| - 1$. The next result, due to Ding, Oporowski, Oxley, and Vertigan [40], shows that a partial converse holds; that is, if $\sqcap(X, Y) \gg k$, then, up to possibly reordering the vertices in Y, there are k-element subsets of X and Y that are coupled.

Lemma 2.4.1 (Ding, Oporowski, Oxley, and Vertigan [40, Theorem 2.3])**.** *There is a function $R_{2.4.1} : \mathbb{Z} \to \mathbb{Z}$ so that for any positive integer k, if G is a graph and $X, Y \subseteq V(G)$ are disjoint sets with $\sqcap(X, Y) \geq R_{2.4.1}(k)$, then there exist k-element subsets $X' \subseteq X$ and $Y' \subseteq Y$ that are coupled in a graph obtained from G by reordering the vertices in Y.*

We will also use the following version of a theorem of Ramsey.

Ramsey's Theorem. *For each integer k, there is a function $R_k : \mathbb{Z} \to \mathbb{Z}$ so that for any positive integer n, every k-edge-coloured clique on at least $R_k(n)$ vertices contains a monochromatic clique of size n.*

The following result can be proven using the multicolour version of Ramsey's Theorem. We omit the proof; the idea is to create an auxiliary graph with $|X|$ vertices and an edge colour for each "configuration" between two vertices in X and two vertices in Y.

Lemma 2.4.2. *There is a function $R_{2.4.2} : \mathbb{Z} \to \mathbb{Z}$ so that for any positive integer k, if X and Y are disjoint sets of vertices in a graph G with $|X| = |Y| \geq R_{2.4.2}(k)$, then there is a k-element subset $X' \subseteq X$ such that X' and $\phi_{X \to Y}(X')$ are either coupled or homogeneous.*

The following lemma is a corollary of Lemma 2.4.2; we omit the proof.

Lemma 2.4.3. *There is a function $R_{2.4.3} : \mathbb{Z} \to \mathbb{Z}$ so that for any positive integer k, if X and Y are disjoint sets of vertices in a graph G with $|X|, |Y| \geq R_{2.4.3}(k)$, then there exist k-element sets $X' \subseteq X$ and $Y' \subseteq Y$ such that X' and Y' are homogeneous.*

For a function $R : \mathbb{Z} \to \mathbb{Z}$ and an integer $n > 1$, we inductively define $R^{(n)}$ to be the function $R(R^{(n-1)})$, where for the base case $R^{(1)} = R$.

2.5 Building a constellation

Roughly speaking, a "large constellation" in a graph is an induced subgraph consisting of many large stars coupled together in a "connected way". The proof of Theorem 2.1.1 then consists of two parts; in this section we prove that, up to local equivalence and reordering vertices, every graph of sufficiently large rank-width contains a large constellation. In the next section we prove that every graph containing a sufficiently large constellation has a vertex-minor isomorphic to the $n \times n$ comparability grid.

Recall that a *coclique* is a set of pairwise non-adjacent vertices.

Definition 2.5.1 (Constellations)**.** *Let G be a graph, let n and k be positive integers, and let m be a non-negative integer. An (n, m, k)-*constellation *in G is a tuple $(H, (W_h : h \in H), K)$ such that*

(i) $H \subseteq V(G)$ is an $(n+m)$-vertex coclique,

(ii) the sets $(W_h : h \in H)$ are disjoint k-vertex cocliques in $G-H$, with orderings induced by the ordering of $V(G)$,

(iii) K is a connected n-vertex graph with $V(K) \subseteq H$,

(iv) for every $h \in H$, the set W_h is complete to $\{h\}$ and anticomplete to $H - \{h\}$,

(v) for distinct $u, v \in H$, the pair (W_u, W_v) is either a coupled half graph or a coupled matching if $uv \in E(K)$, and is anticomplete otherwise.

If $\mathcal{C} = (H, (W_h : h \in H), K)$ is an (n, m, k)-constellation in G, then we write $H(\mathcal{C})$ for H, we write $K(\mathcal{C})$ for K, and for each $h \in H$ we write $W_h^{\mathcal{C}}$ for W_h. We denote the union of the sets $(\{v\} \cup W_v : v \in H)$ by $V(\mathcal{C})$, we denote the union of the sets $(\{v\} \cup W_v : v \in V(K))$ by $A(\mathcal{C})$, and we denote $V(\mathcal{C}) - A(\mathcal{C})$ by $B(\mathcal{C})$. We sometimes use a sequence of constellations $\mathcal{C}_0, \mathcal{C}_1, \ldots$, and in that case we write W_h^0 for $W_h^{\mathcal{C}_0}$, and likewise for $\mathcal{C}_1$, and so on. For $h \in H(\mathcal{C})$ and $X \subseteq W_h^{\mathcal{C}}$, we write $\mathcal{C}|X$ for

$$\left(H, \left(\phi_{W_h^{\mathcal{C}} \to W_z^{\mathcal{C}}}(X) : z \in H\right), K\right).$$

Notice that $\mathcal{C}|X$ is an $(n, m, |X|)$-constellation in G.

This section is devoted to proving that, for positive integers n and k, every graph with sufficiently large rank-width contains, up to local equivalence and reordering vertices, an $(n, 0, k)$-constellation. To build constellations we use "augmentations".

Definition 2.5.2 (Weak augmentations). *For positive integers n, m, and k, a* weak (n,m,k)-augmentation *in a graph G is a tuple $(\mathcal{C}, x, y, X_1, X_2)$ such that $\mathcal{C}$ is an (n,m,k)-constellation; $x \in V(K(\mathcal{C}))$ and $y \in H(\mathcal{C}) - V(K(\mathcal{C}))$; and (X_1, X_2) is a pair of k-vertex subsets of $V(G) - (V(\mathcal{C}))$, with orderings induced by the ordering of $V(G)$, such that $W_x^{\mathcal{C}}$ and X_1 are coupled, $W_y^{\mathcal{C}}$ and X_2 are coupled, and either*

(i) $X_1 = X_2$, or

(ii) X_1 and X_2 are disjoint and coupled, all vertices in X_1 have the same set of neighbours in $B(\mathcal{C})$, and all vertices in X_2 have the same set of neighbours in $A(\mathcal{C})$.

Lemma 2.5.3. *There is a function $k_{2.5.3} : \mathbb{Z}^3 \to \mathbb{Z}$ so that, for all positive integers n, m, k_0, k_1, and k_2 with $k_1 \geq k_{2.5.3}(n, m, k_0)$ and $k_2 \geq k_1$, if $\mathcal{C}$ is an (n, m, k_2)-constellation in a graph G and $\kappa(A(\mathcal{C}), B(\mathcal{C})) \geq k_1$, then there exists a graph that is equivalent to G up to local complementation and reordering vertices and contains a weak (n, m, k_0)-augmentation.*

Proof. For positive integers n, m, and k_0 we define

$$t := n\left(mR^{(3)}_{2.4.1}(k_0) + m + 1\right), \text{ and}$$
$$k_{2.5.3}(n, m, k_0) := L_{2.3.4}(t).$$

Now let k_1 and k_2 be positive integers such that

$$k_2 \geq k_1 \geq k_{2.5.3}(n, m, k_0),$$

and let $\mathcal{C}$ be an (n, m, k_2)-constellation in a graph G with $\kappa(A(\mathcal{C}), B(\mathcal{C})) \geq k_1$. By Lemma 2.3.4, there is a graph G_0 that is locally equivalent to G such that $G_0[V(\mathcal{C})] = G[V(\mathcal{C})]$, and G_0 contains a t-link (X_1, X_2) for $(A(\mathcal{C}), B(\mathcal{C}))$. Up to local equivalence we may assume that $G_0 = G$.

By sub-additivity, $\sqcap(A(\mathcal{C}) - H(\mathcal{C}), X_1) \geq t - n$. Let $X_0 \subseteq A(\mathcal{C}) - H(\mathcal{C})$ be a $(t-n)$-element X_1-independent set. Now let $t' = mR^{(3)}_{2.4.1}(k_0) + m$. Thus

$$|X_0| = t - n = nt'.$$

Thus, by the pigeonhole principle, there exist $x \in V(K(\mathcal{C}))$ and $X_0' \subseteq X_0$ of cardinality t' so that $X_0' \subseteq W_x^{\mathcal{C}}$. Note that $\sqcap(X_0', X_1) = |X_0'|$. By the definition of X_1 and X_2, there exist $X_1' \subseteq X_1$ and $X_2' \subseteq X_2$ so that (X_1', X_2') is a t'-link for $(X_0', B(\mathcal{C}))$. By the same reasoning, there exist a vertex $y \in H(\mathcal{C}) - V(K(\mathcal{C}))$ and $R^{(3)}_{2.4.1}(k_0)$-vertex subsets $X_3'' \subseteq W_y^{\mathcal{C}}$, $X_2'' \subseteq X_2'$, and $X_1'' \subseteq X_1'$ so that (X_1'', X_2'') is a $|X_3''|$-link for (X_0', X_3'').

Next, we apply Lemma 2.4.1 to the sets X_0' and X_1'' so that, after possibly reordering the vertices in X_1'', there exist $R^{(2)}_{2.4.1}(k_0)$-element subsets $Y_0 \subseteq X_0'$ and $Y_1 \subseteq X_1''$ so that Y_0 and Y_1 are coupled. The claim follows by repeating this process one or two more times depending on whether $X_1 = X_2$, and possibly reordering the vertices in X_2'' and X_3''. Note that it is fine to reorder vertices in X_3'' since $y \in H(\mathcal{C}) - V(K(\mathcal{C}))$. □

When taking restrictions of a weak (n, m, k)-augmentation $(\mathcal{C}, x, y, X_1, X_2)$, we need to respect orders between the sets X_1, X_2, and $(W_h^{\mathcal{C}} : h \in V(K(\mathcal{C})) \cup \{y\})$, but not with the sets $(W_z^{\mathcal{C}} : z \in H(\mathcal{C}) - (V(K(\mathcal{C})) \cup \{y\}))$. To be more precise, consider a k'-element subset $Y_1 \subseteq X_1$. Let $Y_2 \coloneqq \phi_{X_1 \to X_2}(Y_1)$ and let $\mathcal{C}'$ be an (n, m, k')-constellation such that $H(\mathcal{C}') = H(\mathcal{C})$; for each $h \in V(K(\mathcal{C})) \cup \{y\}$ we have $W_h^{\mathcal{C}'} = \phi_{X_1 \to W_z^{\mathcal{C}}}(Y_1)$; and for each $z \in H(\mathcal{C}) - (V(K(\mathcal{C})) \cup \{y\})$ the set $W_z^{\mathcal{C}'}$ is a k'-element subset of $W_z^{\mathcal{C}}$. Then $(\mathcal{C}', x, y, Y_1, Y_2)$ is an (n, m, k')-augmentation.

Definition 2.5.4 (Augmentations)**.** *For positive integers n, m, and k, we define an* (n, m, k)-augmentation *to be a weak (n, m, k)-augmentation $(\mathcal{C}, x, y, X_1, X_2)$ such that for each $i \in \{1, 2\}$:*

(i) X_i is either a clique or a coclique,

(ii) for all $h \in V(K(\mathcal{C})) \cup \{y\}$, the sets $W_h^{\mathcal{C}}$ and X_i are either homogeneous or coupled, with orderings induced by the ordering of $V(G)$,

(iii) for all $h \in H(\mathcal{C}) - (V(K(\mathcal{C})) \cup \{y\})$, the sets $W_h^{\mathcal{C}}$ and X_i are homogeneous, and

(iv) for all $h \in H(\mathcal{C})$, the sets $\{h\}$ and X_i are homogeneous.

Lemma 2.5.5. *There is a function $k_{2.5.5} : \mathbb{Z}^3 \to \mathbb{Z}$ so that, for all positive integers n, m, k_0, and k_1 with $k_1 \geq k_{2.5.5}(n, m, k_0)$, if G is a graph containing a weak (n, m, k_1)-augmentation, then G contains an (n, m, k_0)-augmentation.*

Proof. For positive integers n, m, and k_0 we define

$$k_{2.5.5}(n, m, k_0) \coloneqq R_2^{(2)}\left(R_{2.4.2}^{(2n)}\left(R_{2.4.3}^{(2m-2)}\left(k_0 \cdot 2^{2(m+n)}\right)\right)\right).$$

Now consider a weak (n, m, k_1)-augmentation $(\mathcal{C}, x, y, X_1, X_2)$ with $k_1 \geq k_{2.5.5}(n, m, k_0)$.

By applying Ramsey's Theorem first on X_1 and then on the specified subset of X_2, we can get statement (1) to hold. Now, for each $i \in \{1, 2\}$ and $h \in V(K(\mathcal{C})) \cup \{y\}$ so that $W_h^{\mathcal{C}}$ and X_i are not already coupled, we successively apply Lemma 2.4.2 to get statement (2) to hold. Note that we apply the lemma at most $2n$ times since $W_x^{\mathcal{C}}$ and X_1 are already coupled, as are $W_y^{\mathcal{C}}$ and X_2. Then, for each $i \in \{1,2\}$ and for each $h \in H(\mathcal{C}) - (V(K(\mathcal{C})) \cup \{y\})$, we successively apply Lemma 2.4.3 to get statement (3) to hold. Finally we get statement (4) to hold by, for each $i \in \{1,2\}$ and each $h \in H(\mathcal{C})$, successively applying a majority argument to the edges from h to what remains of X_i. $\square$

We can now prove the main result of this section.

Lemma 2.5.6. *There is a function $r_{2.5.6} : \mathbb{Z}^3 \to \mathbb{Z}$ so that, for all positive integers n, m, and k, every graph of rank-width at least $r_{2.5.6}(n,m,k)$ has a vertex-minor which contains an (n,m,k)-constellation, after possibly reordering vertices.*

Proof. For $n = 1$, the result is true by Theorem 2.3.2 with $r_{2.5.6}(1,m,k) := r_{2.3.2}(m+1,k)$. Now assume that for some fixed integer $n \geq 2$, for all positive integers m and k, such a function $r_{2.5.6}(n-1,m,k)$ exists. Now, for fixed m and k, we will show that $r_{2.5.6}(n,m,k)$ exists. Define

$$\begin{aligned} k_1 &:= k_{2.5.3}\left(n-1, m+1, k_{2.5.5}(n-1, m+1, k+4)\right), \\ k_0 &:= g_{2.3.1}(k_1), \text{ and} \\ r_{2.5.6}(n,m,k) &:= \max\left(r_{2.5.6}\big(n-1, m+1, k_0\big), k_1\right). \end{aligned}$$

Towards a contradiction, suppose that G is a graph of rank-width at least $r_{2.5.6}(n,m,k)$ that does not have a vertex-minor containing an (n,m,k)-constellation, after possibly reordering vertices. Choose such a graph with $|V(G)|$ minimum; thus no proper vertex-minor of G has rank-width at least $r_{2.5.6}(n,m,k)$. So, by Lemma 2.3.1, the graph G is $(k_1,\, g_{2.3.1})$-connected.

We may assume that G contains an $(n-1, m+1, k_0)$-constellation $\mathcal{C}_0$. Then, since $\min(|A(\mathcal{C}_0)|,\, |B(\mathcal{C}_0)|) \geq g_{2.3.1}(k_1)$, we have $\kappa(A(\mathcal{C}_0), B(\mathcal{C}_0)) \geq k_1$. Then, by Lemmas 2.5.3 and 2.5.5, there is a graph equivalent to G up to local complementation and reordering vertices that contains an $(n-1, m+1, k+4)$-augmentation.

We choose a graph G_1 that is locally equivalent to G and has an $(n-1, m+1, t)$-augmentation $(\mathcal{C}_1, z_1, z_2, Z_1, Z_2)$ such that:

(i) either

- $Z_1 = Z_2$ and $t = k+2$, or
- $Z_1 \neq Z_2$ and $t = k+4$,

(ii) subject to (1) we have $Z_1 = Z_2$ if possible, and

(iii) subject to (2) the vertex z_2 is complete to Z_2 if possible.

We may assume that $G_1 = G$.

Claim 2.5.6.1. *There is a vertex in $W^1_{z_2} \cup \{z_2\}$ with at least $t-1$ neighbours in either Z_1 or Z_2.*

Proof. Suppose otherwise, then, by the assumption,

- z_2 is anticomplete to $Z_1 \cup Z_2$,
- $(W^1_{z_2}, Z_2)$ is a coupled matching, and
- if $Z_1 \neq Z_2$, then $W^1_{z_2}$ is anticomplete to Z_1.

Note that each vertex in $W^1_{z_2}$ has degree 2 in $G[V(\mathcal{C}_1) \cup Z_1 \cup Z_2]$. Let G' be the graph obtained from G by locally complementing on each vertex in $W^1_{z_2}$. Note that $(\mathcal{C}_1, z_1, z_2, Z_1, Z_2)$ is an $(n-1, m+1, t)$-augmentation in G' and z_2 is complete to Z_2 in G', contrary to our choice of G_1 and $(\mathcal{C}_1, z_1, z_2, Z_1, Z_2)$. □

We break the proof into two cases; there is a lot of overlap in the proofs, but it is less awkward with the cases separated.

Case 1: *There is a vertex $v \in W^1_{z_2} \cup \{z_2\}$ with at least $t-1$ neighbours in Z_1.*

We choose $G_2 \in \{G, G * v\}$ so that the set of neighbours of v in Z_1 is a coclique in G_2. Let w be the first vertex in Z_1 that is a neighbour, in G_2, of v, and let $G_3 \coloneqq G_2 \times vw$. We will show that G_3 contains an (n, m, k)-constellation $\mathcal{C}_3$, giving a contradiction.

Let $H_3 \coloneqq (H(\mathcal{C}_1) - \{z_2\}) \cup \{w\}$, let W^3_w denote a k-element subset of the neighbours, in G_2, of v in $Z_1 - \{w\}$, and, for each $x \in H(\mathcal{C}_1) - \{z_2\}$, let $W^3_x \coloneqq \phi_{Z_1 \to W^1_x}(W^3_w)$. Note that since w is the first neighbour of v in Z_1, the vertex w is either complete or anticomplete to each W^3_x in G_2.

Now let K_3 denote the graph obtained from $K(\mathcal{C}_1)$ by adding the vertex w and all edges wx where $x \in V(K(\mathcal{C}_1))$ and (W^3_w, W^3_x) is coupled; since $(W^3_w, W^3_{z_1})$ is coupled, K_3 is connected. Finally let $\mathcal{C}_3 \coloneqq (H_3, (W^3_x : x \in H_3), K_3)$. We claim that $\mathcal{C}_3$ is an (n, m, k)-constellation in G_3 which follows from the description of pivoting by complementing between three sets and the following observations about adjacencies in G_2:

- v is anticomplete to $V(\mathcal{C}_1) - (\{z_2\} \cup W^1_{z_2})$ and is complete to W^3_w,
- for each $x \in H_3 - \{w\}$, the vertex x is complete or anticomplete to $W^3_w \cup \{w\}$,
- for each $x \in H_3 - V(K_3)$, the set W^3_x is complete or anticomplete to $W^3_w \cup \{w\}$, and
- for each $x \in V(K_3)$, the vertex w is complete or anticomplete to W^3_x.

Case 2: *No vertex in $W^1_{z_2} \cup \{z_2\}$ has at least $t-1$ neighbours in Z_1.*

Then, by the above claim, there is a vertex $v \in W^1_{z_2} \cup \{z_2\}$ with $t-1$ neighbours in Z_2. Thus $Z_2 \neq Z_1$ and, by the definition of an augmentation, v is anticomplete to Z_1.

We choose $G_2 \in \{G, G * v\}$ so that the set of neighbours of v in Z_2 is a coclique in G_2. Let w be the first neighbour, in G_2, of v in Z_2 and let $G_3 := G_2 \times vw$. We will show that G_3 contains an $(n-1, m+1, k+2)$-augmentation $(\mathcal{C}_3, z_1, w, X, X)$ for some $\mathcal{C}_3$ and X, giving a contradiction to our choice of G_1 and $(\mathcal{C}_1, z_1, z_2, Z_1, Z_2)$.

Let $H_3 := (H(\mathcal{C}_1) - \{z_2\}) \cup \{w\}$, let W^3_w denote a $(k+2)$-element subset of the set of neighbours of v in $Z_2 - \{w\}$, and, for each $x \in H(\mathcal{C}_1) - \{z_2\}$, let $W^3_x := \phi_{Z_2 \to W^1_x}(W^3_w)$. By the choice of w to be the first neighbour of v in Z_2, the vertex w is either complete or anticomplete to each W^3_x in G_2.

Finally let $\mathcal{C}_3 := (H_3, (W^3_x : x \in H_3), K(\mathcal{C}_1))$ and let $X := \phi_{Z_2 \to Z_1}(W^3_w)$. Again by the choice of w, the vertex w is either complete or anticomplete to X in G_2. We claim that $(\mathcal{C}_3, z_1, w, X, X)$ is an $(n-1, m+1, k+2)$-augmentation in G_3 which follows from the description of pivoting by complementing between three sets and the following observations about adjacencies in G_2:

- v is anticomplete to both $V(\mathcal{C}_1) - (\{z_2\} \cup W^1_{z_2})$ and X, and is complete to W^3_w,
- for each $x \in H_3 - \{w\}$, the vertex x is complete or anticomplete to $W^3_w \cup \{w\}$,
- for each $x \in H_3 - (V(K(\mathcal{C}_1)) \cup \{w\})$, the set W^3_x is complete or anticomplete to $W^3_w \cup \{w\}$,
- for each $x \in V(K(\mathcal{C}_1))$, the vertex w is complete or anticomplete to W^3_x, and
- X is complete or anticomplete to w.

□

2.6 Extracting a comparability grid

It remains to prove that every graph containing a sufficiently large constellation has a vertex-minor isomorphic to the $n \times n$ comparability grid. Henceforth we will only consider (n, m, k)-constellations with $m = 0$ and will abbreviate these to (n, k)-*constellations*.

We will apply the following well-known Ramsey-type lemma to reduce to constellations whose associated graphs are stars, paths, or cliques.

Lemma 2.6.1. *There is a function $n_{2.6.1} : \mathbb{Z} \to \mathbb{Z}$ such that for every positive integer k, every connected graph on at least $n_{2.6.1}(k)$ vertices has a k-vertex induced subgraph that is either a path, a star, or a clique.*

The following result gives a sufficient condition for a graph to contain arbitrary n-vertex graphs as vertex-minors.

Lemma 2.6.2. *Let $Z = (z_1, z_2, \ldots, z_n)$ be a set of vertices in a graph G, with ordering induced by the ordering of $V(G)$, so that there are distinct components $(A_{i,j} : 1 \leq i < j \leq n)$ of $G - Z$ so that z_i and z_j have neighbours in $A_{i,j}$ and each neighbour of $V(A_{i,j})$ is in $(z_i, z_{i+1}, \ldots, z_j)$. Then every graph with vertex set Z is a vertex-minor of G.*

Proof. Let H be a graph with vertex set Z. We say that a pair (i, j), where $1 \leq i < j \leq n$, is *fixed* if for each $i' \leq i$ and $j' \geq j$ the vertices $z_{i'}$ and $z_{j'}$ are adjacent in either both of or neither of H and G. If all edges are fixed then H is an induced subgraph of G. Among all non-fixed pairs choose (i, j) with i minimum and, subject to that, j is maximum. We will fix (i, j), without unfixing any other pair, by locally complementing in $A_{i,j}$; the result follows by repeating this until all pairs are fixed.

There is an induced path $P = (v_1, v_2, \ldots, v_k)$ in $A_{i,j}$ such that z_i is adjacent to v_1 but not to any of $v_2, \ldots, v_k$ and z_j is adjacent to v_k but not to any of $v_1, \ldots, v_{k-1}$ (if z_i and z_j share a neighbour then it is possible that $k = 1$). Replacing G with $G * v_1 * v_2 * \cdots * v_k$ fixes (i, j) without unfixing any other pair, as required. □

The following two results are applications of Lemma 2.6.2 to constellations.

Lemma 2.6.3. *For any n-vertex graph H, if $\mathcal{C}$ is an $\left(n, \binom{n}{2}\right)$-constellation in a graph G such that $K(\mathcal{C})$ is either a path or a clique, and for each edge uv of $K(\mathcal{C})$ the pair $(W_u^{\mathcal{C}}, W_v^{\mathcal{C}})$ is a coupled matching, then G has a vertex-minor isomorphic to H.*

Proof. We may assume that $V(G) = V(\mathcal{C})$. Let $H(\mathcal{C}) = \{z_1, \dots, z_n\}$ where, if $K(\mathcal{C})$ is a path, then the vertices are in the order $(z_1, \dots, z_n)$ on the path. Note that $G - H(\mathcal{C})$ has $\binom{n}{2}$ components which we label $(G_{i,j} : 1 \le i < j \le n)$; each of these components is isomorphic to $K(\mathcal{C})$. For each $1 \le i < j \le n$, let $A_{i,j}$ denote the (unique) shortest path from the neighbour of z_i in $G_{i,j}$ to the neighbour of z_j in $G_{i,j}$. The result follows by applying Lemma 2.6.2 to the subgraph of G induced on the union of $H(\mathcal{C})$ together with the sets $(V(A_{i,j}) : 1 \le i < j \le n)$. □

Lemma 2.6.4 (Star constellations)**.** *For any n-vertex graph H, if $\mathcal{C}$ is an $\left(\binom{n}{2}+1, n+2\right)$-constellation in a graph G such that $K(\mathcal{C})$ is a star, then G has a vertex-minor isomorphic to H.*

Proof. We may assume that $V(G) = V(\mathcal{C})$. Let $H(\mathcal{C}) = \{h\} \cup \{v_{i,j} : 1 \le i < j \le n\}$, where h is the hub of the star $K(\mathcal{C})$, and let $W_h^{\mathcal{C}} = (z_0, z_1, \dots, z_n, z_{n+1})$. Note that, for each $1 \le i < j \le n$, the graph $G[W_{v_{i,j}}^{\mathcal{C}} \cup \{v_{i,j}\}]$ is a component of $G - (W_h^{\mathcal{C}} \cup \{h\})$. By locally complementing and deleting vertices within the subgraph $G[W_{v_{i,j}}^{\mathcal{C}} \cup \{v_{i,j}\}]$ we will obtain a connected graph $A_{i,j}$ such that z_i and z_j have neighbours in $A_{i,j}$, and each neighbour of $V(A_{i,j})$ is in $(z_i, z_{i+1}, \dots, z_j)$. Then the result will follow by applying Lemma 2.6.2 to the subgraph induced on the union of $\{z_1, \dots, z_n\}$ and the sets $(V(A_{i,j}) : 1 \le i < j \le n)$.

In the case that $(W_{v_{i,j}}^{\mathcal{C}}, W_h^{\mathcal{C}})$ is a coupled matching, we take $A_{i,j}$ to be the path in $G[W_{v_{i,j}}^{\mathcal{C}} \cup \{v_{i,j}\}]$ connecting the neighbours of z_i and z_j. Thus we may assume that $(W_{v_{i,j}}^{\mathcal{C}}, W_h^{\mathcal{C}})$ is a coupled half graph. First suppose that $(W_{v_{i,j}}^{\mathcal{C}}, W_h^{\mathcal{C}})$ is either a down-coupled half graph or the complement of an up-coupled half graph. Then, for each $k \in \{0, 1, \dots, n\}$, there is a vertex $x_k \in W_{v_{i,j}}^{\mathcal{C}}$ whose neighbours in $W_h^{\mathcal{C}}$ are $\{z_0, \dots, z_k\}$. Let $G' = G \times v_{i,j}x_{i-1}$. Then, in G', the set of neighbours of x_j in $\{z_1, \dots, z_n\}$ is $\{z_i, \dots, z_j\}$, and we take $A_{i,j} = G'[x_j]$.

The final case that $(W_{v_{i,j}}^{\mathcal{C}}, W_h^{\mathcal{C}})$ is either an up-coupled half graph or the complement of a down-coupled half graph is similar. In this case, for each $k \in \{1, \dots, n+1\}$, there is a vertex $x_k \in W_{v_{i,j}}^{\mathcal{C}}$ whose set of neighbours in $W_h^{\mathcal{C}}$ is $\{z_k, \dots, z_{n+1}\}$. We set $G' = G \times v_{i,j}x_{j+1}$ and then take $A_{i,j} = G'[x_i]$. □

Next we consider constellations whose associated graphs are cliques. In order to recognize comparability grids we use the following easy characterization.

Lemma 2.6.5. *For any positive integer n, if $(X_1, \dots, X_n)$ is a partition of the vertices of a graph G into n-vertex cliques, with orderings induced by the ordering of $V(G)$, such that, for each $1 \le i < j \le n$, the pair (X_i, X_j) is an up-coupled half graph, then G is isomorphic to the $n \times n$ comparability grid.*

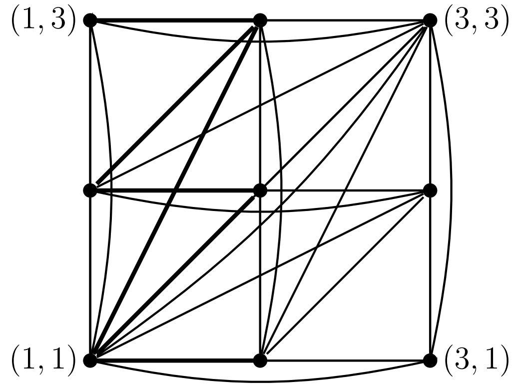

Figure 2.7: The 3×3 comparability grid.

Proof. Recall that the $n \times n$ comparability grid has vertex set $\{(i,j) : i,j \in \{1,2,\ldots,n\}\}$ where there is an edge between vertices (i,j) and (i',j') if either $i \le i'$ and $j \le j'$, or $i \ge i'$ and $j \ge j'$. Relabel the vertices of G so that, for each $i \in \{1,\ldots,n\}$, we have $X_i = ((i,1),(i,2),\ldots,(i,n))$. Then G is the $n \times n$ comparability grid. See Figure 2.7, where the edges between X_1 and X_2 are bolded. □

Lemma 2.6.6. *For any positive integer n, if $(X_1,\ldots,X_{n^2})$ is a partition of the vertices of a graph G into sets of cardinality n^2 such that, for each $1 \le i < j \le n$, the pair (X_i,X_j) is either an up-coupled half graph or the complement of a down-coupled half graph, then there is an induced subgraph of G that is isomorphic to the $n \times n$ comparability grid.*

Proof. Suppose that $X_i = (x_{i,1},\ldots,x_{i,n^2})$ for each $i \in \{1,\ldots,n^2\}$. Now, for each $i,j \in \{1,\ldots,n\}$, let $y_{i,j} := x_{(i-1)n+j,(j-1)n+i}$ and let $Y_i = (y_{i,1},\ldots,y_{i,n})$. Thus $Y_1,\ldots,Y_n$ are cliques and, for each $1 \le i < j \le n$, the pair (Y_i,Y_j) is an up-coupled half graph, so the result follows from Lemma 2.6.5. □

Lemma 2.6.7 (Clique constellations)**.** *There are functions $n_{2.6.7} : \mathbb{Z} \to \mathbb{Z}$ and $k_{2.6.7} : \mathbb{Z} \to \mathbb{Z}$ such that, for any positive integer n, if $\mathcal{C}$ is an $(n_{2.6.7}(n), k_{2.6.7}(n))$-constellation in a graph G such that $K(\mathcal{C})$ is a clique, then G has a vertex-minor isomorphic to the $n \times n$ comparability grid.*

Proof. Recall that the function R_k is defined in Ramsey's Theorem. For a positive integer n we define

$$\begin{aligned} n_{2.6.7}(n) &:= R_3(n^2) \text{ and} \\ k_{2.6.7}(n) &:= \max\left(n^2, \binom{n^2}{2}\right). \end{aligned}$$

Let $\mathcal{C}$ be an $(n_{2.6.7}(n), k_{2.6.7}(n))$-constellation in a graph G such that $K(\mathcal{C})$ is a clique and let $H(\mathcal{C}) = (h_1, \ldots, h_{n_1})$, where $n_1 = n_{2.6.7}(n)$. Towards a contradiction we assume that no vertex-minor of G is isomorphic to the $n \times n$ comparability grid.

By Ramsey's Theorem, there is a subsequence $(v_1, v_2, \ldots, v_{n^2})$ of $(h_1, h_2, \ldots, h_{n_1})$ such that one of the following holds:

(i) For each $1 \leq i < j \leq n^2$, the pair $(W^{\mathcal{C}}_{v_i}, W^{\mathcal{C}}_{v_j})$ is a coupled matching.

(ii) For each $1 \leq i < j \leq n^2$, the pair $(W^{\mathcal{C}}_{v_i}, W^{\mathcal{C}}_{v_j})$ is either an up-coupled half graph or the complement of a down-coupled half graph.

(iii) For each $1 \leq i < j \leq n^2$, the pair $(W^{\mathcal{C}}_{v_i}, W^{\mathcal{C}}_{v_j})$ is either a down-coupled half graph or the complement of an up-coupled half graph.

By possibly reversing the order of the sequence $(v_1, v_2, \ldots, v_{n^2})$ we may assume that we are not in case (iii). However, Lemma 2.6.3 precludes case (i) and Lemma 2.6.6 precludes case (ii). □

It remains to consider constellations whose associated graphs are paths. We say that a graph is an *ordered path* if the graph is a path and the order of the vertices on the path agrees with the ordering of the vertices of the graph; thus every path is isomorphic to an ordered path.

Lemma 2.6.8 (Path constellations)**.** *There are functions $n_{2.6.8} : \mathbb{Z} \to \mathbb{Z}$ and $k_{2.6.8} : \mathbb{Z} \to \mathbb{Z}$ such that, for any positive integer n, if $\mathcal{C}$ is an $(n_{2.6.8}(n), k_{2.6.8}(n))$-constellation in a graph G such that $K(\mathcal{C})$ is a path, then G has a vertex-minor isomorphic to the $n \times n$ comparability grid.*

Proof. For a positive integer n we define

$$\begin{aligned} m &:= n^2, \\ k_3 &:= m \cdot 2^{m-1}, \\ k_2 &:= k_3 + m - 1, \\ k_1 &:= k_2 + m - 1, \\ n_{2.6.8}(n) &:= (n^2 - 1)m, \text{ and} \\ k_{2.6.8}(n) &:= \max\left(k_1, \binom{n^2}{2}\right). \end{aligned}$$

For convenience we also define $n_0 := n_{2.6.8}(n)$ and $k_0 := k_{2.6.8}(n)$. Let $\mathcal{C}$ be an (n_0, k_0)-constellation in a graph G such that $K(\mathcal{C})$ is an ordered path on vertices $(h_1, \ldots, h_{n_0})$. Towards a contradiction we may assume that no vertex-minor of G is isomorphic to the $n \times n$ comparability grid.

Claim 2.6.8.1. *There is a graph G_1 that is locally equivalent to G and has an (m, k_1)-constellation $\mathcal{C}_1$ such that $K(\mathcal{C}_1)$ is an ordered path with vertices $(v_1, \ldots, v_m)$ and, for each $i \in \{1, \ldots, m-1\}$ the pair $(W^1_{v_i}, W^1_{v_{i+1}})$ is a coupled half graph.*

Proof. Let X denote the set of all $i \in \{1, \ldots, n_0 - 1\}$ such that $(W^{\mathcal{C}}_{h_i}, W^{\mathcal{C}}_{h_{i+1}})$ is a coupled matching. Let $(v_1, \ldots, v_t)$ be the restriction of the sequence $(h_1, \ldots, h_{n_0})$ to the elements $\{h_j : j \in \{1, \ldots, n_0\} - X\}$. By Lemma 2.6.3, the set X cannot contain $n^2 - 1$ consecutive integers and hence $t \geq m$. Let $H_1 := \{v_1, \ldots, v_m\}$, let P_1 be the ordered path on $(v_1, \ldots, v_m)$, let $\mathcal{C}_1 := (H_1, (W^{\mathcal{C}}_{v_1}, \ldots, W^{\mathcal{C}}_{v_m}), P_1)$, and let G_1 be the graph obtained from G by locally complementing on each of the vertices in $(W^{\mathcal{C}}_{h_i} : i \in X)$. It is routine to verify that the pair $(G_1, \mathcal{C}_1)$ satisfies the conclusion of the claim. □

Suppose that $A = (a_1, \ldots, a_l, a_{l+1})$ and $B = (b_1, \ldots, b_l, b_{l+1})$ are disjoint sets in a graph and (A, B) is a coupled half graph. If (A, B) is the complement of a down-coupled half graph then $((a_1, \ldots, a_l), (b_2, \ldots, b_{l+1}))$ is an up-coupled half graph, while, if (A, B) is a down-coupled half graph, then $((a_1, \ldots, a_l), (b_2, \ldots, b_{l+1}))$ is the complement of an up-coupled half graph. Starting with the first elements of $W^1_{v_1}$ and then choosing elements appropriately from each of $W^1_{v_2}, \ldots, W^1_{v_m}$ in turn we obtain the following result.

Claim 2.6.8.2. *There is an (m, k_2)-constellation $\mathcal{C}_2$ in G_1 such that $K(\mathcal{C}_2)$ is an ordered path on vertices $(v_1, \ldots, v_m)$ and, for each $i \in \{1, \ldots, m-1\}$ the pair $(W^2_{v_i}, W^2_{v_{i+1}})$ is an up-coupled half graph or the complement of an up-coupled half graph.*

By pivoting we can further reduce to the case where all pairs are up-coupled half graphs.

Claim 2.6.8.3. *There is a graph G_3 that is obtained from G_1 by pivoting and has an (m, k_3)-constellation $\mathcal{C}_3$ such that $K(\mathcal{C}_3)$ is an ordered path on vertices $(u_1, u_2, \ldots, u_m)$ and, for each $i \in \{1, \ldots, m-1\}$ the pair $(W^3_{u_i}, W^3_{u_{i+1}})$ is an up-coupled half graph.*

Proof. We will prove by induction on $m - t$, where $1 \leq t \leq m$, that if a graph G contains an $(m, k_3 + m - t)$-constellation $\mathcal{C}_2$ such that $K(\mathcal{C}_2)$ is an ordered path on vertices $(w_1, w_2, \ldots, w_m)$ and for each $i \in \{1, \ldots, t-1\}$, the pair $(W^2_{w_i}, W^2_{w_{i+1}})$ is an up-coupled half graph, and for each $i \in \{t, \ldots, m-1\}$, the pair $(W^2_{w_i}, W^2_{w_{i+1}})$ is an up-coupled half graph or the complement of an up-coupled half graph, then there is a graph G_3 that is obtained from G by pivoting and has an (m, k_3)-constellation $\mathcal{C}_3$ as in the claim. The case where $t = 1$ implies the claim since $k_2 = k_3 + m - 1$. The base case where $t = m$ holds by deleting excess vertices from each set $W^2_{w_i}$.

Now we may assume that $t < m$. We may also assume that the pair $(W^2_{w_t}, W^2_{w_{t+1}})$ is the complement of an up-coupled half graph, as otherwise we may delete one vertex from each set $W^2_{w_i}$ and apply induction. Let w be the first vertex in $W^2_{w_{t+1}}$. Let $G_3 = G \times ww_{t+1}$, let $H_3 = (H(\mathcal{C}_2) - \{w_{t+1}\}) \cup \{w\}$, and let K_3 be the graph obtained from $K(\mathcal{C}_2)$ by relabeling w_{t+1} to w. Let W^3_w be the set obtained from $W^2_{w_{t+1}}$ by deleting w, and, for each $h \in H(\mathcal{C}_2) - \{w_{t+1}\}$, let W^3_h be the set obtained from W^2_h by deleting its first vertex. Finally, let $\mathcal{C}_3 = (H_3, (W^3_h : h \in H_3), K_3)$.

Consider the neighbours of w and w_{t+1} in $G[V(\mathcal{C}_3) \cup \{w_{t+1}\}]$. The neighbourhood of w_{t+1} is exactly $W^2_{w_{t+1}}$. The vertex w is complete to $W^3_{w_t}$ and either complete or anticomplete to $W^3_{w_{t+2}}$, if $t + 2 \leq m$. These are the only neighbours of w other than w_{t+1}. Thus $\mathcal{C}_3$ is an $(m, k_3 + m - t - 1)$-constellation in G_3 so that all pairs are coupled in the same way as in G, except for $(W^3_{w_t}, W^3_w)$, which is an up-coupled half graph in G_3, and $(W^3_w, W^3_{w_{t+2}})$, which may be complemented. The claim follows by the induction hypothesis. $\square$

For each $s \in \{1, \ldots, m\}$, we let L_s denote the graph with vertex set $\{u_1, \ldots, u_m\}$ and edge set

$$\{u_iu_j : 1 \leq i < j \leq s\} \cup \{u_su_{s+1}, u_{s+1}u_{s+2}, \ldots, u_{m-1}u_m\}.$$

Thus L_1 is a path and L_m is a complete graph. For each $s \in \{1, \ldots, m\}$ we let $d_s := m2^{m-s}$; thus $d_1 = k_3$ and $d_m = m = n^2$.

Claim 2.6.8.4. *For each $s \in \{1, \ldots, m\}$, there is a graph G'_s that is locally equivalent to G and has disjoint d_s-vertex cocliques $(X^s_1, \ldots, X^s_m)$ such that*

(i) for each $i \in \{1, \ldots, m\}$, $X^s_i \subseteq W^3_{u_i}$,

(ii) for $1 \leq i < j \leq m$, the pair (X^s_i, X^s_j) is an up-coupled half graph if $u_iu_j \in E(L_s)$ and is anticomplete otherwise, and

(iii) for each $i \in \{s+1, \ldots, m\}$, the vertex u_i is complete to X^s_i and anticomplete to each of $X^s_1, \ldots, X^s_{i-1}$ and to each $X^s_{i+1} \cup \{u_{i+1}\}, \ldots, X^s_m \cup \{u_m\}$.

Proof. The proof is by induction on s; when $s = 1$ the conclusion is satisfied by $G'_1 := G_3$ and $X^1_i = W^3_{u_i}$ for each $i \in \{1, \ldots, m\}$. For some $s \in \{2, \ldots, m\}$ suppose that there exist G'_{s-1} and $(X^{s-1}_1, \ldots, X^{s-1}_m)$ as claimed; we will determine G'_s and $(X^s_1, \ldots, X^s_m)$.

We let G'_s be the graph obtained from G'_{s-1} by locally complementing on each vertex in $X^{s-1}_{s-1} \cup \{u_s\}$. Suppose that, for each $i \in \{1, \ldots, m\}$, we have $X^{s-1}_i = (x^i_1, \ldots, x^i_{d_{s-1}})$, and let $X^s_i := (x^i_1, x^i_3, \ldots, x^i_{(2d_s)-1})$. We claim that G'_s and $(X^s_1, \ldots, X^s_m)$ satisfy the result; this follows from the following observations about adjacencies in G'_{s-1}:

- for each $i, j \in \{1, 3, \ldots, d_{s-1} - 1\}$ and each $a, b \in \{1, \ldots, s-2\}$, the vertices x^a_i and x^b_j have an even number of common neighbours in $X^{s-1}_{s-1} \cup \{u_s\}$,
- for each $i, j \in \{1, 3, \ldots, d_{s-1} - 1\}$ and each $a, b \in \{s, s+1, \ldots, m\}$, the vertices x^a_i and x^b_j have an even number of common neighbours in $X^{s-1}_{s-1} \cup \{u_s\}$, and
- for each $i, j \in \{1, 3, \ldots, d_{s-1} - 1\}$ and each $a \in \{1, \ldots, s-2\}$, the vertices x^a_i and x^s_j have an odd number of common neighbours in $X^{s-1}_{s-1} \cup \{u_s\}$ if and only if $j \geq i$.

□

We obtain the final contradiction to Lemma 2.6.8 by applying Lemma 2.6.6 to G'_m and $(X^m_1, \ldots, X^m_m)$. □

We can now combine the above results to prove our main result, Theorem 2.1.1, which we restate here for convenience.

Theorem 2.1.1. *There is a function $f : \mathbb{Z} \to \mathbb{Z}$ so that for any positive integer n, every graph of rank-width at least $f(n)$ has a vertex-minor isomorphic to the $n \times n$ comparability grid.*

Proof. For a positive integer n we define

$$\begin{aligned} k_1 &:= \max\left(n^2 + 2, k_{2.6.7}(n), k_{2.6.8}(n)\right), \\ n_1 &:= \max\left(\binom{n^2}{2} + 1, n_{2.6.7}(n), n_{2.6.8}(n)\right), \text{ and} \\ f(n) &:= r_{2.5.6}(n_{2.6.1}(n_1), 0, k_1). \end{aligned}$$

Let G be a graph with rank-width at least $f(n)$. By Lemmas 2.5.6 and 2.6.1, there is a graph G_1, equivalent to G up to local complementation and reordering vertices, that contains an (n_1, k_1)-constellation $\mathcal{C}$ such that $K(\mathcal{C})$ is either a star, a clique, or a path. Now the result follows by Lemmas 2.6.4, 2.6.7, and 2.6.8. □

Chapter 3

The Local Structure Theorem

3.1 Representing a circle graph and its "attachments"

The rest of this thesis is dedicated to proving the Local Structure Theorem. Recall that, informally, the theorem says that for any graph in a proper vertex-minor-closed class with a prime circle graph containing a comparability grid, the rest of the graph "almost attaches" to the circle graph in a way that is "mostly compatible" with the comparability grid. In this chapter we will formalize the theorem and discuss our proof approach.

More precisely, the Local Structure Theorem is about a graph G with sets $B \subseteq C \subseteq V(G)$ so that B induces a comparability grid and C induces a prime circle graph. (In light of the Grid Theorem for Vertex-Minors, it is enough to prove the Structural Conjecture for graphs with a large comparability grid H as a vertex-minor. We might as well locally complement to obtain H as an induced subgraph; then our overall approach is to "grow a circle graph around H".)

Kotzig [83] and Bouchet [10] showed how to represent the edges/non-edges with an end in C by signatures on a connected 4-regular graph. This beautiful representation is called the "labelled tour graph", and is invariant under locally complementing at vertices in C. Our proof approach is based on this representation; moreover, the labelled tour graph helps us measure "how compatible the attachments are with B".

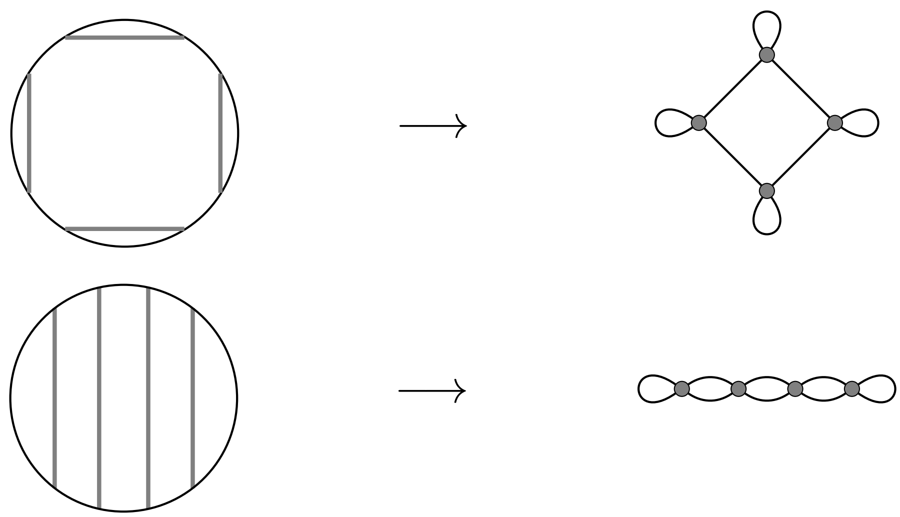

Figure 3.1: Two chord diagrams with the same circle graph (left) and the corresponding tour graphs (right).

Representing a circle graph by its tour graph

Let G be a circle graph, and fix a chord diagram $\mathcal{C}$ for G so that no two chords share an end (such a chord diagram exists). Then $\mathcal{C}$ "breaks up" the unit circle into $2|\mathcal{C}|$ arcs. That is, there are $2|\mathcal{C}|$ arcs I of the unit circle so that

(i) each end of I is the end of a chord in $\mathcal{C}$, and

(ii) no interior point of I is the end of a chord in $\mathcal{C}$.

A *tour graph of* G has vertex-set $V(G)$; it has an edge for each such arc I, and that edge is incident to the vertices whose chords intersect I. Tour graphs are allowed to have loops and multiple edges, and a single circle graph can have more than one tour graph; see Figure 3.1. The circle becomes an Eulerian circuit of the tour graph. Conversely, a chord diagram can be uniquely recovered from any connected 4-regular graph with a specified Eulerian circuit.

Kotzig [83] proved the following wonderful theorem.

Theorem 3.1.1 (Kotzig [83])**.** *Two circle graphs have a tour graph in common if and only if they are locally equivalent, in which case they have all of the same tour graphs.*

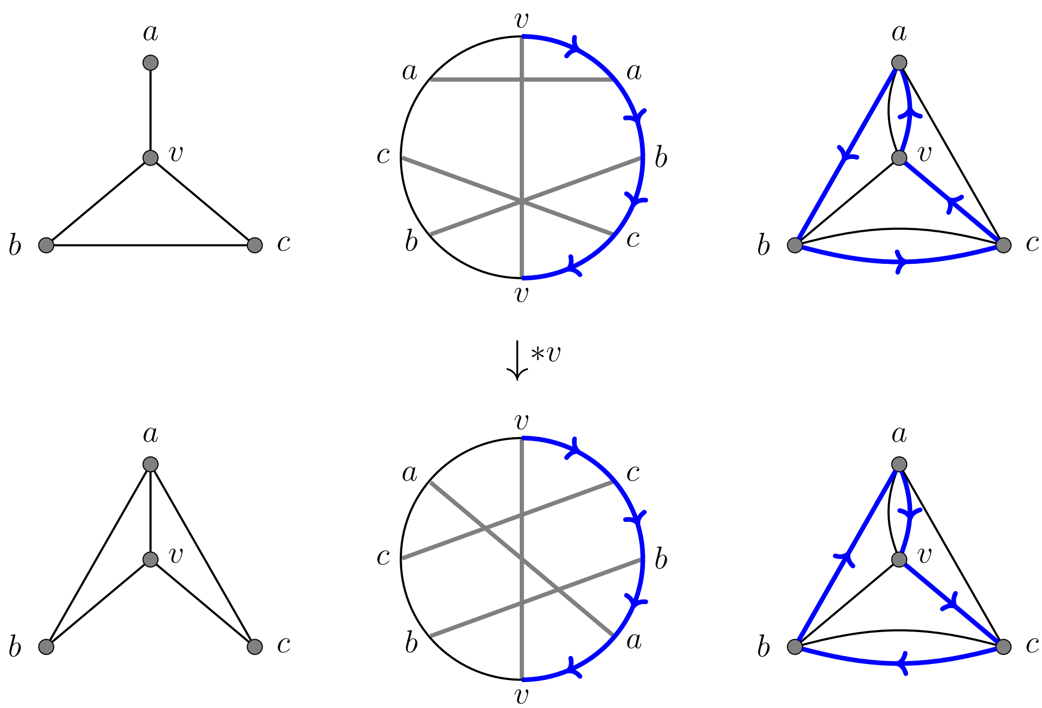

Figure 3.2: A circle graph (left), chord diagram (middle), and tour graph (right), before and after locally complementing at v.

So each local equivalence classes of circle graphs can be succinctly described by a collection of connected 4-regular graphs.

To give some intuition for the theorem, consider locally complementing at a vertex v in a circle graph with a fixed chord diagram. In the chord diagram, the corresponding operation is to flip one of the two arcs of the circle that has the same ends as the chord v (in Figure 3.2, we flip the bold blue arc). In the tour graph, the circle corresponds to an Eulerian circuit, and the two arcs of the circle with the same ends as v correspond to two proper subcircuits which begin and end at v. The corresponding operation in the tour graph is to reverse one of those subcircuits (in Figure 3.2, we reverse the bold blue subcircuit).

Throughout the rest of the thesis, we frequently work with Eulerian graphs; a graph is *Eulerian* if it is connected and every vertex has even degree (or, equivalently, if it has an Eulerian circuit). A *trail* is a walk with no repeated edges, and a *circuit* is a closed trail (that is, a trail which begins and ends at the same vertex). The *ends* of a trail are the vertices (or vertex, if the trail is also a circuit) at which it begins and ends.

In the last paragraph, given an Eulerian circuit of the tour graph, we were interested in the "two proper subcircuits" with v as their end. More formally, given a 4-regular graph with a vertex v and an Eulerian circuit $\hat{T}$, the *circuits of v in $\hat{T}$* are the two circuits C_1 and C_2 so that

(i) both C_1 and C_2 have at least one edge and have v as their end, and

(ii) $\hat{T}$ can be cyclically reordered so that $\hat{T} = C_1, C_2$ (as a sequence of edges).

These two circuits C_1 and C_2 are well-defined.

Representing an additional vertex by its signature

Let G be a graph with a vertex x so that $G - x$ is a circle graph. We will see that the neighbourhood of x can be represented by a certain "decoration" on the tour graph called a signature. In this subsection we give an informal description to motivate the representation.

So, let $\hat{G}$ be a tour graph of $G - x$ with an Eulerian circuit $\hat{T}$ so that the corresponding chord diagram $\mathcal{C}$ represents $G - x$. (Recall that a chord diagram can be uniquely recovered from $\hat{G}$ and $\hat{T}$; this was discussed when tour graphs were first defined.) If we can add x as a chord to $\mathcal{C}$, then we can specify the ends of that chord by two edges in the tour graph (see Figure 3.3). Furthermore, a vertex $v \neq x$ is adjacent to x in G if and only if each of the circuits of v in $\hat{T}$ contains one of those two edges (again see Figure 3.3). In general, we can represent the neighbourhood of x in G by a "hyperchord" of the chord diagram: that is, by a specified set of edges in the tour graph. This set of edges will be the "signature of x".

First, for each neighbour v of x, choose an arbitrary end of the chord of v, and add it to a set of points P. Then include an edge in the signature if its corresponding arc (on the circle) has exactly one end in P; see Figure 3.4. In the depicted example, we do not end up with the same set of edges from Figure 3.3. However, a vertex $v \neq x$ is still adjacent to x in G if and only if each of the circuits of v in $\hat{T}$ contains an odd number of edges in the signature. (The signature has an even number of edges, so the "parity" of the two circuits is the same.) To see that the neighbourhood of x in G is indeed as we have described, imagine walking around the circle, beginning anywhere and only going around once. Each time a point in P is passed, "switch" the two "incident" edges in/out of the signature. In terms of the "parity" of the circuits of a vertex $v \neq x$ in $\hat{T}$, only the vertex v whose chord has that point as an end is affected.

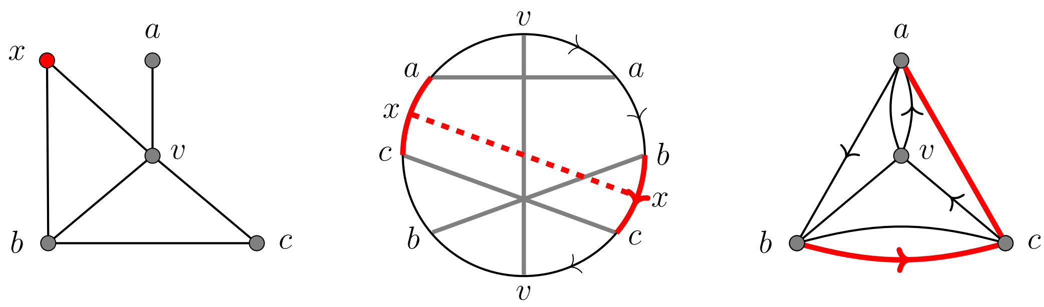

Figure 3.3: A circle graph (left), chord diagram (middle), and tour graph (right); the chord of x is dashed red and the two corresponding edges in the tour graph are bold red. The arrows represent one of the circuits of v in $\hat{T}$.

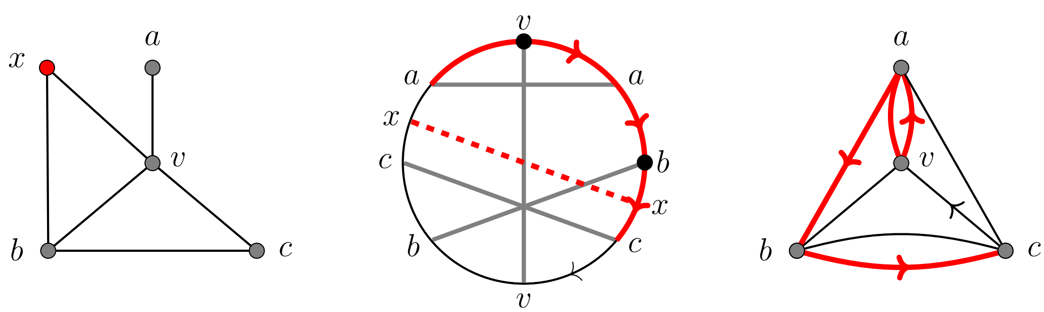

Figure 3.4: A circle graph (left), chord diagram (middle), and tour graph (right); the points in P are black points on the circle.

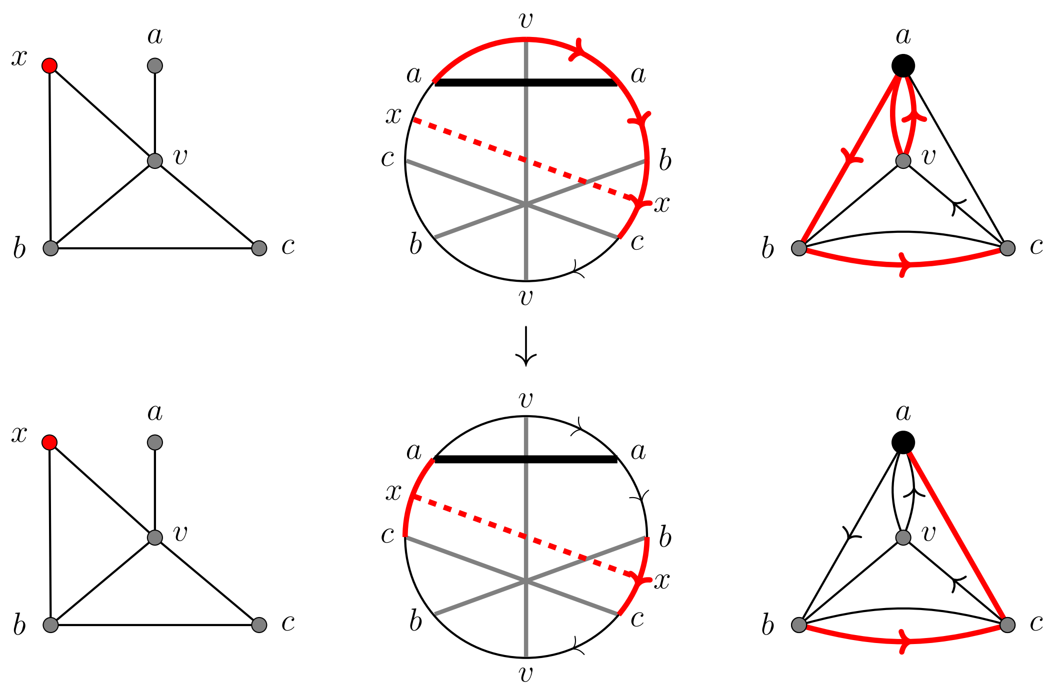

Figure 3.5: A circle graph (left), chord diagram (middle), and tour graph (right), before and after "switching" the non-loop edges incident to a in/out of the specified set.

Similarly, there is another operation which does not affect the "parity of circuits" at all. For any vertex a of the tour graph, we can "switch" all non-loop edges incident to a in/out of the specified set; see Figure 3.5. This operation will be called "shifting at a"; shifting at a neighbour of x is like choosing the other end of its chord to add to P. Bouchet [10] proved that G itself is a circle graph if and only if there exist $\hat{G}$ and $\hat{T}$ so that, after performing some shiftings, there are at most two edges in the signature. We will give the formal statement of this result in Lemma 3.3.2.

3.2 Labelled tour graphs

When we have a graph that contains a fixed circle graph, each vertex outside of the circle graph will correspond to a signature on the tour graph. In this section we give the formal definitions regarding this representation, which we call the "labelled tour graph". We also discuss when this representation is unique.

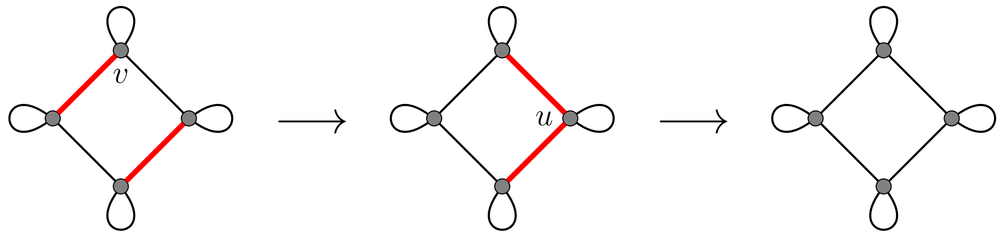

Figure 3.6: Shifting a signature at a vertex v and then a vertex u; the non-zero edges are bold red.

Labelled graphs

A *signature* on a graph is a function from its edge-set to the binary field, which we denote by $\mathbb{F}_2$. This is non-standard; as discussed in the last subsection, a signature is typically specified by the set of edges which are sent to 1. However, this formulation as a function is much more convenient in our case because we have many signatures.

The *weight* of an edge is the corresponding element in $\mathbb{F}_2$, and the *weight* of a circuit is the sum (over $\mathbb{F}_2$) of the weights of its edges. An edge or circuit is *zero* or *non-zero* depending on its weight. Signatures are used to specify which circuits of a graph are zero/non-zero; so there is an equivalence relation on signatures as follows.

First, *shifting at* a vertex means to add 1 to the weight of each incident non-loop edge (see Figure 3.6). A *shifting* of a signature λ is any signature that can be obtained from λ by performing a sequence of shiftings at vertices. Equivalently, a shifting is obtained from λ by adding 1 to the weight of each edge in a cut. Thus shifting is an equivalence relation that does not change the weight of any circuit. There is also a converse; Harary [70] proved that if each circuit of a graph has the same weight according to two signatures, then the signatures are shiftings of each other.

We need to keep track of many signatures at once. So, for a finite set V, we write $\mathbb{F}_2^V$ for the binary vector space where each vector is indexed by the elements of V. (If V is empty, then $\mathbb{F}_2^V$ just has the element 0. For $k \in \mathbb{N}$, we also write $\mathbb{F}_2^k$ for the binary vector space of dimension k.) From now on, we will not specify that the set V is finite, because this will always be the case.

An $\mathbb{F}_2^V$*-labelled graph* is a tuple (G, λ) so that G is a graph (which may have loops and multiple edges) and λ is a function from the edge-set of G to $\mathbb{F}_2^V$. We call λ an $\mathbb{F}_2^V$*-labelling on* G. For an edge e, the *weight* of e is $\lambda(e)$. Likewise, the *weight* of a trail T, which we denote by $\lambda(T)$, is the sum of the weights of its edges. If V consists of a single element, then (G, λ) is a *signed graph*. When we do not want to specify the set V, we call (G, λ) a *labelled graph* and λ a *labelling*.

For each $A \subseteq V$, we write λ_A for the $\mathbb{F}_2^A$-labelling that is obtained from λ by restricting the weights to $\mathbb{F}_2^A$. (If A is empty, then λ_A labels every element 0.) When we want to be more specific about the weight of an edge or trail, we talk about its λ_A-*weight* or, equivalently, its $\mathbb{F}_2^A$-*weight*. For a trail T, we write $\lambda_A(T)$ for its λ_A-*weight*, which is the sum of the λ_A-weights of its edges. For an element $v \in V$, we write λ_v for $\lambda_{\{v\}}$ and $\mathbb{F}_2^v$ for $\mathbb{F}_2^{\{v\}}$ for convenience. Thus λ is fully specified by the collection of signatures $(\lambda_v : v \in V)$; we call those the *signatures of* λ. Finally, a *shifting of* λ is any labelling that can be obtained from λ by replacing each of its signatures λ_v by a shifting of λ_v.

Sometimes we also need to combine labellings. Given a graph G with an $\mathbb{F}_2^V$-labelling λ and an $\mathbb{F}_2^W$-labelling μ, for disjoint sets V and W, we write $\lambda \times \mu$ for the $\mathbb{F}_2^{V \cup W}$-labelling where each edge e has weight $(\lambda(e), \mu(e))$. Likewise we write $\mathbb{F}_2^V \times \mathbb{F}_2^W$ for $\mathbb{F}_2^{V \cup W}$.

Labelled tour graphs

It is convenient to have a name for a graph that contains a fixed circle graph. So a *circle-structure* is a tuple (G, C) so that G is a graph and C is a set of vertices of G which induces a circle graph. We call G *its graph* and $G[C]$ *its circle graph*. Two circle-structures $(\widetilde{G}, \widetilde{C})$ and (G, C) are *locally equivalent* if $V(\widetilde{G}) = V(G)$, $\widetilde{C} = C$, and $\widetilde{G}$ can be obtained from G by locally complementing at vertices in C.

We will see how to represent the local equivalence class of a circle-structure (G, C), up to changing edges/non-edges with neither end in C. Surprisingly, it is oftentimes possible to find vertex-minors of G without knowing about such edges/non-edges; we show how to do so in Section 3.5. For some intuition, suppose that there is a vertex $v \in C$ with exactly two neighbours x and y, both of which are outside of C. Then locally complementing at v "exchanges the edge/non-edge" between x and y; so it does not matter whether there was an edge/non-edge originally.

Now, let (G, C) be a circle-structure, and fix a chord diagram for $G[C]$ so that no two chords share an end. Let $\hat{G}$ and $\hat{T}$ be the corresponding tour graph and Eulerian circuit, respectively. A *labelled tour graph of* (G, C) is an $\mathbb{F}_2^{V(G)-C}$-labelled graph $(\hat{G}, \hat{\lambda})$ so that the signatures of $\hat{\lambda}$ are as follows.

For each vertex $x \in V(G) - C$, the *signature of* x *on* $\hat{G}$ *with respect to* $\hat{T}$ is the unique signature, up to shifting, so that $\hat{T}$ has weight zero and x is adjacent to a vertex $v \in C$ if and

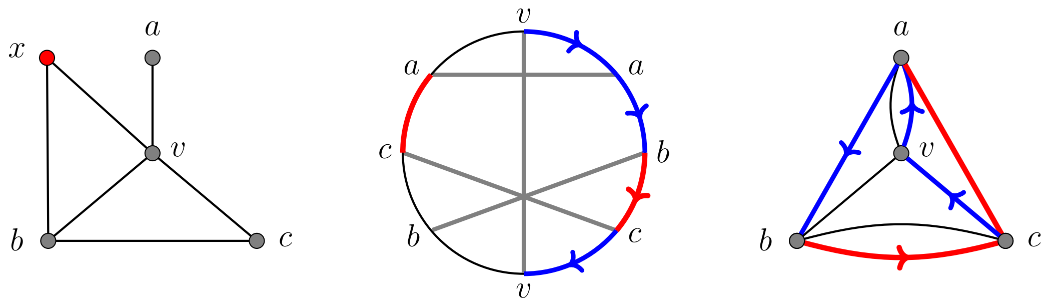

Figure 3.7: A circle graph (left), chord diagram (middle), and tour graph (right); the non-zero edges (according to the signature of x) are bold red.

only if each of the two circuits of v in $\hat{T}$ is non-zero (see Figure 3.7). Such a signature exists because we could begin with the uniformly zero signature and then, for each neighbour v of x, add 1 to the two edges "incident" to one of the ends of the chord v (the choice of end does not matter up to shifting at v). The signature is unique (given $\hat{G}$ and $\hat{T}$) because the specified circuits generate the cycle space; see the paper by Bouchet [10], who introduced this definition.

Bouchet [10] generalized the theorem of Kotzig [83] to labelled tour graphs as follows.

Theorem 3.2.1 (Bouchet [10])**.** *Two circle-structures (G, C) and $(\widetilde{G}, \widetilde{G})$ have a labelled tour graph in common if and only if $(\widetilde{G}, \widetilde{G})$ can be obtained from a circle-structure that is locally equivalent to (G, C) by changing edges/non-edges with neither end in C. Moreover, if they have one labelled tour graph in common, then they have all of the same labelled tour graphs.*

Bouchet worked with a single vertex outside of the circle graph, but Theorem 3.2.1 is equivalent to Bouchet's work. See Lee's thesis [86] for a nice explanation of this theorem.

We use labelled tour graphs to succinctly describe the edges/non-edges with an end in the circle graph. We could also go in the other direction; the obvious necessary conditions are also sufficient for a labelled graph to be a labelled tour graph. Formally, if V is a set and $(\hat{G}, \hat{\lambda})$ is a connected 4-regular $\mathbb{F}_2^V$-labelled graph so that V is disjoint from $V(\hat{G})$ and the sum (over $\mathbb{F}_2^V$) of the weights of the edges of $\hat{G}$ is zero, then $(\hat{G}, \hat{\lambda})$ is a labelled tour graph of a circle-structure.

Prime graphs

There is a minor connectivity-like condition which guarantees that a circle-structure has a unique labelled tour graph. A circle-structure is *prime* if its circle graph is prime, and a graph is *prime* if it is 2-rank-connected (recall the definition from Section 1.5; a graph is 2-*rank-connected* if it has at least four vertices, is connected, and its only cuts of cut-rank 1 have one vertex on one side). The name "prime" comes from Cunningham [29], who introduced primeness and showed how to "decompose" a graph into its "prime parts".

Theorem 3.2.2 (Bouchet [5] and Gabor, Supowit, and Hsu [52]). *Any prime circle-structure has a unique labelled tour graph, up to renaming edges in a way that preserves the labelling.*

So from now on, we can talk about *the labelled tour graph* and *the labelling* or *the signatures* of a prime circle-structure.

Bouchet [8] also gave the corresponding connectivity condition on tour graphs. A 4-regular graph is *internally* 6-*edge-connected* if it is simple and loopless, and any cut with at most four edges has at most one vertex on one side. (Thus any internally 6-edge-connected graph is 4-edge-connected.)

Theorem 3.2.3 (Bouchet [5, 8]). *The tour graph of a prime circle graph is internally* 6-*edge-connected, and each internally* 6-*edge-connected* 4-*regular graph is the tour graph of a prime circle graph.*

3.3 Stating the theorem

Here is the formal statement of the Local Structure Theorem; we will give the remaining definitions afterwards. The theorem says that any prime circle-structure whose graph has a forbidden vertex-minor can be "perturbed" so that each signature is "small relative to" a fixed comparability grid.

The Local Structure Theorem (Geelen, McCarty, Wollan). *For any $t \in \mathbb{N}$, there exists $\ell_t \in \mathbb{N}$ so that if (G, C) is a prime circle-structure whose graph does not have all $(t+1)$-vertex graphs as vertex-minors, up to isomorphism, and if B is a subset of C that induces a comparability grid, then there is a rank-t perturbation of (G, C) so that each signature is ℓ_t-small relative to B.*

So there are two definitions remaining: a "rank-t perturbation of a circle-structure" and when a signature is "ℓ_t-small relative to B". A perturbation of a circle-structure is not allowed to change its circle graph; so the signatures of the perturbed circle-structure are well-defined. Moreover, the rank of the perturbation in the theorem is, informally, optimal; we believe that this will be important for proving the Structural Conjecture.

Perturbations of circle-structures

Perturbations of circle-structures are different from perturbations of graphs; for the rest of the thesis we only work with perturbations of circle-structures. For $t \in \mathbb{N}$, a *rank-t perturbation of a circle-structure* (G, C) is a circle-structure $(\widetilde{G}, C)$ so that $\widetilde{G}[C] = G[C]$, $\widetilde{G} - C = G - C$, and the submatrix of $Adj_{\widetilde{G}} + Adj_G$ with rows $V(G) - C$ and columns C has rank at most t. (So using the prior definition on graphs, $\widetilde{G}$ is a rank-$2t$ perturbation of G; it is natural to give a different definition for circle-structures because we would like to maintain the circle graph, and we do not care about edges/non-edges with neither end in C.) Thus each perturbation of a prime circle-structure is also prime.

We need a better understanding of how a perturbation changes the signatures. Consider the row space of the submatrix of $Adj_{\widetilde{G}} + Adj_G$ with rows $V(G) - C$ and columns C. Informally, this row space specifies at most t signatures on the tour graph, and each new signature is obtained from the old signature by adding some of those t signatures (over $\mathbb{F}_2$). This formulation is equivalent due to the following lemma which says that "summing two rows sums the corresponding signatures". Primeness is not necessary, but it makes the statement simpler.

Lemma 3.3.1. *If (G, C) is a prime circle-structure and $x, y, z \in V(G) - C$ are such that the neighbourhood of z in C is the symmetric difference of the neighbourhoods of x and y in C, then the signature of z is the sum (over $\mathbb{F}_2$) of the signatures of x and y.*

Proof. Let $(\hat{G}, \hat{\lambda})$ be the labelled tour graph of (G, C), and consider the Eulerian circuit $\hat{T}$ of $\hat{G}$ which yields the circle graph induced by C. A vertex $v \in C$ is adjacent to a vertex $w \in V(G) - C$ if and only if each of the two circuits of v in $\hat{T}$ has non-zero $\hat{\lambda}_w$-weight; denote this weight by $\hat{\lambda}^v_w$. So $\hat{\lambda}^v_w$ is the indicator function for whether or not w is adjacent to v. Thus $\hat{\lambda}^v_x + \hat{\lambda}^v_y = \hat{\lambda}^v_z$. One signature for z which satisfies this property is $\hat{\lambda}_x + \hat{\lambda}_y$. So the lemma follows from uniqueness of signatures. □

The size of a signature

The *size* of a signature λ is the minimum, over all its shiftings λ', of the number of edges with non-zero λ'-weight. Informally, the size measures how "compatible" a vertex is with a circle graph; a vertex is "fully compatible" if and only if its signature has size at most 2.

Lemma 3.3.2 (Bouchet [10])**.** *For any circle-structure (G, C) and vertex $x \in V(G) - C$, the set $C \cup \{x\}$ induces a circle graph if and only if there exists a labelled tour graph for (G, C) so that the signature of x has size at most 2.*

Intuitively, those two edges specify where to place the ends of the chord of x; recall Figure 3.3 and the related discussion.

There is a dual notion to the size of a signature: the maximum number of edge-disjoint non-zero circuits. Certainly this number is at most the size of a signature. In the other direction we have the following theorem (a *k-edge-cut* is a cut with precisely k edges).

Theorem 3.3.3 (Kawarabayashi and Kobayashi [77])**.** *For each $k \in \mathbb{N}$, there exists an integer f_k so that the size of any signature on a graph with no 3-edge-cut and no k pairwise edge-disjoint non-zero circuits is at most f_k.*

So in particular the theorem holds for graphs with an Eulerian circuit (because each cut has an even number of edges). This type of property, where there is an approximate duality between packing and covering (in this case, packing and covering edge-disjoint non-zero circuits), is called the Erdös-Pósa property after [49]. The condition on 3-edge-cuts is necessary; see the construction of Lovász and Schrijver in the paper by Reed [107].

There is also a version of Theorem 3.3.3 for hitting circuits (a circuit *hits* a set of vertices B if it contains an edge incident to a vertex in B). Kakimura, Kawarabayashi, and Kobayashi [76] proved that, informally, edge-disjoint non-zero hitting circuits have the Erdös-Pósa property on graphs with no 3-edge-cut. This motivates our main definition; informally, we will say that a signature is "small relative to B" if we can constructively show that there is no large collection of edge-disjoint non-zero circuits hitting B. In particular, it suffices to find a small cut where one side contains all of B and has few non-zero edges (possibly in a shifting).

Here is the formal definition. Let (G, C) be a prime circle-structure, and let $B \subseteq C$. For convenience, let $(\hat{G}, \hat{\lambda})$ denote the labelled tour graph of (G, C). Then, for a vertex $x \in V(G) - C$ and an integer ℓ, the signature $\hat{\lambda}_x$ is *ℓ-small relative to B* if there exists a set $X \subseteq C$ which contains B and a shifting $\hat{\lambda}'_x$ of $\hat{\lambda}_x$ so that *(i)* there are at most ℓ edges of $\hat{G}$ with exactly one end in X, and

(ii) there are at most ℓ edges of $\hat{G}$ with both ends in X that have non-zero $\hat{\lambda}'_x$-weight.

These conditions guarantee that there is no collection of more than $3\ell/2$ edge-disjoint non-zero circuits hitting B (because at most $\ell/2$ of the circuits can "use an edge in the cut", and at most ℓ can use an edge with both ends in X that has non-zero $\hat{\lambda}'_x$-weight). So, informally, a signature is "mostly compatible with B" if it is small relative to B.

3.4 Complete immersion minors

Labelled tour graphs represent local equivalence classes of circle-structures. By considering vertex-deletion, we will obtain a notion of containment that corresponds to "vertex-minors of circle-structures". First, an *induced substructure* of a circle-structure (G, C) is a circle-structure $(G[X], C \cap X)$, where $X \subseteq V(G)$. Then a *vertex-minor* of (G, C) is an induced substructure of a circle-structure that is locally equivalent to (G, C). This notion of containment is transitive.

Here is the informal intuition for the corresponding notion of containment in labelled tour graphs. To delete a vertex from outside of the circle graph, we just forget about its signature. To delete a vertex from inside the circle graph, we "split off" its four incident half-edges into two new edges according to the relevant Eulerian circuit (see Figure 3.8). The weight of each new edge e is obtained by summing the weights of the two edges which we "combined" to obtain e. There are three partitions of the four incident half-edges into two parts of size two; they correspond to the three ways of removing a vertex from Lemma 1.6.5 (delete, locally complement and delete, and pivot-delete). However, since we must split off edges with respect to an Eulerian circuit, we are not allowed to disconnect the tour graph.

Continuing this informal description, a "complete immersion minor" is any labelled graph that can be obtained by shifting and successively "splitting off vertices by combining their incident edges" and forgetting about signatures. In order to give the formal definition, however, it is more convenient to look at things another way. Each edge of the smaller graph comes from a trail in the larger graph; this trail must have the appropriate ends and appropriate weight. Moreover, the collection of trails must partition the edge-set of the larger graph. If instead the trails are edge-disjoint (but do not necessarily use every edge), then we obtain the standard definition of (weak) immersion minors; see Figure 3.9.

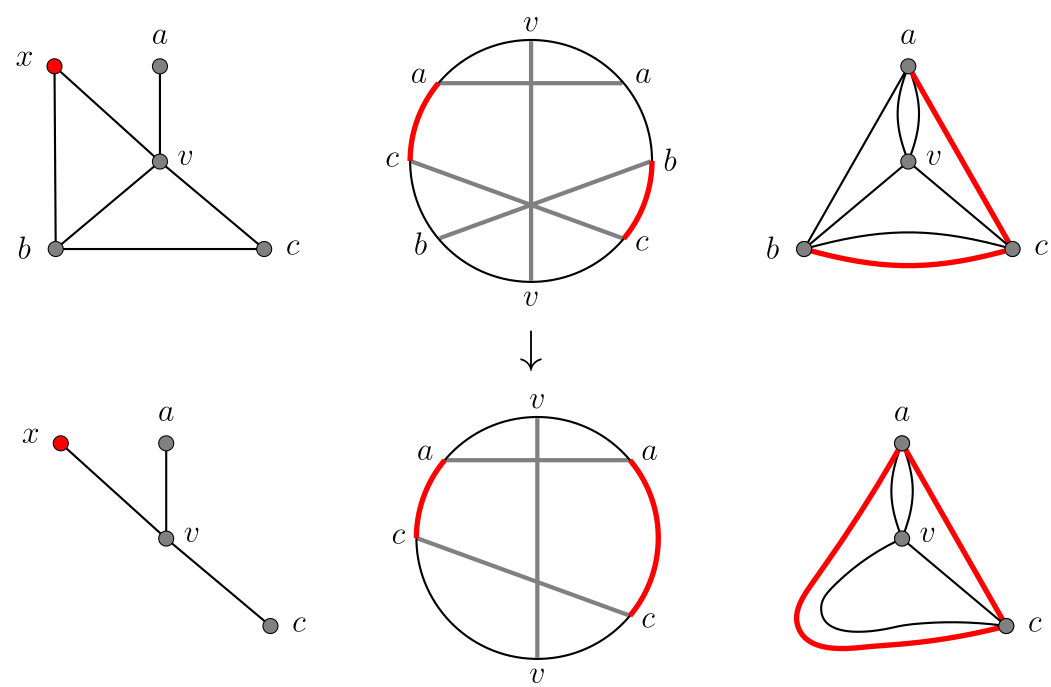

Figure 3.8: A circle graph (left), chord diagram (middle), and tour graph (right), before and after removing b.

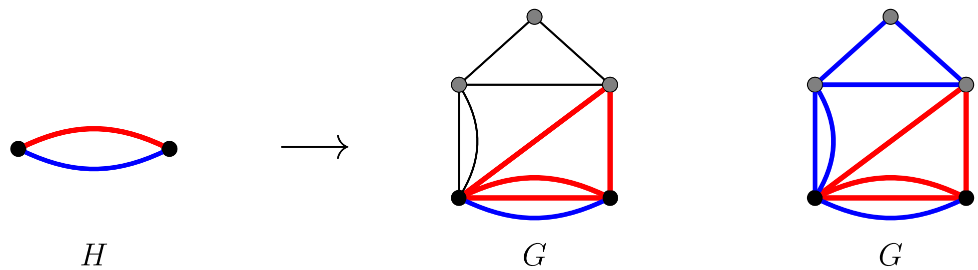

Figure 3.9: A depiction of an immersion minor and a complete immersion minor, where the colors/thickness denote the edges in the two trails in $\mathcal{T}$, and $W = V = \emptyset$.

Formally, let W and V be sets, and let (H, μ) and (G, λ) be an $\mathbb{F}_2^W$-labelled and $\mathbb{F}_2^V$-labelled graph, respectively. Then (H, μ) is an *immersion minor* of (G, λ) if $V(H) \subseteq V(G)$, $W \subseteq V$, and there exists a shifting λ' of λ and a collection $\mathcal{T} = (T_e : e \in E(H))$ of edge-disjoint trails of G such that for each $e \in E(H)$, the trail T_e has the same ends as e and satisfies $\lambda'_W(T_e) = \mu(e)$. (We allow H and G to have loops and multiple edges. So if e is a loop, then T_e is a circuit.) If additionally every edge of G is in a trail in $\mathcal{T}$, then (H, μ) is a *complete* immersion minor of (G, λ). If $\lambda' = \lambda$, then (H, μ) is a *shifting-free* (complete) immersion minor of (G, λ). We also sometimes say that (G, λ) *(completely) immerses* (H, μ). For (unlabelled) graphs H and G, we use these same definitions by implicitly referring to the corresponding $\mathbb{F}_2^{\emptyset}$-labelled graphs.

If (G, λ) (completely) immerses (H, μ) and μ' is a shifting of μ, then (G, λ) also (completely) immerses (H, μ'); we can simply perform the same sequence of shiftings in (G, λ). Furthermore, for unlabelled Eulerian graphs, immersion minors and complete immersion minors are equivalent.

Lemma 3.4.1. *If H and G are Eulerian graphs so that G immerses H, then G also completely immerses H.*

Proof. Let $\mathcal{T} = (T_e : e \in E(H))$ be a collection of trails which shows that G immerses H; choose $\mathcal{T}$ to use as many edges of G as possible. Each component of the graph obtained from G by removing all edges in a trail in $\mathcal{T}$ is Eulerian. Thus $\mathcal{T}$ must use every edge of G, as otherwise we could "add a circuit" onto a trail in $\mathcal{T}$. $\square$

This lemma does not necessarily hold for signed Eulerian graphs (G, λ) and (H, μ); informally, the weight of a trail could change when "adding on a circuit". For instance, if $\sum_{e \in E(H)} \mu(e) \neq \sum_{e \in E(G)} \lambda(e)$, then (H, μ) is not a complete immersion minor of (G, λ).

Complete immersion minors are the right notion of containment for labelled tour graphs.

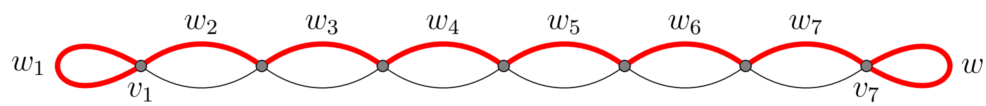

Figure 3.10: A depiction of the word graph of $W = w_1, \dots, w_7$, where w is the evaluation of W.

Theorem 3.4.2 (Bouchet [10] and Kotzig [83])**.** *The following hold for any circle-structure (G, C) any labelled tour graph $(\hat{G}, \hat{\lambda})$ of (G, C).*

(i) Any complete immersion minor of $(\hat{G}, \hat{\lambda})$ is a labelled tour graph of a vertex-minor of (G, C).

(ii) Any vertex-minor of (G, C) has at least one labelled tour graph that is a complete immersion minor of $(\hat{G}, \hat{\lambda})$.

This theorem motivates our overall approach to proving the Local Structure Theorem; from now on we mainly work with complete immersion minors.

3.5 Finding vertex-minors

We have seen how to find vertex-minors of a circle-structure (G, C); now we show how how to find vertex-minors of G. In particular, we show (in Lemma 3.5.3) that if any labelled tour graph of (G, C) completely immerses a particular type of labelled graph, then G has every t-vertex graph as a vertex-minor, up to isomorphism.

We construct labelled graphs corresponding to words. For a set A, a *word over* $\mathbb{F}_2^A$ is a sequence $W = w_1, \dots, w_n$ of elements of $\mathbb{F}_2^A$. The elements $w_1, \dots, w_n$ are *letters*. A *subword* of W is a word that can be obtained from W by deleting zero or more letters at its beginning and end; if we only delete letters at the end (respectively, beginning) then we obtain a *prefix* (respectively, *suffix*) of W. We particularly care about "evaluations" of prefixes; the *evaluation* of a word is the sum of its elements over $\mathbb{F}_2^A$ (where the empty word evaluates to 0). For a letter w_i of W, the *prefix value of w_i* is the evaluation of the corresponding prefix $w_1, \dots, w_i$. We will occasionally also be interested in *words over* $\mathbb{F}_2^k$ for $k \in \mathbb{N}$, which are defined in the same way.

Now, for a set A and a non-empty word $W = w_1, \dots, w_n$ over $\mathbb{F}_2^A$, the *word graph of W* (see Figure 3.10) is the $\mathbb{F}_2^A$-labelled graph with n vertices $v_1, \dots, v_n$ so that *(i)* there is a loop at v_1 of weight w_1,

(ii) for each $i \in \{1, \ldots, n-1\}$, there are two edges between v_i and v_{i+1}, one of weight zero and the other of weight w_{i+1},

(iii) there is a loop at v_n whose weight is the evaluation of W, and

(iv) there are no other edges.

The weights of the edges sum to zero and the graph is 4-regular and connected; so the word graph is the labelled tour graph of some circle-structure.

The following lemma says that those circle-structures can be easily understood; the neighbourhood of each vertex v_i is the "support" of the prefix value of w_i. The *support* of an element $w \in \mathbb{F}_2^A$ is the set of all $a \in A$ so that the entry of w at a is non-zero. So in particular, the following lemma says that $\{v_1, \ldots, v_n\}$ is a stable set in the circle-structure, and locally complementing at vertices in the circle graph (which has vertex-set $\{v_1, \ldots, v_n\}$) does not change the neighbourhood of any v_i.

Lemma 3.5.1. *For any set A and any non-empty word $W = w_1, \ldots, w_n$ over $\mathbb{F}_2^A$, a circle-structure (G, C) has the word graph of W as a labelled tour graph if and only if*

(i) $A = V(G) - C$,

(ii) $C = \{v_1, \ldots, v_n\}$ *is the vertex-set of the word graph, and*

(iii) for each $i \in \{1, \ldots, n\}$, the neighbourhood of v_i in G is precisely the support of the prefix value of w_i.

Proof. Fix an Eulerian circuit $\hat{T}$ of the word graph of W; it suffices to show that the corresponding circle-structure, which we denote by (G, C), satisfies condition *(iii)*. By considering the corresponding chord diagram, we can see that $\{v_1, \ldots, v_n\}$ is a stable set in G. So it just remains to consider when a vertex v_i, for $i \in \{1, \ldots, n\}$, is adjacent to a vertex $a \in A$ (in the graph G). Such vertices are adjacent if and only if each of the two circuits of v in $\hat{T}$ has non-zero weight according to the signature of a. The $\mathbb{F}_2^A$-weight of each of these two subcircuits is precisely the prefix value of w_i, so the lemma follows. □

Next we will show how to use Lemma 3.5.1 and Theorem 3.4.2 (on complete immersion minors) to find vertex-minors in the graph of a circle-structure. In light of Lemma 3.5.1, it is natural to consider words "with every possible prefix value". So for a set A, an *A-universal* word is a word over $\mathbb{F}_2^A$ so that every element of $\mathbb{F}_2^A$ is the prefix value of a letter. It straightforward to construct small such words.

Lemma 3.5.2. *For any set A, there is an A-universal word with $2^{|A|}$ letters.*

Proof. We begin with the empty word (which has no prefix values since it has no letters), and "add prefix values" one at a time. Suppose that the current word evaluates to α and we want to add an element $\alpha' \in \mathbb{F}_2^A$ as a prefix value. Then we add the element $\alpha+\alpha' \in \mathbb{F}_2^A$ at the end of the word. This process results in the desired word after adding $|\mathbb{F}_2^A| = 2^{|A|}$ letters. □

As one would hope, A-universal words let us find every graph with vertex-set A as a vertex-minor.

Lemma 3.5.3. *For any set A and any circle-structure (G, C) with at least one labelled tour graph that completely immerses the word graph of an A-universal word, the graph G has every graph with vertex-set A as a vertex-minor.*

Proof. By Theorem 3.4.2, there exists a vertex-minor $(\widetilde{G}, \widetilde{C})$ of (G, C) so that the word graph of an A-universal word is a labelled tour graph of $(\widetilde{G}, \widetilde{C})$. By Lemma 3.5.1 and A-universality, for each $N \subseteq A$, there exists a vertex in $\widetilde{C}$ whose neighbourhood, in $\widetilde{G}$, is precisely N.

Now consider an arbitrary graph with vertex-set A. There are some pairs of vertices x, y on which this graph differs from the subgraph of $\widetilde{G}$ induced on A (that is, where exactly one of the two graphs has an edge between x and y). We can correct this difference by locally complementing at the vertex in $\widetilde{C}$ with neighbourhood $\{x, y\}$. So the graph $\widetilde{G}$, and thus also the graph G, has every graph with vertex set A as a vertex-minor. □

This lemma motivates our overall approach to proving the Local Structure Theorem; we will work entirely in the labelled tour graph, and we will find the word graph of an A-universal word as a complete immersion minor.

Vertex-minors and the Growth Rate Theorem

While we will only need Lemma 3.5.3 in order to prove the Local Structure Theorem, we believe that far fewer prefix values are needed in order to force every t-vertex graph as a vertex-minor, up to isomorphism. In particular, notice that in the proof of Lemma 3.5.3, we only actually used $\binom{|A|+1}{2}$ of the prefix values. Informally, we believe that for every integer t, there exists $c \in \mathbb{R}$ so that $c|A|^2 + 1$ prefix values suffice.

In light of Lemma 3.5.1 about circle-structures and word graphs, we can state this conjecture in terms of graphs with a bipartition (A, C) so that C is a stable set containing many vertices with distinct neighbourhoods in A. Moreover, we conjecture that C does not need to be a stable set.

Conjecture 3.5.4. *For any $t \in \mathbb{Z}^+$, there exists $c \in \mathbb{R}$ so that any graph with a bipartition (A, C) so that there are more than $c|A|^2$ vertices in C with distinct neighbourhoods in A has every t-vertex graph as a vertex-minor, up to isomorphism.*

We believe that we can prove Conjecture 3.5.4 when C is a stable set. We omit the proof, but the key point is the following corollary of the Growth Rate Theorem for minor-closed classes of matroids, which was ultimately proven by Geelen, Kung, and Whittle [63]. For a matroid M, we write $E(M)$ for the set of elements of M and $r(M)$ for the rank of M.

Corollary 3.5.5 (Geelen, Kung, and Whittle [63])**.** *For any $t \in \mathbb{Z}^+$, there exists $c \in \mathbb{R}$ so that any simple binary matroid M with $|E(M)| \geq c(r(M))^2$ has every t-element binary matroid as a minor, up to isomorphism.*

The connection comes from Bouchet's [7] work relating pivot-minors of bipartite graphs to minors of binary matroids (recall Theorem 1.6.1). As long as one side is a stable set, pivot-minors still "behave appropriately between the two sides". In this manner, we believe that we can reduce Conjecture 3.5.4 to the "A-universal case" when C is a stable set. Again motivated by the Growth Rate Theorem, we also conjecture that the bound $c|A|^2$ in Conjecture 3.5.4 can be replaced by the bound $c(\omega|A|)$ for graphs G with clique number at most ω.

3.6 The proof approach

Our proof of the Local Structure Theorem relies on Lemma 3.5.3; to explain this connection, consider a prime circle-structure (G, C) whose graph does not have all $(t+1)$-vertex graphs as vertex-minors, up to isomorphism. By Lemma 3.5.3, there is no set $A \subseteq V(G) - C$ of size $t+1$ so that the labelled tour graph of (G, C) completely immerses the word graph of an A-universal word. So we "inductively grow" a set $A \subseteq V(G) - C$ so that the labelled tour graph of (G, C) does completely immerse such a word graph; we will eventually get stuck and no longer be able to grow A. At that point, we will find a rank-$|A|$ perturbation of the circle-structure which makes each signature small relative to the fixed subset of C which induces a comparability grid. This perturbation will be "defined by" the signatures of the vertices in A; we give the relevant definitions next.

Keeping track of the perturbation

Let (G, C) be a prime circle-structure. For $a, v \in V(G) - C$, a circle-structure $(\widetilde{G}, \widetilde{C})$ is *obtained from* (G, C) *by adding* a *to* v if $V(\widetilde{G}) = V(G)$, and $\widetilde{C} = C$, and the edge-set of $\widetilde{G}$ is obtained from the edge-set of G by exchanging edges/non-edges between v and any neighbour of a in C. So, equivalently, the submatrix of $Adj_{\widetilde{G}}$ with rows $V(G) - C$ and columns C is obtained from the submatrix of Adj_G with the same rows and columns by adding row a to row v, and all other incidences remain the same. Lemma 3.3.1 implies that the signature of v in $(\widetilde{G}, \widetilde{C})$ is the sum of the signatures of v and a in (G, C).

In general, we will specify a set of vertices that we are allowed to "add". So for a circle-structure (G, C) and a set $A \subseteq V(G) - C$, a circle-structure $(\widetilde{G}, \widetilde{C})$ is *obtained from* (G, C) *by perturbing to* A if $(\widetilde{G}, \widetilde{C})$ can be obtained from (G, C) by adding vertices in A to vertices in $V(G) - C$. Thus $(\widetilde{G}, \widetilde{C})$ is a rank-$|A|$ perturbation of (G, C). If we only add vertices in A to a single vertex $v \in V(G) - C$, then we say that $(\widetilde{G}, \widetilde{C})$ is *obtained from* (G, C) *by perturbing* v *to* A. This is how we will find a perturbation.

Growing the set

Let (G, C) be a prime circle-structure, let $B \subseteq C$ be a set which induces a comparability grid, and let $(\hat{G}, \hat{\lambda})$ denote the labelled tour graph of (G, C). We have shown how to keep track of a perturbation by perturbing to a set $A \subseteq V(G) - C$. Next we discuss how to build this set A. If A grows to size $t + 1$, then we will find every $(t + 1)$-vertex graph as vertex-minor, up to isomorphism. If instead we get stuck while $|A| \leq t$, then we will perturb to A to make every signature small relative to B.

In order to guarantee that this strategy will work, we need to "certify" that each vertex in A needs to be there. A good analogy is how in graph minors, the apex vertices are chosen to be adjacent to many vertices far apart on a grid minor. Our graph will not be a grid minor, but a labelled "grid-like graph" that is a complete immersion minor of $(\hat{G}, \hat{\lambda})$. This graph, which we denote by (H, λ), needs to relate to B; we will choose it so that $V(H) \subseteq B$. Intuitively, this is sufficient because for any small edge-cut of $\hat{G}$, almost all of $V(H)$, and therefore almost all of B, is on one of the sides (we will eventually prove results to this effect in Lemmas 5.1.1 and 5.1.2). Since $\hat{G}$ is 4-regular, we can move a few vertices to the other side of a cut without making it too large; so, informally, a signature is small relative to B if and only if it is small relative to $V(H)$.

We also need to relate the signatures of vertices in A to the "grid-like graph" (H, λ). We will have many edge-disjoint collections of cycles whose λ_A-weights generate $\mathbb{F}_2^A$. (We will use a different, more explicit definition, but this idea still motivates our approach.) We call a grid-like graph with these properties "A-rich". Such a graph "certifies" that we cannot perturb to any proper subset of A in order to make each signature small relative to B. Moreover, these graphs are "universal"; for any two $\mathbb{F}_2^A$-labelled, A-rich grid-like graphs whose weights sum to the same element, the larger one completely immerses the smaller one (as long as it is in fact much larger; we will prove a result to this effect in Proposition 5.2.1). So in particular we will be able to find a word graph of an A-universal word as a complete immersion minor.

Now suppose that we are given an "A-rich grid-like graph" (H, λ) with $V(H) \subseteq B$ as a complete immersion minor of $(\hat{G}, \hat{\lambda})$. For each $v \in V(G) - C - A$, we need to know whether we can add v to A or perturb v to A. Our first step will be to "clean up the grid-like graph with respect to v". To explain this idea, it is easiest to just imagine that there is a subspace N of $\mathbb{F}_2^A \times \mathbb{F}_2^v$ of dimension $|A|$ so that every edge of (H, λ) has $\lambda_A \times \lambda_v$-weight in N. By A-richness, restricting this subspace to A does not lower its dimension. So N tells us how to perturb v to A.

After perturbing as specified, every edge of H will have $\mathbb{F}_2^v$-weight zero; if the following statement holds, then we will add v to A.

(*) There is a collection of trails of $\hat{G}$ so that their edge-sets partition the edge-set of $\hat{G}$, their ends are in $V(H)$, and many have non-zero $\mathbb{F}_2^v$-weight.

In order to tell whether or not such a collection exists, we will identify all vertices in $V(H)$ to a new vertex b and consider the resulting signed graph; it will be 4-edge-connected by Theorem 3.2.3 of Bouchet [5, 8].

Equivalently, in the resulting Eulerian signed graph, we will ask for the largest size of a circuit-decomposition where each circuit is non-zero and hits $\{b\}$; a *circuit-decomposition* is a collection of circuits whose edge-sets partition the edge-set of the graph. The next chapter is dedicated to proving a precise min-max theorem for this problem. As a corollary of the min-max theorem, we will prove that either (*) holds, or the "perturbed signature" is small relative to $V(H)$ (and therefore also small relative to B). We will formalize this outline in Chapter 5 to prove the Local Structure Theorem.

Chapter 4

Decomposing a rooted Eulerian signed graph

4.1 Introduction

In this chapter we prove a precise min-max theorem (Theorem 4.2.1) for the following problem. We allow graphs to have loops and multiple edges throughout this chapter.

Problem 4.1.1. *Given a signed Eulerian graph (G, λ) with a vertex b, what is the maximum size of a circuit-decomposition where each circuit is non-zero and hits b?*

Recall that an *Eulerian graph* is a connected graph where each vertex has even degree, a circuit *hits* b if it has an edge incident to b, and a *circuit-decomposition* is a collection of circuits whose edge-sets partition the edge-set of the graph. See Section 3.2 for the relevant definitions on signed graphs and Section 3.6 for an explanation of how this problem is related to the Local Structure Theorem.

Máčajová and Škoviera [92] also studied Problem 4.1.1, especially for regular graphs, and used their results to characterize signed Eulerian graphs with flow number three in [91]. We hope that our min-max theorem can be used to solve their conjectures (see [92, Conjectures 1 and 2]), but we focus on the connection with vertex-minors. Recall from Section 3.3 that our motivation for the definition of "small signatures" was the problem of packing edge-disjoint non-zero circuits. We obtain the following relationship between "packing" and "decomposing".

Corollary 4.1.2. *For any signed* 4*-edge-connected Eulerian graph with a vertex* b*, if there are* ℓ *edge-disjoint non-zero circuits hitting* b*, then there is a circuit decomposition of size* $\lceil \ell/2 \rceil$ *where each circuit is non-zero and hits* b*.*

The bound is best possible, and 4-edge-connectivity is necessary.

The following corollary of the min-max theorem is the only result from this chapter that will be used in the proof of the Local Structure Theorem. However, we believe that the full min-max theorem (Theorem 4.2.1) will be useful in proving the Structural Conjecture; see Section 4.7 for a brief discussion. Notice that shifting does not change the answer to Problem 4.1.1, and the conditions below closely resemble the conditions for a signature to be "small relative to B".

Corollary 4.1.3. *If* (G, γ) *is an Eulerian* 4*-edge-connected signed graph with a vertex* b*, and there is no circuit decomposition of size larger than* ℓ *where each circuit is non-zero and hits* b*, then there exist a shifting* γ' *of* γ *and a set of vertices* X *which contains* b *so that*

(i) there are at most 4ℓ *edges with exactly one end in* X*, and*

(ii) there are at most ℓ *edges with both ends in* X *that have non-zero* λ'*-weight.*

Corollary 4.1.2 is already enough to conclude a version of Corollary 4.1.3 (with worse bounds) without using the min-max theorem. This is because we could use the theorem of Kakimura, Kawarabayashi, and Kobayashi [76] which was mentioned in Section 3.3 (informally, that edge-disjoint non-zero hitting circuits have the Erdös-Pósa property in graphs with no 3-edge-cuts). The lemma of Churchley [24, Lemma 3.5] could also be used in the same manner.

Churchley [24, Lemma 3.5] observed that a min-max theorem for the "packing version" of Problem 4.1.1 follows from a theorem of Chudnovsky, Geelen, Gerards, Goddyn, Lohman, and Seymour [22] on vertex-disjoint non-zero "rooted" paths in group-labelled graphs. Our proof of the min-max theorem is based on [22], which, in turn, is based on a short proof of the Tutte–Berge Formula using the matching matroid (see [22]). We will show that there is a matroid, which we call the "flooding matroid", underlying Problem 4.1.1.

4.2 The min-max theorem

Throughout the chapter we will be interested in "rooted Eulerian signed graphs"; so we call an *RES-graph* a tuple (G, γ, b) so that (G, γ) is an Eulerian signed graph and b is a

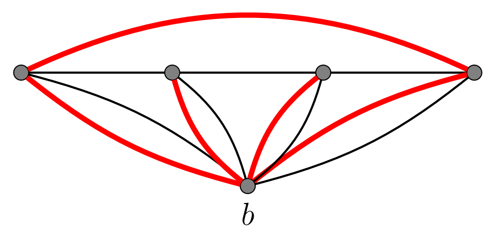

Figure 4.1: An RES-graph with four edge-disjoint non-zero circuits hitting b but flooding number three. The non-zero edges are bold red; we will use this convention in all of the figures in this chapter.

vertex of G. We call b the *root* of (G, γ, b). We write $\tilde{\nu}(G, \gamma, b)$ for the maximum size of a circuit-decomposition where each circuit is non-zero and hits b; we call this the *flooding number* of (G, γ, b). If no such decomposition exists, then we consider the flooding number to be zero.

If, after shifting, there is a small edge-cut so that the side containing b has few non-zero edges, then the flooding number must be small. Formally, for a set of edges F and a shifting γ' of γ, we write $\gamma'(F)$ for the number of non-zero edges in F according to γ'. For a set of vertices X, we write $E(X)$ for the set of edges with both ends in X and $\delta(X)$ for the set of edges with exactly one end in X. Using this notation we can state a pretty good upper bound on the flooding number;

$$\tilde{\nu}(G, \gamma, b) \leq \min_{\gamma', X} \left(\gamma'(E(X)) + \frac{1}{2}|\delta(X)| \right), \tag{4.1}$$

where the minimum is taken over all shiftings γ' of γ and all sets of vertices X that contain b.

If we did not require the edge-sets of the circuits to partition the edge-set of G, but just to be disjoint, then inequality (4.1) would be tight; this fact was observed by Churchley [24, Lemma 3.5] following from [22]. For the flooding number, however, inequality (4.1) is not tight. Intuitively, this is because parity matters; since we are interested in circuit-decompositions, the flooding number has the same parity as $\gamma(E(G))$. So in Figure 4.1, for instance, the flooding number must be odd. Therefore, while that example has $\deg(b)/2 = 4$ edge-disjoint non-zero circuits hitting b, its flooding number is just 3. (We write $\deg(v)$ for the degree of a vertex v.)

It turns out that parity is the only problem; the min-max theorem says that if we subtract one for each component of $G - X$ where "the parity is wrong", then inequality (4.1) becomes tight. So, let (G, γ, b) be an RES-graph, and let γ' be a shifting of γ. A set of vertices Y is γ'-*odd* if the parity of $\gamma'(E(Y) \cup \delta(Y))$ is different from the parity of $|\delta(Y)|/2$. Then, for a set of vertices X that contains b, we write $\mathrm{odd}_{\gamma'}(G - X)$ for the number of components of $G - X$ whose vertex-set is γ'-odd. Now we can state the min-max theorem.

Theorem 4.2.1. *For any RES-graph* (G, γ, b)*,*

$$\tilde{\nu}(G, \gamma, b) = \min_{\gamma', X} \left(\gamma'(E(X)) + \frac{1}{2}|\delta(X)| - \mathrm{odd}_{\gamma'}(G - X) \right),$$

where the minimum is taken over all shiftings γ' *of* γ *and all sets of vertices* X *that contain* b*.*

A *certificate* for an RES-graph (G, γ, b) is a tuple (X, γ') as in Theorem 4.2.1 so that equality holds. So, to prove the theorem, we need to show that a certificate exists.

We complete this section by proving the two corollaries of Theorem 4.2.1 (assuming that the theorem holds). They are restated below for convenience.

Corollary 4.1.2. *For any signed* 4*-edge-connected Eulerian graph with a vertex* b*, if there are* ℓ *edge-disjoint non-zero circuits hitting* b*, then there is a circuit decomposition of size* $\lceil \ell/2 \rceil$ *where each circuit is non-zero and hits* b*.*

Proof. Let (G, γ, b) be a 4-edge-connected RES-graph with a certificate (X, γ'). Since G is 4-edge-connected, $\mathrm{odd}_{\gamma'}(G - X) \leq \frac{1}{4}|\delta(X)|$. So

$$\begin{aligned} 2\tilde{\nu}(G, \gamma, b) &= 2\gamma'(E(X)) + |\delta(X)| - 2\mathrm{odd}_{\gamma'}(G - X) \\ &\geq 2\gamma'(E(X)) + \frac{1}{2}|\delta(X)| \\ &\geq \ell, \end{aligned}$$

since each of the ℓ edge-disjoint non-zero circuits must use either a non-zero edge in $E(X)$, or at least two edges in $\delta(X)$. Then $\tilde{\nu}(G, \gamma, b) \geq \lceil \ell/2 \rceil$ since $\tilde{\nu}(G, \gamma, b)$ is an integer. □

This proof of Corollary 4.1.2 also shows how to construct an example where the bound is tight. Informally, we take $\lceil \ell/2 \rceil$ disjoint copies of graphs like the one depicted in Figure 4.1, except where the vertex b has degree 4. Then we glue them together by identifying all of the copies of b. The resulting RES-graph has ℓ edge-disjoint non-zero circuits hitting b, and flooding number precisely $\lceil \ell/2 \rceil$.

Corollary 4.1.3. *If (G, γ) is an Eulerian 4-edge-connected signed graph with a vertex b, and there is no circuit decomposition of size larger than ℓ where each circuit is non-zero and hits b, then there exist a shifting γ' of γ and a set of vertices X which contains b so that*

(i) there are at most 4ℓ edges with exactly one end in X, and

(ii) there are at most ℓ edges with both ends in X that have non-zero λ'-weight.

Proof. Let (G, γ, b) be a 4-edge-connected RES-graph with $\tilde{\nu}(G, \gamma, b) \leq \ell$, and let (X, γ') be a certificate. Since G is 4-edge-connected,

$$\ell \geq \tilde{\nu}(G, \gamma, b) \geq \gamma'(E(X)) + \frac{1}{4}|\delta(X)|,$$

and the corollary holds. □

4.3 Preliminaries

Our approach to Theorem 4.2.1 is based on the observation that we do not need to require every circuit in the circuit-decomposition to be non-zero. So for an RES-graph (G, γ, b), a *flooding* is a collection of $\deg(b)/2$-many circuits that each hit b. A flooding is *optimal* if it contains as many non-zero circuits as possible. This gives an alternate definition of the flooding number as follows.

Lemma 4.3.1. *For any RES-graph (G, γ, b), the maximum number of non-zero circuits in a flooding is equal to $\tilde{\nu}(G, \gamma, b)$.*

Proof. By "splitting up" an Eulerian circuit at b, we can find a flooding that contains at least $\tilde{\nu}(G, \gamma, b)$-many non-zero circuits (since each non-zero circuit "splits up" into one non-zero circuit and one zero circuit). In the other direction, all of the zero circuits in an optimal flooding can be "combined" with one of its non-zero circuits (we may assume a non-zero circuit exists since certainly $\tilde{\nu}(G, \gamma, b) \geq 0$) to obtain a circuit-decomposition where each circuit is non-zero and hits b. This completes the proof. □

We will work with this alternate definition of the flooding number from now on.

Informally, our approach to Theorem 4.2.1 is to consider why the zero circuits in an optimal flooding cannot be "turned into" non-zero circuits. Notice that each optimal flooding has the same number of zero circuits; we will define a matroid of this rank. The bases of this matroid are obtained by selecting a "representative" for each zero circuit in an optimal flooding. A representative specifies a "split" of the circuit into two subtrails, each with b as one end, and the weight of the two subtrails (they must have the same weight since the circuit is a zero circuit). Our goal is to reduce to the case that this matroid has rank at most 1; then $(\{b\}, \gamma)$ will already be a certificate by a short parity argument.

Notation

We need some terminology to talk about trails more carefully. We think of RES-graphs as having half-edges; so an *edge* is an unordered pair of half-edges and an *arc* is an ordered pair of half-edges. (We use arcs to resolve technical issues with loops.) If $\{h_1, h_2\}$ is an edge, then there are two corresponding arcs, (h_1, h_2) and (h_2, h_1). The *tail* (respectively *head*) of an arc (h_1, h_2) is the vertex that is incident to h_1 (respectively h_2). A *trail* is a sequence of arcs so that the corresponding edges are all distinct and the head of each arc, other than the last one, is the tail of the next. The *tail* of a trail T is the tail of its first arc, and the *head* of T is the head of its last arc. If T has the same head and tail, say v, then T is a *circuit* or a *v-circuit*.

It is convenient to have some notation about how to "combine" trails. If T_1 and T_2 are trails so that the head of T_1 is the tail of T_2, then we can compose them into a new trail denoted (T_1, T_2). We simply write $\gamma(T_1, T_2)$ for its weight. Likewise, we can reverse a trail T to obtain another trail denoted T^{-1}. We also use this notation if f is an arc; so f^{-1} is the arc with the same edge, but in the reverse direction. As an example, we have $(T_1, T_2)^{-1} = (T_2^{-1}, T_1^{-1})$. A *subtrail* of T is any trail which can be obtained from T by deleting zero or more arcs at its beginning and end. It is a *proper subtrail* if it is not just T, a *prefix* if no arcs are deleted at the beginning, and a *suffix* if no arcs are deleted at the end. A *subcircuit* is a subtrail which also happens to be a circuit.

If C is a circuit whose sequence of arcs is $a_1, \ldots, a_t$, then any circuit of the form $a_i, a_{i+1}, \ldots, a_t, a_1, a_2, \ldots, a_{i-1}$ is *obtained from C by cyclically reordering its arcs.* This operation does not necessarily maintain the same subcircuits (because subcircuits are not allowed to "wrap around").

Proof approach

One direction of the min-max theorem (Theorem 4.2.1) is easy. So we will prove that direction now. Afterwards, we will outline the proof of the other direction.

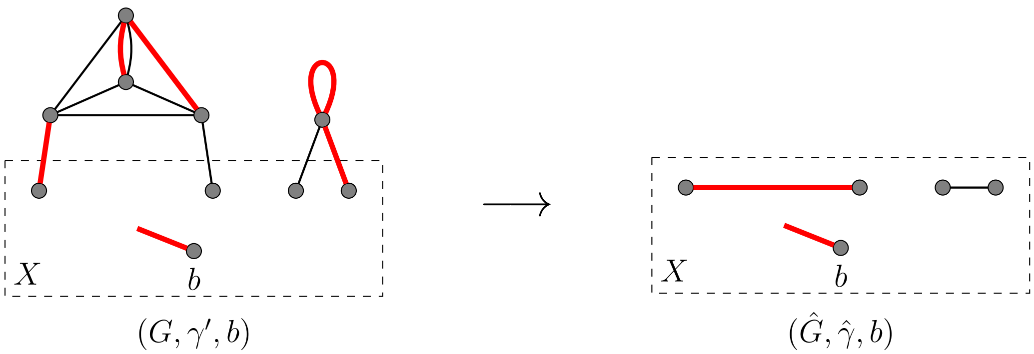

Figure 4.2: A depiction for Lemma 4.3.2. The flooding $\mathcal{C}$ is not drawn since in this particular example, any choice of $\mathcal{C}$ yields the same RES-graph $(\hat{G}, \hat{\gamma}, b)$.

Lemma 4.3.2. *For any RES-graph (G, γ, b), shifting γ' of γ, and set of vertices X that contains b,*

$$\tilde{\nu}(G, \gamma, b) \leq \gamma'(E(X)) + \frac{1}{2}|\delta(X)| - \mathrm{odd}_{\gamma'}(G - X). \tag{4.2}$$

Proof. Let $\mathcal{C}$ be an optimal flooding (so it contains precisely $\tilde{\nu}(G, \gamma, b)$-many non-zero circuits). We will use $\mathcal{C}$ and X to define another RES-graph $(\hat{G}, \hat{\gamma}, b)$ with vertex-set X. For an example, see Figure 4.2. This RES-graph $(\hat{G}, \hat{\gamma}, b)$ will have a flooding $\hat{\mathcal{C}}$ with the same number of non-zero circuits as $\mathcal{C}$, but the number of non-zero edges will be at most the right-hand side of inequality (4.2). This will complete the proof as every non-zero circuit in $\hat{\mathcal{C}}$ must have a non-zero edge.

So, let $(\hat{G}, \hat{\gamma}, b)$ be the graph obtained from (G, γ', b) by first deleting all vertices not in X and then adding $\frac{1}{2}|\delta(X)|$-many new edges as follows. We add an edge for each subtrail T of a circuit in $\mathcal{C}$ so that the tail and head of T are in X, the first and last edge of T are in $\delta(X)$, and no other edges of T are in $\delta(X)$. For each such trail T, we add an edge of weight $\gamma'(T)$ whose ends are the tail and head of T. So indeed $(\hat{G}, \hat{\gamma}, b)$ has a flooding $\hat{\mathcal{C}}$ with the same number of non-zero circuits as $\mathcal{C}$.

Now we need to show that the number of non-zero edges of $(\hat{G}, \hat{\gamma}, b)$ is equal to the right-hand side of inequality (4.2). So consider the vertex set Y of a component of $G - X$ so that Y is γ'-odd. We will show that there exists a trail T as above so that T is incident to a vertex in Y and satisfies $\gamma'(T) = 0$. First of all, observe that there are $|\delta(Y)|/2$ such trails T that are incident to a vertex in Y; moreover, the sum of their weights is $\gamma'(E(Y) \cup \delta(Y))$. So, since the parity of $\gamma'(E(Y) \cup \delta(Y))$ is different from the parity of

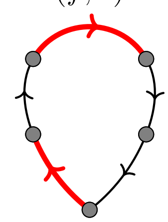

Figure 4.3: A b-circuit which is represented by $(f, 0)$, where f is the third arc of the circuit.

$|\delta(Y)|/2$, at least one of these trails T much have $\gamma'(T) = 0$. This completes the proof of Lemma 4.3.2. □

Now we outline the rest of the proof of Theorem 4.2.1. We begin by defining the "flooding matroid", whose rank is the number of non-zero circuits in an optimal flooding. We will prove that the relevant structure is in fact a matroid in Section 4.4.

Let (G, γ, b) be an RES-graph. For a zero circuit C, a *representative for* C is a tuple (f, α) so that f is an arc of C and $\alpha \in \{0, 1\}$ is the weight of the prefix of C whose last arc is f. See Figure 4.3 for an example. A *system of representatives* for a flooding $\mathcal{C}$ is a set B that consists of one representative for each zero circuit in $\mathcal{C}$. We define the *flooding matroid* $M(G, \gamma, b)$ by its ground set and its bases. The ground set of $M(G, \gamma, b)$ is the set of all tuples (f, α) so that f is an arc of (G, γ, b) and $\alpha \in \{0, 1\}$. A set B is a basis of $M(G, \gamma, b)$ if it is a system of representatives for an optimal flooding. (If the flooding number is equal to $\deg(b)/2$, then we view the empty set as a system of representatives for an optimal flooding; this guarantees that $M(G, \gamma, b)$ always has a basis.)

To prove Theorem 4.2.1, we will reduce to the case that $M(G, \gamma, b)$ has rank at most 1. Then a parity argument will immediately imply that $(\{b\}, \gamma)$ is a certificate. The key step in the reduction is to show that if (G, γ, b) is a counterexample to Theorem 4.2.1 which is, in a certain sense, "minimal", then for each arc f of $G - b$, both $(f, 0)$ and $(f, 1)$ are non-loop elements of $M(G, \gamma, b)$. We will show that any "minimal counterexample" to Theorem 4.2.1 has this property in Section 4.5. Then we will complete the proof of the theorem in Section 4.6 using the transitivity of parallel pairs and the following key lemma. (The proof of the lemma does not use the fact that $M(G, \gamma, b)$ is a matroid, just the definition of its bases.)

Lemma 4.3.3. *For any RES-graph (G, γ, b) and arcs f_0 and f_1 with the same head, there is no basis of $M(G, \gamma, b)$ which contains both $(f_0, 0)$ and $(f_1, 1)$.*

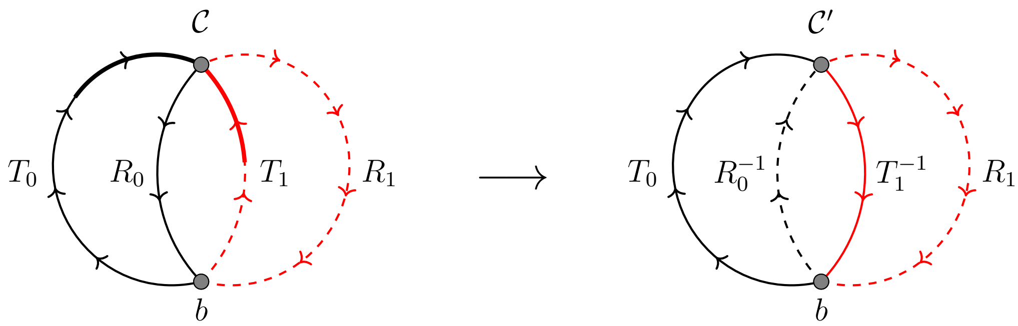

Figure 4.4: A depiction for Lemma 4.3.3. The bold arcs on the left represent f_0 and f_1. The colours represent the weights of the trails, and the dashed arcs show which circuits are in the floodings.

Proof. Suppose to the contrary that there is such a basis. Then there exists an optimal flooding $\mathcal{C}$ that contains distinct zero circuits C_0, C_1 so that $(f_0, 0)$ is a representative for C_0 and $(f_1, 1)$ is a representative for C_1. Thus there are trails T_0, R_0, T_1, R_1 so that $C_0 = (T_0, R_0)$, $C_1 = (T_1, R_1)$, the trail T_0 has weight 0, the trail T_1 has weight 1, and T_0 and T_1 have the same head. See Figure 4.4 (left) for a depiction.

We can obtain another flooding $\mathcal{C}'$ from $\mathcal{C}$ by replacing C_0 and C_1 with the circuits (T_0, T_1^{-1}) and (R_0^{-1}, R_1). This flooding is depicted in Figure 4.4 (right). This contradicts the optimality of $\mathcal{C}$ as the new two circuits are both non-zero. □

Transitions

We complete this section by introducing an important reduction operation involving transitions; a *transition* is a set of two half-edges which are incident to the same vertex, say v. We say the transition is *at* v. The *transitions of a trail* T are the transitions $\{h_1, h_2\}$ so that T has two consecutive arcs of the form (h_1', h_1) and (h_2, h_2'). Thus a trail with ℓ arcs is fully determined by its first arc and its $\ell - 1$ transitions. The *transitions of a flooding* are the transitions of its circuits.

Informally, the reduction operation that we are interested in removes a transition by combining two incident edges into a single edge. This is usually called "splitting off" or "lifting". This operation provides a powerful reduction technique for problems involving edge-connectivity (see the theorems of Lovász [88] and Mader [90]). Moreover, complete immersion minors could equivalently be defined via this operation. The formal definition will sound technical, but the point is that we are only allowed to split off transitions when

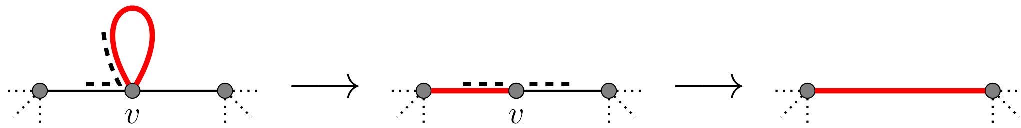

Figure 4.5: This figure depicts what happens when a transition is split off. The transitions are depicted by thick dashed lines which indicate their half-edges; we will use this convention in all of the figures.

the resulting graph is still connected. We can split off a transition at a vertex of degree 2, other than the root, by combining the two incident edges and then deleting that vertex; refer to Figure 4.5.

Let (G, γ, b) be an RES-graph . Let $\{h_1, h_2\}$ be a transition at a vertex v so that $\{h_1, h_2\}$ is not a loop edge and, where $\{h_1, h_1'\}$ and $\{h_2, h_2'\}$ are the two edges that contain h_1 and h_2 (respectively), there is no 2-edge-cut consisting of $\{h_1, h_1'\}$ and $\{h_2, h_2'\}$. Then *splitting off* $\{h_1, h_2\}$ deletes the edges $\{h_1, h_1'\}$ and $\{h_2, h_2'\}$, and then adds a new edge $\{h_1', h_2'\}$ of weight $\gamma(\{h_1, h_1'\}) + \gamma(\{h_2, h_2'\})$. So this operation results in another RES-graph . If instead $\{h_1, h_2\}$ was a transition at a vertex $v \neq b$ of degree 2, then we define the RES-graph obtained by *splitting off* $\{h_1, h_2\}$ in the same way, except we also delete v (which would otherwise have become an isolated vertex). So, in particular, we are allowed to split off any transition of a flooding of (G, γ, b).

4.4 The flooding matroid

This section is dedicated to proving that the flooding matroid is in fact a matroid. To do so, we will prove that the basis exchange axiom holds in Lemma 4.4.2. The proof will first reduce to the 4-edge-connected case; notice that an RES-graph (G, γ, b) is 4-edge-connected if and only if there is no set $Y \subseteq V(G) - \{b\}$ so that $|\delta(Y)| = 2$.

After that, we will find a transition to split off that "maintains" two different bases. We will need the following key lemma.

Lemma 4.4.1. *For any 4-edge-connected RES-graph (G, γ, b), vertex $v \neq b$, half-edge h incident to v, and basis B of $M(G, \gamma, b)$, the following holds.*

(*) For more than half of the transitions T at v which include h, there exists an optimal flooding $\mathcal{C}$ so that T is a transition of $\mathcal{C}$ and B is a system of representatives for $\mathcal{C}$.

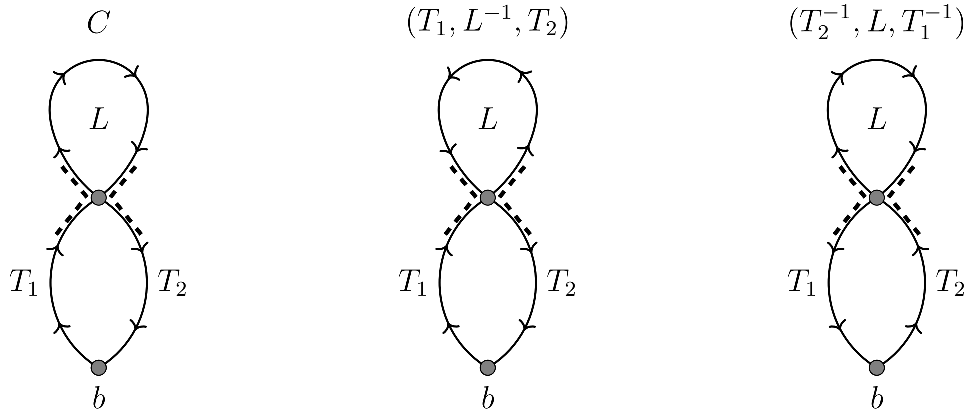

Figure 4.6: A depiction of the transitions $\{h_1, h_2\}$ and $\{h, h'\}$, as well as the relevant circuits in *Case 1* of Lemma 4.4.1.

Proof. We will say that a transition is *valid* if it is a transition at v that includes h. We will say that a valid transition T *works* if there exists an optimal flooding $\mathcal{C}$ so that T is a transition of $\mathcal{C}$ and B is a system of representatives for $\mathcal{C}$. So we are trying to show that more than half of the valid transitions work.

Notice that there are exactly $\deg(v) - 1$ valid transitions, which is an odd number. Now, fix an optimal flooding $\mathcal{C}$ for which B is a system of representatives. There is a unique half-edge h' so that $\{h, h'\}$ is a transition of $\mathcal{C}$. So $\{h, h'\}$ works. Thus it suffices to show that half of the other valid transitions also work. We will do this by proving that for each transition $\{h_1, h_2\} \neq \{h, h'\}$ of $\mathcal{C}$ at v, either $\{h, h_1\}$ or $\{h, h_2\}$ works. We need to consider how these two transitions $\{h_1, h_2\}$ and $\{h, h'\}$ of $\mathcal{C}$ "break up" the circuits of $\mathcal{C}$ into subtrails. We will consider two separate cases. We begin with the hardest one.

Case 1: There exists a circuit $C \in \mathcal{C}$ so that $\{h, h'\}$ and $\{h_1, h_2\}$ are both transitions of C.

Then there are trails T_1, L, T_2 so that $C = (T_1, L, T_2)$ and exactly one half-edge from each of the two transitions is contained in L, which is a v-circuit (see Figure 4.6, left). We can obtain a new flooding from $\mathcal{C}$ by deleting C and adding the circuit (T_1, L^{-1}, T_2) instead (see Figure 4.6, middle). If B is still a system of representatives, then we are done; so we may assume otherwise. Then C must be a zero circuit, and the arc of the representative for C must be in L. Now consider instead replacing C by (T_2^{-1}, L, T_1^{-1}) (see Figure 4.6, right). Again we may assume that B is no longer a system of representatives. It follows that $\gamma(T_1) \neq \gamma(T_2)$, and therefore that L is a non-zero circuit.

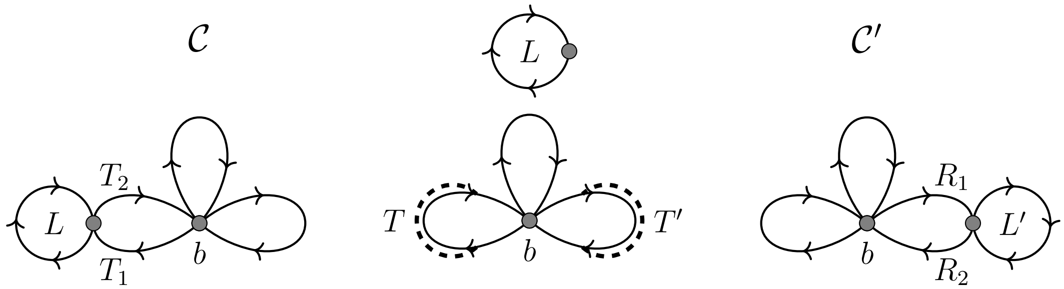

Figure 4.7: A depiction of how we use 4-edge-connectivity to find the new flooding $\mathcal{C}'$ in *Case 1* of Lemma 4.4.1.

Next we will use 4-edge-connectivity to "find another place to put L". This is the only time that we will use 4-edge-connectivity. Refer to Figure 4.7 for the following definitions. First of all, let us consider only the transitions of the circuit (T_1, T_2) and of the circuits in $\mathcal{C} - \{(T_1, L, T_2)\}$. Let T denote the transition that has exactly one half-edge from each of T_1 and T_2. Since (G, γ, b) is 4-edge-connected, there exists another transition $T' \neq T$ which is at a vertex u that is incident to an edge of L. Let L' be a u-circuit that is obtained from L by cyclically reordering its arcs. Let R_1 and R_2 be the trails so that T' has exactly one half-edge from each of R_1 and R_2 and (R_1, R_2) is a circuit under consideration (it is possible that $(R_1, R_2) = (T_1, T_2)$). By attaching L' onto T', we obtain a flooding $\mathcal{C}'$ so that T is a transition of $\mathcal{C}'$ and $(R_1, L', R_2) \in \mathcal{C}'$.

We claim that $\mathcal{C}'$ is an optimal flooding and (R_1, L', R_2) is a zero circuit. To see this, notice that when we removed L, we gained the non-zero circuit (T_1, T_2). So we must have lost a non-zero circuit when we added L' to (R_1, R_2). So our system of representatives B does not contain an element whose arc is in (R_1, R_2). Thus if B has a representative for (R_1, L', R_2), then we are done. Otherwise, since L is a non-zero circuit that contains an arc of an element in B, the set B must contain a representative for (R_2^{-1}, L', R_1^{-1}), and again we are done. This completes the first case.

Case 2: There exist distinct circuits $C, C' \in \mathcal{C}$ so that $\{h, h'\}$ is a transition of C and $\{h_1, h_2\}$ is a transition of C'.

Notice that the transitions $\{h, h'\}$ and $\{h_1, h_2\}$ "split" C and C' into four subtrails. More formally, there are unique trails T_1, T_2, R_1, R_2 so that $C = (T_1, T_2)$ and each of T_1, T_2 has exactly one half-edge in $\{h, h'\}$, and $C' = (R_1, R_2)$ and each of R_1, R_2 has exactly one half-edge in $\{h_1, h_2\}$. There are three ways to partition $\{T_1, T_2, R_1, R_2\}$ into two parts of size two, as depicted in Figure 4.8. Each of these three ways yields a unique flooding of (G, γ, b), up to reversing the two circuits that contain any of the half-edges h, h', h_1, h_2.

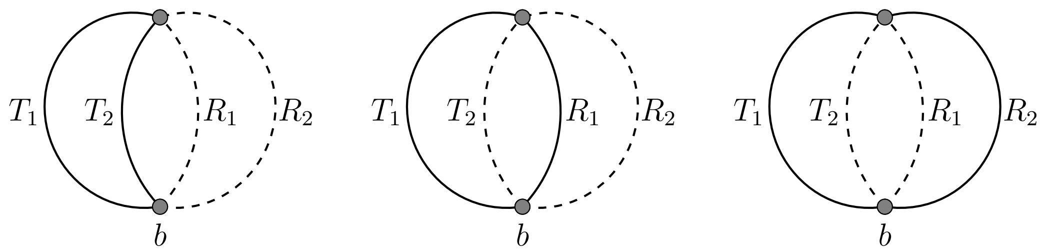

Figure 4.8: The three ways to partition $\{T_1, T_2, R_1, R_2\}$ into two parts of size two, from *Case 2* of Lemma 4.4.1.

We will call these two circuits the *new circuits* of the flooding.

One of these three floodings is our original flooding $\mathcal{C}$; we are interested in the other two floodings, which we denote by $\mathcal{C}_1$ and $\mathcal{C}_2$. Note that reversing the new circuits of $\mathcal{C}_1$ (respectively $\mathcal{C}_2$) does not affect whether or not $\mathcal{C}_1$ (respectively $\mathcal{C}_2$) is optimal. However, it might affect whether or not B is a system of representatives for $\mathcal{C}_1$ (respectively $\mathcal{C}_2$). There is an obvious choice to make though; if L is a new circuit of $\mathcal{C}_1$ (respectively $\mathcal{C}_2$) so that more elements of B have arcs in L^{-1} than in L, then replace L with L^{-1}. We claim that, with this choice, there exists $i \in \{1, 2\}$ so that $\mathcal{C}_i$ is an optimal flooding and B is a system of representatives for $\mathcal{C}_i$. This will complete the proof of Lemma 4.4.1. We now break into cases.

Case 2.1: Both (T_1, T_2) and (R_1, R_2) are non-zero circuits.

Then, as a multi-set, $\{\gamma(T_1), \gamma(T_2), \gamma(R_1), \gamma(R_2)\} = \{0, 0, 1, 1\}$. So for some $i \in \{1, 2\}$, both of the new circuits of $\mathcal{C}_i$ are non-zero. Then $\mathcal{C}_i$ is an optimal flooding and B is a system of representatives for $\mathcal{C}_i$.

Case 2.2: Exactly one of (T_1, T_2), (R_1, R_2) is a non-zero circuit.

Then, as a multi-set, $\{\gamma(T_1), \gamma(T_2), \gamma(R_1), \gamma(R_2)\}$ is either $\{0, 0, 0, 1\}$ or $\{0, 1, 1, 1\}$. So in fact both $\mathcal{C}_1$ and $\mathcal{C}_2$ are optimal. Let f be the arc of the element of B that represents whichever of (T_1, T_2), (R_1, R_2) is a zero circuit. Then f is in a zero circuit in either $\mathcal{C}_1$ or $\mathcal{C}_2$, and B is a system of representatives for that flooding.

Case 2.3: Both (T_1, T_2) and (R_1, R_2) are zero circuits.

As $\mathcal{C}$ is optimal, it follows that, as a multi-set, $\{\gamma(T_1), \gamma(T_2), \gamma(R_1), \gamma(R_2)\}$ is either $\{0, 0, 0, 0\}$ or $\{1, 1, 1, 1\}$. So both $\mathcal{C}_1$ and $\mathcal{C}_2$ are optimal. Now let f_1 (respectively f_2) be the arc of the element in B that represents (T_1, T_2) (respectively (R_1, R_2)). Then f_1 and

f_2 are in distinct circuits in either $\mathcal{C}_1$ or $\mathcal{C}_2$, and B is a system of representatives for that flooding.

This completes all possible cases and therefore the proof of Lemma 4.4.1. □

Now we are ready to prove that the basis exchange axiom holds, which is the final lemma of this section. As the flooding matroid always has a basis (possibly the empty set), this proves that the flooding matroid is in fact a matroid.

Lemma 4.4.2. *For any RES-graph (G, γ, b), any bases B_1 and B_2 of $M(G, \gamma, b)$, and any $b_1 \in B_1 - B_2$, there exists $b_2 \in B_2 - B_1$ so that $(B_1 - \{b_1\}) \cup \{b_2\}$ is a basis of $M(G, \gamma, b)$.*

Proof. Going for a contradiction, suppose that the lemma is false. Then choose a counterexample so that (G, γ, b) has as few edges as possible, and, subject to that, as many vertices as possible. This may seem strange but will prove to be convenient later. Such a choice is possible since an Eulerian graph with m edges has at most m vertices.

Our aim is to apply Lemma 4.4.1. So we need a vertex other than b to split off at, and we need (G, γ, b) to be 4-edge-connected. We take care of these things in the next two claims.

Claim 4.4.2.1. *There exists a vertex other than b.*

Proof. Suppose otherwise. Then every zero circuit in a flooding consists of a single loop f and must be represented by $(f, 0)$. We can reverse such a circuit to obtain a zero circuit represented by $(f^{-1}, 0)$. Then the element b_1 must be of the form $(f, 0)$, and we can take b_2 to be the element $(f^{-1}, 0) \in B_2 - B_1$. □

The next claim is actually the hardest part of the proof.

Claim 4.4.2.2. *The graph (G, γ, b) is 4-edge-connected.*

Proof. Otherwise, there exists a set $Y \subseteq V(G) - \{b\}$ with $|\delta(Y)| = 2$. Let $(\hat{G}, \hat{\lambda}, b)$ be the RES-graph that is obtained from (G, γ, b) by deleting all vertices in Y and then adding a new edge $\hat{e}$ whose ends are the neighbours of Y (possibly $\hat{e}$ is a loop) and whose weight is the sum of the weights of the edges in $E(Y) \cup \delta(Y)$. Note that $\tilde{\nu}(\hat{G}, \hat{\lambda}, b) = \tilde{\nu}(G, \gamma, b)$. The proof of the claim is fairly straightforward from here; we apply Lemma 4.4.2 to the graph $(\hat{G}, \hat{\lambda}, b)$, which has fewer edges than (G, γ, b). Unfortunately though, it is rather technical to state this precisely. We will begin by giving some definitions related to B_1 and B_2. So let $i \in \{1, 2\}$.

First of all, fix an optimal flooding $\mathcal{C}_i$ of (G,γ,b) so that B_i is a system of representatives for $\mathcal{C}_i$. Then there exists an optimal flooding $\hat{\mathcal{C}}_i$ of $(\hat{G},\hat{\gamma},b)$ so that $\mathcal{C}_i$ is obtained from $\hat{\mathcal{C}}_i$ by replacing the arc $\hat{f}_i$ that corresponds to $\hat{e}$ with a trail T_i with edge-set $E(Y) \cup \delta(Y)$. (To be more clear, $\hat{f}_i$ is the unique arc that corresponds to $\hat{e}$ and is also in a circuit of $\hat{\mathcal{C}}_i$.) Now, if no element of B_i has an arc in T_i, then B_i is also a system of representatives for $\hat{\mathcal{C}}_i$. Otherwise, let $(f_i,\alpha_i) \in B_i$ be the element whose arc is in T_i; then there exists $\hat{\alpha}_i \in \{0,1\}$ so that $(B_i - \{(f_i,\alpha_i)\}) \cup \{(\hat{f}_i,\hat{\alpha}_i)\}$ is a system of representatives for $\hat{\mathcal{C}}_i$. In either case, let $\hat{B}_i$ denote the system of representatives for $\hat{\mathcal{C}}_i$ that we have obtained. This completes the definitions.

Next we will apply Lemma 4.4.2 to $(\hat{G},\hat{\lambda},b)$, which has fewer edges than (G,γ,b). So let $\hat{b}_1$ be the element of $\hat{B}_1$ that corresponds to b_1. It is possible that $\hat{b}_1$ is in $\hat{B}_2$. In this case, $\hat{b}_1 = (\hat{f}_1,\hat{\alpha}_1) = (\hat{f}_2,\hat{\alpha}_2)$, and $(B_1 - \{(\hat{f}_1,\hat{\alpha}_1)\}) \cup \{(f_2,\alpha_2)\}$ is a basis of $M(G,\gamma,b)$. To see this, note that it is a system of representatives for the flooding that is obtained from $\hat{\mathcal{C}}_1$ by replacing the arc $\hat{f}_1$ with the trail T_2.

So we may assume that $\hat{b}_1 \in \hat{B}_1 - \hat{B}_2$. Then there exists $\hat{b}_2 \in \hat{B}_2 - \hat{B}_1$ so that $(\hat{B}_1 - \{\hat{b}_1\}) \cup \{\hat{b}_2\}$ is a basis of $M(\hat{G},\hat{\gamma},b)$. If $\hat{b}_2$ is in B_2 as well, then $(B_1 - \{b_1\}) \cup \{\hat{b}_2\}$ is a basis of $M(G,\gamma,b)$; we replace $\hat{f}_1$ or its inverse by T_1 or its inverse. Otherwise, $\hat{b}_2 = (\hat{f}_2,\hat{\alpha}_2)$, and instead $(B_1 - \{b_1\}) \cup \{(f_2,\alpha_2)\}$ is a basis of $M(G,\gamma,b)$; we replace $\hat{f}_2$ by T_2. We note that (f_2,α_2) is not in B_1 simply because $(B_1 - \{b_1\}) \cup \{(f_2,\alpha_2)\}$ is an optimal flooding and therefore has the same size as B. This completes the proof of Claim 4.4.2.2. □

Now, fix a vertex $v \neq b$ and a half-edge h incident to v. By Lemma 4.4.1 applied to B_1 and B_2, there exists a transition $\{h,h'\}$ at v so that there are optimal floodings $\mathcal{C}_1$ and $\mathcal{C}_2$ so that $\{h,h'\}$ is a transition of both $\mathcal{C}_1$ and $\mathcal{C}_2$ and B_1, B_2 (respectively) is a system of representatives for $\mathcal{C}_1, \mathcal{C}_2$ (respectively). Let $(\hat{G},\hat{\gamma},b)$ be the RES-graph that is obtained from (G,γ,b) by adding a new vertex v' and making the half-edges h and h' incident to v' instead of v. Then B_1 and B_2 are both bases of $M(\hat{G},\hat{\gamma},b)$. Moreover, Lemma 4.4.2 holds for $(\hat{G},\hat{\gamma},b)$ since it has the same number of edges as (G,γ,b) but more vertices. It follows that the lemma holds for (G,γ,b) as well. This is a contradiction and completes the proof of Lemma 4.4.2. □

4.5 The reduction step

Recall that our aim is to reduce to the case that the flooding matroid has rank at most 1. As mentioned before, the key step is to first reduce to the case that every arc of $G - b$ is

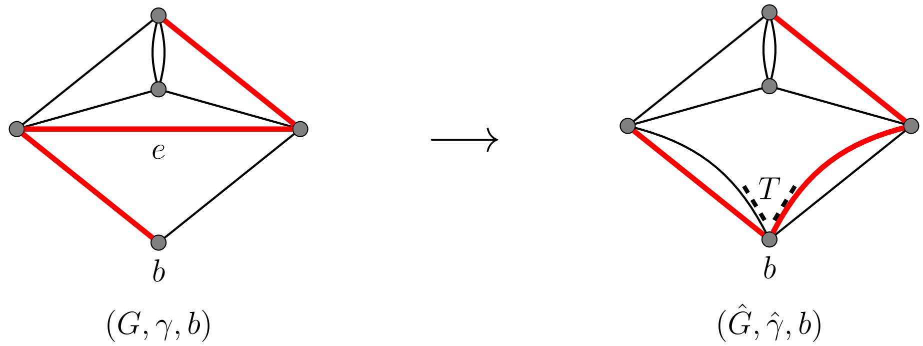

Figure 4.9: An RES-graph to the left, and, to the right, an RES-graph that is obtained from it by pulling e down to b.

the arc of two non-loop elements of $M(G, \gamma, b)$. That is what we will do in this section. We begin with some definitions.

Let (G, γ, b) be an RES-graph and let e be an edge of $G - b$. We say that an RES-graph $(\hat{G}, \hat{\gamma}, b)$ is *obtained from* (G, γ, b) *by pulling* e *down to* b if there exists a transition T so that (G, γ, b) is obtained from $(\hat{G}, \hat{\gamma}, b)$ by splitting off T, and e is the new edge of (G, γ, b). So pulling an edge down to b is the "inverse" of splitting off a transition at b. See Figure 4.9 for a depiction. Notice that $\hat{G} - b$ has fewer edges than $G - b$; this is the sense in which we are "reducing" to a smaller graph.

We call an RES-graph (G, γ, b) *critical* if for each edge e of $G - b$, the flooding number of any RES-graph that is obtained from (G, γ, b) by pulling e down to b is at least $\tilde{\nu}(G, \gamma, b) + 2$. This section is dedicated to proving two lemmas. The first essentially shows that any counterexample to Theorem 4.2.1 with $|E(G - b)|$ minimal must be critical. The second shows that for any critical graph, every edge not incident to b is in two non-loop elements of the flooding matroid.

Lemma 4.5.1. *If* (G, γ, b) *is an RES-graph,* e *is an edge of* $G - b$*, and* $(\hat{G}, \hat{\gamma}, b)$ *is obtained from* (G, γ, b) *by pulling* e *down to* b*, then if* $(\hat{G}, \hat{\gamma}, b)$ *has a certificate and* $\tilde{\nu}(\hat{G}, \hat{\gamma}, b) \leq \tilde{\nu}(G, \gamma, b) + 1$*, the RES-graph* (G, γ, b) *also has a certificate.*

Proof. First of all, observe that $\gamma(E(G))$ has the same parity as $\hat{\gamma}(E(\hat{G}))$. So the flooding numbers also have the same parity and $\tilde{\nu}(\hat{G}, \hat{\gamma}, b) \leq \tilde{\nu}(G, \gamma, b)$.

Now, observe that by performing the same sequence of shiftings in (G, γ, b), we may assume that there exists a set $X \subseteq V(\hat{G})$ that contains b so that $(X, \hat{\gamma})$ is a certificate for

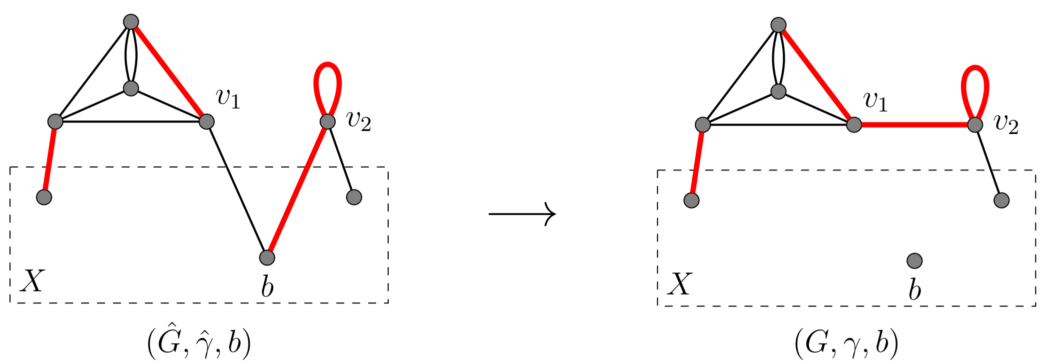

Figure 4.10: A depiction for *Case 4* of Lemma 4.5.1, where Y_1 and Y_2 are both odd in $(\hat{G}, \hat{\gamma}, b)$.

$(\hat{G}, \hat{\gamma}, b)$. We will show that (X, γ) is a certificate for (G, γ, b). Let v_1 and v_2 be the ends of the edge e of $G - b$. We now break into cases.

Case 1: Both v_1 and v_2 are in X.

Then if $\gamma(e) \neq 0$, then one of the two new edges of $(\hat{G}, \hat{\gamma}, b)$ is also non-zero. So indeed (X, γ) is a certificate for (G, γ, b).

Case 2: Exactly one of v_1, v_2 is in X.

Then the other of v_1, v_2 is in a component of $\hat{G} - X$; write Y for the vertex-set of that component. The only way (X, γ) might not be a certificate for (G, γ, b) is if Y is odd in $(\hat{G}, \hat{\gamma}, b)$ but not in (G, γ, b). However, if this occurs, then the new edge of (G, γ, b) which is not incident to a vertex in Y has non-zero weight. Therefore, it contributed to $\hat{\gamma}(E(X))$ but not to $\gamma(E(X))$; so again (X, γ) is a certificate for (G, γ, b).

Case 3: The vertices v_1 and v_2 are in the same component of $\hat{G} - X$.

Then $|\delta_G(X)| = |\delta_{\hat{G}}(X)| - 2$, and regardless of whether the vertex set of that component is odd in (G, γ, b), we still have that (X, γ) is a certificate for (G, γ, b).

Case 4: The vertices v_1 and v_2 are in different components of $\hat{G} - X$.

Let Y_1 and Y_2 be the vertex-sets of those components. Again we have that $|\delta_G(X)| = |\delta_{\hat{G}}(X)| - 2$. So the only possible problem is if Y_1 and Y_2 are odd in $(\hat{G}, \hat{\gamma}, b)$ and $Y_1 \cup Y_2$ is not odd in (G, γ, b). However, this is not possible; note that $|\delta_G(Y_1 \cup Y_2)|/2 = |\delta_{\hat{G}}(Y_1)|/2 + |\delta_{\hat{G}}(Y_2)|/2 - 1$, while the parity of the number of relevant non-zero edges just sums (as in Figure 4.10). So indeed (X, γ) is a certificate for (G, γ, b).

This completes all of the cases and therefore also the proof of Lemma 4.5.1. □

Here is the final lemma of this section.

Lemma 4.5.2. *If (G, γ, b) is an RES-graph that is critical, then every arc of $G - b$ is in two non-loop elements of the flooding matroid $M(G, \gamma, b)$.*

Proof. Let f be an arc that is not incident to b, and let e be the corresponding edge. We are trying to prove that for $i = 0, 1$, there exists an optimal flooding of (G, γ, b) that contains a zero circuit represented by (f, i).

Let $(\hat{G}, \hat{\gamma}_0, b)$ be an RES-graph that is obtained from (G, γ, b) by pulling e down to b. Such an RES-graph exists and, moreover, by adding 1 to both of the new edges of $(\hat{G}, \hat{\gamma}_0, b)$, we obtain another RES-graph $(\hat{G}, \hat{\gamma}_1, b)$ that is also obtained from (G, γ, b) by pulling e down to b. Since (G, γ, b) is critical, the flooding number of each of these two new RES-graphs is at least $\tilde{\nu}(G, \gamma, b) + 2$.

Now, for $i = 0, 1$, fix an optimal flooding $\hat{\mathcal{C}}_i$ of $(\hat{G}, \hat{\gamma}_i, b)$. We claim that for $i = 0, 1$, neither of the two new edges of $(\hat{G}, \hat{\gamma}_i, b)$, which we will denote by $\hat{e}_1$ and $\hat{e}_2$, are in a zero circuit in $\hat{\mathcal{C}}_i$. Going for a contradiction, suppose otherwise. Then $\hat{e}_1$ and $\hat{e}_2$ must be in the same circuit $\hat{C} \in \hat{\mathcal{C}}_i$, as otherwise we would obtain a contradiction to the fact that the flooding number went up. Then in (G, γ, b), the circuit $\hat{C}$ becomes a circuit C which contains e. As (G, γ, b) is connected, this circuit C can be attached back onto some other circuit in $\hat{\mathcal{C}}_i - \{\hat{C}\}$. This again contradicts the fact that the flooding number went up.

Now suppose that for some $i = 0, 1$, the edges $\hat{e}_1$ and $\hat{e}_2$ are in the same circuit $\hat{C} \in \hat{\mathcal{C}}_i$. Then $\hat{C}$ is non-zero, and, similarly to before, we can obtain an optimal flooding $\mathcal{C}$ of (G, γ, b) that contains a zero circuit with a non-zero subcircuit C containing f. That is, there exist trails T_1 and T_2 so that $\mathcal{C}$ contains a zero circuit (T_1, C, T_2), and C is a non-zero circuit that contains the arc f. We can obtain another optimal flooding of (G, γ, b) by replacing (T_1, C, T_2) with the circuit (T_2^{-1}, C, T_1^{-1}). Then one of (T_1, C, T_2), (T_2^{-1}, C, T_1^{-1}) is represented by $(f, 0)$, the other is represented by $(f, 1)$. So in this case the claim holds.

Finally, for the last case, suppose that the edges $\hat{e}_1$ and $\hat{e}_2$ are in distinct, non-zero circuits in both $\hat{\mathcal{C}}_0$ and $\hat{\mathcal{C}}_1$. By relabeling $\hat{e}_1$ and $\hat{e}_2$, we may assume that f is an ordered pair of half-edges (h_1, h_2), where h_1 is a half-edge of $\hat{e}_1$ and h_2 is a half-edge of $\hat{e}_2$. Then the flooding $\hat{\mathcal{C}}_0$ of $(\hat{G}, \hat{\gamma}_0, b)$ yields an optimal flooding of (G, γ, b) so that $(f, 1 + \hat{\gamma}_0(\hat{e}_2))$ represents a zero circuit. Likewise, $\hat{\mathcal{C}}_1$ yields an optimal flooding of (G, γ, b) so that $(f, 1 + \hat{\gamma}_1(\hat{e}_2))$ represents a zero circuit. As $\hat{\gamma}_0(\hat{e}_2) \neq \hat{\gamma}_1(\hat{e}_2)$, this completes the proof of Lemma 4.5.2. □

4.6 Completing the proof

This section is dedicated to completing the proof of Theorem 4.2.1, which is restated below for convenience.

Theorem 4.2.1. *For any RES-graph* (G, γ, b)*,*

$$\tilde{\nu}(G, \gamma, b) = \min_{\gamma', X} \left(\gamma'(E(X)) + \frac{1}{2}|\delta(X)| - \text{odd}_{\gamma'}(G - X) \right),$$

where the minimum is taken over all shiftings γ' *of* γ *and all sets of vertices* X *that contain* b*.*

Proof. Suppose for a contradiction that the theorem is false. Let (G, γ, b) be a counterexample so that $|E(G - b)|$ is as small as possible and, subject to that, so that $|E(G)|$ is as small as possible. Recall that we have already shown one direction of the inequality in Lemma 4.3.2.

From the last two sections, we also know that every arc of $G - b$ is in two non-loop elements of the flooding matroid $M(G, \gamma, b)$. This is because, by the choice of (G, γ, b), every graph $(\hat{G}, \hat{\gamma}, b)$ that is obtained from (G, γ, b) by pulling an edge in $E(G - b)$ down to b has a certificate. Thus, using Lemma 4.5.1 (which, informally, says when we can "lift" certificates from $(\hat{G}, \hat{\gamma}, b)$ to (G, γ, b)), it follows that (G, γ, b) is critical. Thus, by Lemma 4.5.2, every arc of $G - b$ is in two non-loop elements of the flooding matroid $M(G, \gamma, b)$.

The goal is to show that $M(G, \gamma, b)$ has rank 1 and thereby find a certificate. However, first we need to prove a straightforward claim.

Claim 4.2.1.1. *There are no loops at* b*, the graph* $G - b$ *is connected, and* $E(G - b)$ *is non-empty.*

Proof. First of all, there is no loop at b; otherwise, any certificate for the graph obtained from (G, γ, b) by deleting that loop also yields a certificate for (G, γ, b).

Now suppose for a contradiction that the graph $G - b$ is not connected. Then there are RES-graphs (G_1, γ_1, b) and (G_2, γ_2, b) so that $G = G_1 \cup G_2$, the only vertex in common between G_1 and G_2 is b, and both $V(G_1) - \{b\}$ and $V(G_2) - \{b\}$ are non-empty. Then by the choice of the counterexample (G, γ, b), it follows that for $i = 1, 2$, there is a certificate (X_i, γ_i') for (G_i, γ_i, b). We may assume that γ_1' and γ_2' are obtained without shifting at b; any time we wish to shift at b, we can instead shift at every vertex other than b.

Thus, there exists a shifting γ' of γ which, for $i = 1, 2$, agrees with γ'_i on $E(G_i)$. Then $\tilde{\nu}(G, \gamma, b) = \tilde{\nu}(G_1, \gamma_1, b) + \tilde{\nu}(G_2, \gamma_2, b)$, and $(X_1 \cup X_2, \gamma')$ is a certificate for (G, γ, b). This is a contradiction, which shows that $G - b$ is connected.

Finally, suppose for a contradiction that $E(G - b)$ is empty. From the last two paragraphs, this must mean that (G, γ, b) has two vertices and no loops. Let $\mathcal{C}$ be an optimal flooding of (G, γ, b). If $\mathcal{C}$ has no zero circuits, then $\tilde{\nu}(G, \gamma, b) = \deg(b)/2$ and $(\{b\}, \gamma)$ is a certificate. So we may assume that $\mathcal{C}$ contains a zero circuit C. Then, after possibly shifting at b, we may assume that both edges of C have weight zero. Then, since every other circuit of $\mathcal{C}$ "hits" C at the vertex other than b, this must mean that every zero circuit in $\mathcal{C}$ has both of its edges of weight zero (otherwise $\mathcal{C}$ would not be optimal). It follows that in this case $(V(G), \gamma)$ is a certificate. This is again a contradiction and completes the proof of Claim 4.2.1.1. □

The next claim almost completes the proof.

Claim 4.2.1.2. *The matroid $M(G, \gamma, b)$ has rank 1.*

Proof. Let F be the set of elements of $M(G, \gamma, b)$ whose arcs are not incident to b. We claim that F is a rank 1 set with no loops. We know that every arc of $G - b$ is in two non-loop elements of the flooding matroid $M(G, \gamma, b)$. Thus, for each arc f of $G - b$, the four elements in F whose arc is f or f^{-1} are all parallel. So by Lemma 4.3.3, transitivity of parallel pairs, and the fact that $G - b$ is connected and has an edge, it follows that F has rank 1.

Now consider an arc f whose tail is b. Note that if $(f^{-1}, 0)$ is a non-loop, then so is $(f, 0)$. The same holds for $(f^{-1}, 1)$ and $(f, 1)$. Furthermore, by Claim 4.2.1.1, there exists an arc of $G - b$ with the same head as f. Then as before, it follows that all of the non-loop elements of $M(G, \gamma, b)$ with f or f^{-1} as an arc are in the parallel class of F. It follows that $M(G, \gamma, b)$ has rank 1. □

Now, since the flooding matroid has rank 1, any optimal flooding $\mathcal{C}$ has exactly one zero circuit. Then the parity of $\gamma(E(G))$ is equal to the parity of $|\mathcal{C}| - 1$, which is different from the parity of $|\mathcal{C}| = \deg(b)/2$. So $\text{odd}_\gamma(G - b) = 1$ and $(\{b\}, \gamma)$ is a certificate. This is a contradiction, which completes the proof of Theorem 4.2.1. □

4.7 Possible generalization to group-labelled graphs

To understand the structure of graphs with a forbidden vertex-minor, we would like a generalization of Theorem 4.2.1 to $\mathbb{F}_2^k$-labelled graphs. A precise theorem seems necessary because, intuitively, a signature of size 2 does not help at all while a signature of size 4 does. (Recall from Lemma 3.3.2 of Bouchet [10] that a vertex can be "added" to a prime circle graph if and only if its signature has size at most 2.) We believe that we can prove such a generalization to group-labelled graphs using different proof techniques. (We just care about the additive group structure of $\mathbb{F}_2^k$.) Moreover, we believe that our proof gives a polynomial-time algorithm to find an optimal flooding, assuming that we can determine if two words evaluate to the same element in polynomial time.

Here is the more general version of the problem. A *group-labelled graph* is a tuple (G, γ) so that G is a directed graph and γ is a function from the arcs of G to a (not necessarily abelian or finite) group Γ. Trails are allowed to traverse arcs in either direction. The weight of a trail is the sum of the contributions of its arcs in order, where an arc e contributes $\gamma(e)$ if it is traversed in the forwards direction and $\gamma(e)^{-1}$ if it is traversed in the backwards direction. As before, for a vertex b, we are interested in the maximum size of a circuit-decomposition where each circuit is non-zero and begins and ends at b.

We are allowed to shift at any vertex v and by any group element α. *Shifting at v by α* changes the label of each non-loop arc e incident to v to

$$\begin{cases} \gamma(e) - \alpha & \text{if } e \in \delta^-(v) \\ \alpha + \gamma(e) & \text{if } e \in \delta^+(v). \end{cases}$$

Thus shifting just conjugates the weight of a circuit, so it does not change whether or not a circuit is non-zero.

In the group-labelled case we must first "decompose" on certain 2-edge-cuts. A *decomposable 2-edge-cut* is a set $X \subseteq V(G) - \{b\}$ so that $|\delta(X)| = 2$ and there is a shifting γ' of γ so that the group generated by $\{\gamma'(e) : e \in E(X)\}$ is isomorphic to the additive group of $\mathbb{F}_2^k$ for some non-negative integer k. Let $v_1, v_2 \in V(G) - X$ be the ends of the two arcs in $\delta(X)$. Then each trail that has arc-set $\delta(X) \cup E(X)$, tail v_1, and head v_2 has the same γ'-weight. We can *decompose on X* by deleting X and adding a new arc of that weight from v_1 to v_2; this does not change the flooding number.

We believe that Theorem 4.2.1 holds for group-labelled graphs with no decomposable 2-edge-cut, after subdividing each arc of $G - b$ once (this step needs to be done for technical reasons as shown in Figure 4.11; informally, odd components should really "count labelled

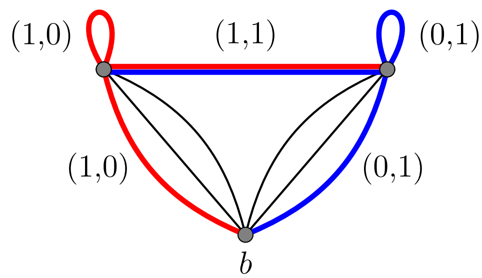

Figure 4.11: A depiction of a graph labelled over the additive group of $\mathbb{F}_2^2$ for which we need to subdivide at least the arc labelled $(1,1)$ to obtain a certificate.

half-edges"). As before, a set of vertices X is *odd* (with respect to a particular shifting) if there exists an element α of order 2 so that every non-zero arc with an end in X has weight α, and the parity of the number of non-zero arcs in $E(X) \cup \delta(X)$ is different from the parity of $|\delta(X)|/2$.

The main difficulty in proving such a generalization is that it is possible for the flooding matroid (which we believe still exists) to have rank 1 and yet for $V(G) - b$ to not be odd with respect to γ. So it is substantially more difficult to find the desired shifting.

Chapter 5

Completely immersing a rich grid-like graph

5.1 Rich grid-like graphs

In this chapter we prove the Local Structure Theorem, which is restated below for convenience; see Chapter 3 for the definitions and a discussion of labelled tour graphs.

The Local Structure Theorem (Geelen, McCarty, Wollan)**.** *For any $t \in \mathbb{N}$, there exists $\ell_t \in \mathbb{N}$ so that if (G, C) is a prime circle-structure whose graph does not have all $(t+1)$-vertex graphs as vertex-minors, up to isomorphism, and if B is a subset of C that induces a comparability grid, then there is a rank-t perturbation of (G, C) so that each signature is ℓ_t-small relative to B.*

Recall the outline of our proof approach from Section 3.6. Informally, we will grow a set $A \subseteq V(G) - C$ so that the labelled tour graph of (G, C) completely immerses a so-called "A-rich grid-like graph" whose vertex-set is contained in B. The set A cannot grow to have size $t + 1$ by Lemma 3.5.3 on A-universal words. So at some point we will get stuck; at that point we will be able to perturb each vertex to A to make its signature small relative to B.

In this section we introduce "A-rich grid-like graphs" and prove some basic lemmas about them. This chapter will use some of the notation from the last chapter; in particular, for a set of vertices X, we write $\delta(X)$ for the set of edges with exactly one end in X.

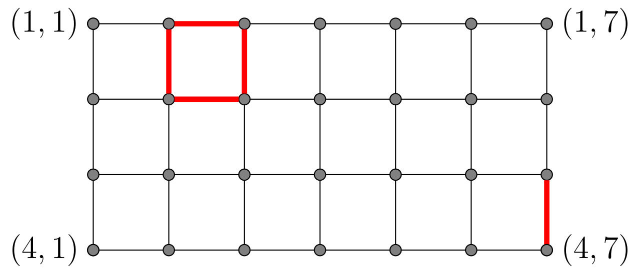

Figure 5.1: A 4×7 grid with a face on row 1 and column 2 (bold red), its four corner vertices (labelled), and a vertical boundary edge (also bold red).

Grid-like graphs

We begin by introducing some notation on grids. This will take a little while since we will be working with grids and grid-like graphs throughout the chapter. Refer to Figure 5.1 for the following discussion.

For positive integers m and n, an $m \times n$ *grid* is the graph with vertex set $\{(i, j) : i \in \{1, \ldots, m\}, j \in \{1, \ldots, n\}\}$, where there is an edge between vertices (i, j) and (i', j') whenever $|i - i'| + |j - j'| = 1$. We think of the vertex set as having m rows and n columns. Thus we call m the *height* and n the *width*, and we will refer to the *row* and *column* of a vertex. A *face* is a cycle of length 4; its *row* (respectively, *column*) is the smallest row (respectively, column) of one of its vertices.

As in the figure, we consider the vertex $(1, 1)$ to be the *top-left vertex*. We will view the columns as being ordered from left to right and the rows as being ordered from top to bottom (or equivalently, from smallest to largest). The *corner vertices* are the top-left, top-right, bottom-left, and bottom-right vertices. The *boundary vertices* are the vertices that are on the same row or column as a corner vertex, and the *boundary edges* are the edges that are in only one face. An edge between vertices (i, j) and $(i + 1, j)$ is *vertical*; other edges are *horizontal*. Finally, a *subgrid* of a grid is just a subgraph that is isomorphic to a grid. We will use all of this same terminology for subgrids.

We will usually work with 4-regular graphs. So we say that a graph G is *grid-like* if G is 4-regular and there exists $F \subseteq E(G)$ so that $G - F$ is isomorphic to a grid (see Figure 5.2). Thus every end of an edge in F is a boundary vertex of that grid. We call $G - F$ the *underlying grid of* G and use all of the prior definitions on grids for G as well (by implicitly referring to the underlying grid). With this set-up, we are implicitly fixing F as well as an

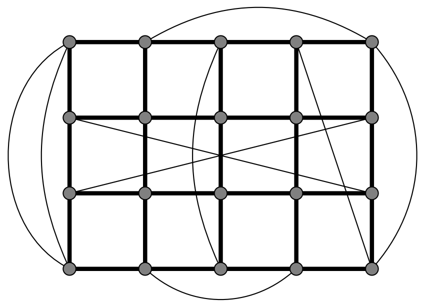

Figure 5.2: A 4×5 grid-like graph and its underlying grid (bold).

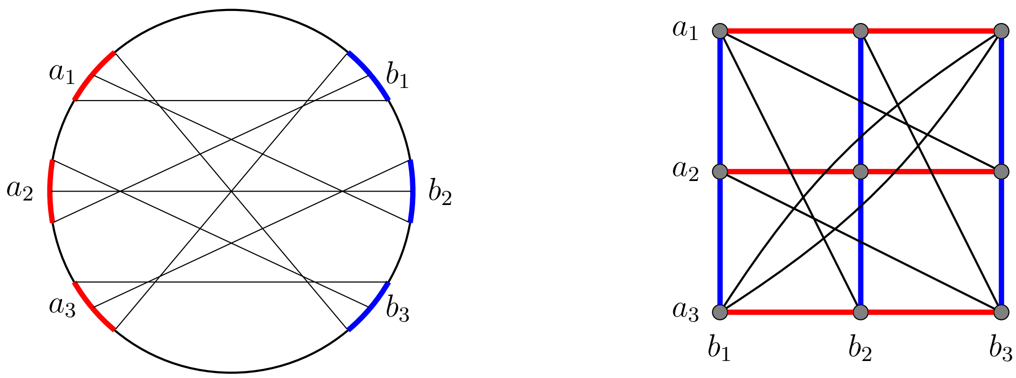

Figure 5.3: A chord diagram for the 3×3 comparability grid (left) and the corresponding tour graph (right).

isomorphism from $G - F$ to an actual grid. This is an abuse of notation, but we believe that the chapter is easiest to read this way.

Now let G be a grid-like graph. An *outside edge* is one that is not part of the underlying grid. The *distance* between two vertices is the length of a shortest path between them in the underlying grid (where the *length* of a path is the number of edges it contains). Then the *distance* between two sets of vertices S and T is the smallest distance between a vertex in S and a vertex in T. This completes all of the basic definitions.

One reason we are interested in these graphs is that the tour graphs of comparability grids are grid-like (the bound $n \geq 5$ in the following lemma is there to simplify the proof).

Lemma 5.1.1. *Any tour graph of an $n \times n$ comparability grid with $n \geq 5$ is an $n \times n$ grid-like graph.*

Proof. We considered a chord diagram for the $n \times n$ comparability grid in Lemma 2.2.1; we placed distinct points $b_1, b_2, \ldots, b_n, a_n, a_{n-1}, \ldots, a_1$ in clockwise order around a circle and included every chord with one end in $\{a_1, \ldots, a_n\}$ and one end in $\{b_1, \ldots, b_n\}$. In order to define the tour graph, however, we need a chord diagram where no two chords share an end. We can perturb the chords at each of $b_1, b_2, \ldots, b_n, a_n, a_{n-1}, \ldots, a_1$ slightly so as to achieve this; see Figure 5.3 (left).

The corresponding tour graph is depicted in Figure 5.3 (right); each vertex is labelled by the tuple (a_i, b_j) of points that were originally the ends of its chord. Then, at each of $a_1, a_2, \ldots, a_n$, we have a corresponding row of horizontal edges; see the chords in bold red in Figure 5.3. Similarly, each of $b_1, b_2, \ldots, b_n$ yields a column of vertical edges. So this particular tour graph is an $n \times n$ grid-like graph.

To consider other possible tour graphs, we need to consider cuts with at least two vertices on each side and at most 4 edges; Bouchet [5, 8] showed how to obtain every tour graph by changing such cuts in a certain fashion. By inspecting the underlying grid, since $n \geq 5$, every such cut has at most four vertices on one side. Now consider the outside edges; there is one matching between the top and bottom rows, and one matching between the leftmost and rightmost columns. Then, again by inspection, there is exactly one such cut: the cut where one side is $\{(a_n, b_1), (a_1, b_n)\}$. These two vertices have the same neighbourhood in the comparability grid, and the tour graph is unique up to switching their labels. So indeed any tour graph of the $n \times n$ comparability grid is an $n \times n$ grid-like graph. □

Richness

We are now ready to define rich grid-like graphs. Let V be a set, let $A \subseteq V$, and let (G, λ) be an $\mathbb{F}_2^V$-labeled grid-like graph. Then (G, λ) is A-*rich* with *richness* k if each edge of the underlying grid has λ_A-weight zero, and there exist k pairwise disjoint sets of $|A|$ outside edges whose λ_A-weights generate $\mathbb{F}_2^A$. We call those k pairwise disjoint sets the *batches* of (G, λ).

Note that every grid-like graph is $\emptyset$-rich with arbitrarily many batches (because the empty set is an $\emptyset$-rich batch). This will be our base case. Then we will sacrifice height, width, and richness to grow A. The base case holds since by Lemma 5.1.1, the tour graphs of the $n \times n$ comparability grid, for $n \geq 5$, are $n \times n$ grid-like graphs.

In fact, we could have proven the Local Structure Theorem for any set B which induces, instead of a comparability grid, a circle graph whose tour graphs are $n \times n$ grid-like graphs. Moreover, working with an $n \times n$ grid-like graph as a complete immersion minor is roughly equivalent to working with a tangle; each tangle "controls" a grid minor [114, (2.3)], and, for 4-edge-connected graphs, having a large grid as a minor is equivalent to having a large grid as an immersion minor [21]. We have already seen that for unlabelled Eulerian graphs, immersion minors and complete immersion minors are equivalent (Lemma 3.4.1).

We will not need any of the theory on tangles, however. We just need the fact that, for any small cut of a grid, almost all of the vertices are on one side.

Lemma 5.1.2. *For any grid G of height and width at least n and any $X \subseteq V(G)$ with $|\delta(X)| < n$, either X or $V(G) - X$ has size at most $|\delta(X)|^2$.*

Proof. Each row and each column that contains a vertex on each side of the cut contributes at least one to $|\delta(X)|$. So by symmetry between the two sides, we can assume that there is a column with no vertices in X. Then there is also a row with no vertices in X. So X is contained in at most $|\delta(X)|$ rows and columns, and $|X| \leq |\delta(X)|^2$, as desired. □

5.2 Universality

In this section we prove that A-rich grid-like graphs are "universal", and then use universality to find vertex-minors. It is well-known that unlabelled grids are "universal"; Dvořák and Klimošová [45] proved that every class of 4-edge-connected graphs of unbounded tree-width contains all graphs of maximum degree 4 as (strong) immersion minors.

Informally, the universality proposition says that for any $\mathbb{F}_2^A$-labelled A-rich grid-like graphs (G, λ) and (H, μ) whose edge-weights sum to the same element, if (G, λ) has sufficient height, width, and richness, then it completely immerses (H, μ). In fact we prove a stronger statement; if (G, λ) has sufficient richness, then the underlying grid of H can be "placed" on any subgrid of G with sufficient distance from the boundary. Intuitively, (G, λ) must then immerse (H, μ) because we can "freely draw" edge-disjoint trails from the boundary of H to the boundary of G. We can then "make the immersion complete" by "adding on each component of unused edges"; the edge-weights of each of these components will sum to zero since all non-zero edges of (G, λ) are on the outside.

With this intuition in mind, we are ready to state the proposition.

Proposition 5.2.1 (Universality). *There is a function $f_{5.2.1} : \mathbb{Z}^+ \times \mathbb{N} \to \mathbb{Z}^+$ so that for any set A, if (G, λ) and (H, μ) are non-empty $\mathbb{F}_2^A$-labelled A-rich grid-like graphs whose edge-weights sum to the same element, so that* *(i)* (G, λ) *has richness at least* $f_{5.2.1}(|V(H)|, |A|)$*, and*

(ii) *the underlying grid of* H *is a subgrid of* G *so that* $V(H)$ *has distance at least* $f_{5.2.1}(|V(H)|, |A|)$ *from the boundary vertices of* G*,*

then (G, λ) *completely immerses* (H, μ)*, where each edge of the underlying grid of* H *is sent to the same edge of* G*.*

Proof. Let $t \in \mathbb{Z}^+$, let A be a set, and define

$$\begin{aligned} d &:= 2t(|A|+1)+1, \\ k &:= 6t+8d, \text{ and} \\ f_{5.2.1}(t, |A|) &:= k. \end{aligned}$$

Now, let (G, λ) and (H, μ) be labelled graphs that satisfy the conditions of the proposition, where t is the number of vertices of H. Let $\hat{H}$ denote the underlying grid of H.

First of all, we claim that it suffices to find a collection $\mathcal{T} = (T_e : e$ is an outside edge of $H)$ of pairwise edge-disjoint trails of G so that

(i) for each outside edge $e = uv$ of H, the trail T_e has ends u and v and has $\lambda(T_e) = \mu(e)$, and

(ii) no trail in $\mathcal{T}$ uses any edge of $\hat{H}$ or any boundary edge of G.

So, suppose that such a collection $\mathcal{T}$ exists; we will show that then the proposition holds.

Consider the components of the graph obtained from (G, λ) by deleting the edges of $\hat{H}$ and the edges of each trail in $\mathcal{T}$. Since there is a component that includes all boundary edges of (G, λ), and only the outside edges of (G, λ) can be non-zero, it follows that there is at most one component with a non-zero edge. Furthermore, since the sum of the weights of the edges is the same in the two graphs, the weights of the edges of that component sum to zero. So each component is an Eulerian graph whose edge-weights sum to zero; thus we can add on each component to a trail in $\mathcal{T}$ without changing its weight. In this manner, we obtain another collection of trails with the appropriate weights so that the edge-sets of the trails partition $E(G) - E(\hat{H})$. Thus (G, λ) completely immerses (H, μ), as desired.

Now it just remains to show that such a collection $\mathcal{T}$ exists. Here is our overall approach. First of all, H has at most $2t$ edges and therefore at most $2t$ outside edges. For each of those outside edges, we will use one of the k batches of (G, λ) to build the corresponding

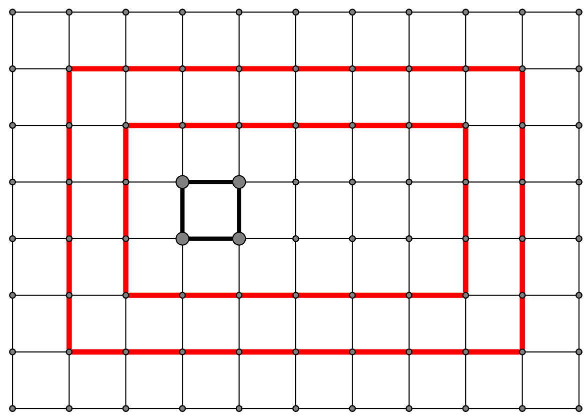

Figure 5.4: The underlying grid of G, the concentric cycles at distances 1 and 2 (bold red), and the subgrid $\hat{H}$ (also bold, with large vertices).

trail in $\mathcal{T}$. We have plenty of batches to choose from. Each batch has size $|A|$, and we will use $|A|+1$ "consecutive concentric cycles around $\hat{H}$" to build the corresponding trail.

To state the definition of a concentric cycle, let (m, n) be the bottom-right vertex of G. For each $i \in \{1, \ldots, d-1\}$, the *concentric cycle at distance* i is the cycle consisting of the boundary edges of the subgrid with top-left vertex $(1+i, 1+i)$ and bottom-right vertex $(m-i, n-i)$; see Figure 5.4. These cycles exist and are disjoint from the boundary of $\hat{H}$. Since $d = 2t(|A|+1)+1$, we can partition these concentric cycles into $2t$ parts of size $|A|+1$. Furthermore, we can assume that the distances of each part are of the form $i, i+1, \ldots, i+|A|$. Now fix an arbitrary injection from the outside edges of H to these $2t$ parts.

Next we show how to choose batches. Consider just those batches whose ends have distance at least d from the corners of G and are not on a row or column of $\hat{H}$. There are at least $2t$ such batches; at most $8d$ outside edges are too close to a corner of G, and at most $4t$ outside edges are on a row or column of $\hat{H}$. So we can fix an injection from the outside edges of H to these $2t$ such batches.

Now we show how to obtain $\mathcal{T}$; see Figure 5.5. Consider an outside edge $e = uv$ of H with an associated batch F and associated concentric cycles at distances $i, i+1, \ldots, i+|A|$. The trail corresponding trail T_e begins at u, and its first edge is an edge of $E(G) - E(\hat{H})$ that is incident to u. Then it proceeds along this direction until hitting the first concentric cycle, at distance $i+|A|$. If $\mu(e) = 0$, then it proceeds around this concentric cycle and then returns to v in the same manner.

So now assume that $\mu(e) \neq 0$. Then there are distinct edges $f_1, \ldots, f_\ell \in F$ so that

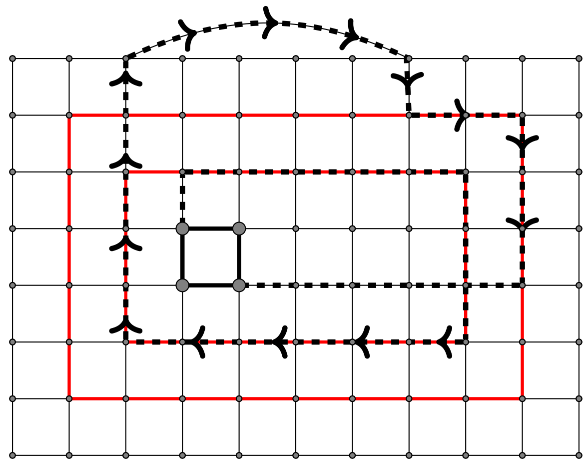

Figure 5.5: A depiction of a possible trail T_e (bold dashed) where the ends of e are the top-left and bottom-right vertices of $\hat{H}$.

$\mu(e) = \lambda(f_1) + \ldots + \lambda(f_\ell)$. Then the trail proceeds along the concentric cycle, and then directly out to an end of f_1, then along f_1, and finally back to the next concentric cycle, at distance $i + |A| - 1$. This process repeats until the trail has used all of the edges $f_1, \ldots, f_\ell$ and returned to the concentric cycle at distance $i + |A| - \ell$. Finally, the trail proceeds along that concentric cycle, and then back to v in the same manner. Due to the choice of the batches, these trails exist and satisfy all of the conditions that were outlined earlier. So this completes the proof of Proposition 5.2.1. □

We complete this section by showing how to use universality to find vertex-minors.

Lemma 5.2.2. *There is a function $f_{5.2.2} : \mathbb{N} \to \mathbb{Z}^+$ so that if A is a set and (G, C) is a prime circle-structure whose labelled tour graph completely immerses an A-rich grid-like graph of richness, height, and width at least $f_{5.2.2}(|A|)$, then G contains every graph with vertex set A as a vertex-minor.*

Proof. We write $f_{5.2.1}$ for the function from Proposition 5.2.1. Now, define $f_{5.2.2}(|A|) :=$ $2f_{5.2.1}(2^{|A|}, |A|) + 2^{|A|}$. Suppose that the conditions of the lemma hold for a circle-structure (G, C). Recall that, by Lemma 3.5.2, there exists an A-universal word W with $2^{|A|}$ letters. Moreover, the word graph of W is an A-rich grid-like graph of height 1 and width $2^{|A|}$ whose edge-weights sum to zero; see Figure 5.6. So by Proposition 5.2.1 on universality, the labelled tour graph of (G, C) completely immerses the word graph of W. Thus by Lemma 3.5.3 on A-universal words, G has every graph with vertex-set A as a vertex-minor.

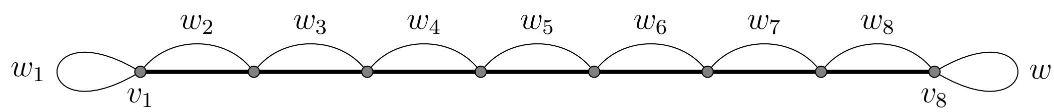

Figure 5.6: The word graph of $W = w_1, \ldots, w_8$ as an A-rich grid-like graph, with the underlying grid (bold), outside edges (labelled), and two corner vertices (also labelled).

□

5.3 Finding a clean toroidal grid

In the next few sections, we show that either we can grow A, or we can perturb to A so that each signature is small relative to the vertex-set of the fixed comparability grid. In order to do so, we attempt to add each vertex v to A one at a time. Our goal is to apply Corollary 4.1.3 on floodings from the last chapter; this result will tell us whether to add v to A or perturb v to A. However, before we apply the corollary, we reduce to the case of an "(A, v)-clean toroidal grid"; this is a particular type of A-rich grid-like graph.

To motivate the "cleaning" step, suppose that we have an A-rich grid-like graph so that every edge has $\lambda_A \times \lambda_v$-weight in a subspace N of $\mathbb{F}_2^A \times \mathbb{F}_2^v$ of dimension $|A|$. By A-richness, the λ_A-weights of each batch generate $\mathbb{F}_2^A$; so, intuitively, this subspace N tells us how to perturb v to A. We perturb v to A as specified and then apply Corollary 4.1.3 on floodings. Equivalently, we apply Corollary 4.1.3 to the signature which is obtained by labelling edges in the quotient space $\mathbb{F}_2^A \times \mathbb{F}_2^v / N$ (which is isomorphic to $\mathbb{F}_2$). The actual definition will be somewhat different, but we will still have a particular subspace N which "tells us how to perturb v to A".

There are two main reasons we reduce to the case of "toroidal grids". First, they are 4-edge-connected, and Corollary 4.1.3 only holds for 4-edge-connected graphs. Second, they help us "find many disjoint A-rich grids", informally.

Clean toroidal grids

A *toroidal grid* (see Figure 5.7) is a grid-like graph H so that for each column c (respectively, row r), there is an outside edge joining the topmost and bottom-most boundary vertices on column c (respectively, leftmost and rightmost boundary vertices on row r). We call

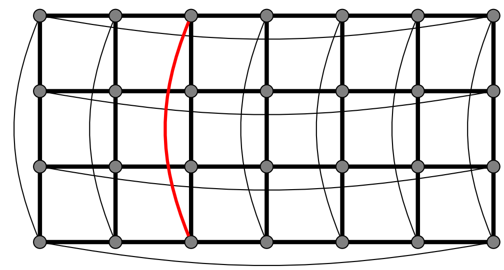

Figure 5.7: A 4×7 toroidal grid with its underlying grid (bold) and an outside-vertical edge on column 3 (bold red).

these edges *outside-vertical* on *column* c (respectively, *outside-horizontal* on *row* r). We use all of the same terminology for toroidal grids as for grid-like graphs.

As mentioned before, toroidal grids are 4-edge-connected.

Lemma 5.3.1. *Any toroidal grid of height and width at least* 2 *is* 4*-edge-connected.*

Proof. Consider a cut with at least one vertex on each side. Each row and each column that has a vertex on either side contributes at least two edges towards the cut. There must be at least two such columns or rows since the height and width is at least 2. □

Now we can define "(A, v)-clean toroidal grids". So, let A and V be sets with $A \subseteq V$, let $v \in V - A$, and let (G, λ) be an $\mathbb{F}_2^V$-labelled toroidal grid. Then (G, λ) is (A, v)-*clean* with *richness* k (see Figure 5.8) if each edge of the underlying grid has $\lambda_A \times \lambda_v$-weight zero, and there exist pairwise disjoint sets $F_1, \ldots, F_k$ so that

(i) (G, λ) is A-rich with batches $F_1, \ldots, F_k$,

(ii) there exist disjoint subgrids $H_1, \ldots, H_k$ of the same height as G so that for each $i \in \{1, \ldots, k\}$, every edge in F_i is an outside-vertical edge with its ends in $V(H_i)$, and

(iii) the $\lambda_A \times \lambda_v$-weights of each of $F_1, \ldots, F_k$ generate the same subspace.

Again we call $F_1, \ldots, F_k$ the *batches*. We call the subspace that is generated by the $\lambda_A \times \lambda_v$-weights of each batch the *subspace* of (G, λ).

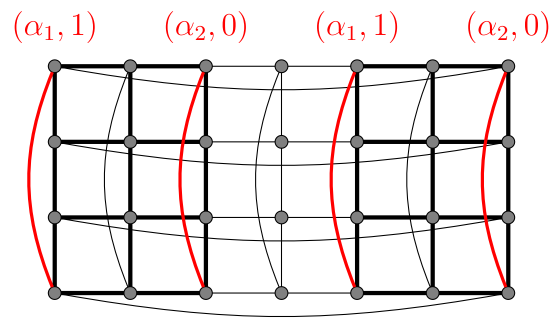

Figure 5.8: An (A, v)-clean toroidal grid with two batches of size $|A| = 2$ (in red, labelled above by their $\lambda_A \times \lambda_v$-weights), and the subgrids H_1 and H_2 (bold).

Finding a complete immersion

To find an (A, v)-clean toroidal grid as a complete immersion, we will first find a grid where every edge has λ_v-weight zero. This is the hardest part of the proof; the rest will follow quickly from Proposition 5.2.1 on universality. To find this well-behaved grid, we need a lemma on decomposing words.

A *decomposition* of a word W is a sequence $z_1, \ldots, z_k$ of non-empty subwords so that $W = z_1, \ldots, z_k$. (It is important that the subwords in a decomposition must be non-empty.) The following lemma can be proven by applying the pigeonhole principle to the sequence of prefix values of a word (see Section 3.5 for the relevant definitions on words).

Lemma 5.3.2. *There is a function* $f_{5.3.2} : (\mathbb{Z}^+)^2 \to \mathbb{Z}^+$ *so that for any* $m, k \in \mathbb{Z}^+$, *every word over* $\mathbb{F}_2^m$ *of length* $f_{5.3.2}(m, k)$ *has a decomposition* $w_0, z_1, \ldots, z_k, w_{k+1}$ *so that* z_i *evaluates to zero for each* $i \in \{1, \ldots, k\}$.

We will apply this lemma to a word $w_1, \ldots, w_n$ over $\mathbb{F}_2^m$ that is obtained from an $(m+1) \times (n+1)$ grid; each letter w_i will represent the λ_v-weights of the m faces that are on column i. It is enough to consider the faces due to the following lemma.

Lemma 5.3.3. *If* (G, γ) *is a labelled grid so that each face has weight zero, then there exists a shifting* γ' *of* γ *so that for each edge* e, $\gamma'(e) = 0$.

Proof. The set of facial cycles is a basis for the cycle space of G. So every circuit has weight zero since, if C_1 and C_2 are the edge-sets of two circuits, then their symmetric difference has weight

$$\lambda(C_1) + \lambda(C_2) - 2\lambda(C_1 \cap C_2) = \lambda(C_1) + \lambda(C_2).$$

So Harary's theorem [70] on balanced signed graphs says that the desired shifting exists. □

Now we are ready to prove the main proposition of this section; it tells us that we can always find an (A, v)-clean toroidal grid as a complete immersion minor of a sufficiently large and rich A-rich grid-like graph.

Proposition 5.3.4. *There is a function* $f_{5.3.4} : \mathbb{Z}^+ \times \mathbb{N} \to \mathbb{Z}^+$ *so that for any* $m \in \mathbb{Z}^+$, $A \subseteq V$, *and* $v \in V - A$, *every* $\mathbb{F}_2^V$*-labelled,* A*-rich grid-like graph of richness, height, and width at least* $f_{5.3.4}(m, |A|)$ *completely immerses an* (A, v)*-clean toroidal grid of richness, height, and width at least* m.

Proof. We write $f_{5.2.1}$ for the function from Proposition 5.2.1 and $f_{5.3.2}$ for the function from Lemma 5.3.2. Let $m \in \mathbb{Z}^+$, let $a \in \mathbb{N}$, and define

$$\begin{aligned} k &:= (a+1)m2^a, \\ t &:= (m+1)(k+1), \\ n &:= (a+2)f_{5.2.1}(t, a) + f_{5.3.2}(m, k) + m + 1, \text{ and} \\ f_{5.3.4}(m, a) &:= 2f_{5.2.1}(n^2, a) + n. \end{aligned}$$

Suppose that the conditions of the proposition hold for some set A of size a and some A-rich grid-like graph (G, λ).

Our aim is to apply Lemma 5.3.2, the Ramsey-type theorem on words. Since each letter of the word will come from a column of faces, we will want to "combine columns". So it is helpful to have the following claim.

Claim 5.3.4.1. *The graph* (G, λ) *completely immerses an* A*-rich grid-like graph* (H, μ) *of richness* $f_{5.2.1}(t, a)$ *and height and width* n *so that for each column* c *of* H*, there is an edge of* μ_A*-weight zero whose ends are the two boundary vertices on column* c.

Proof. Notice that (G, λ) has an $n \times n$ subgrid whose vertex-set has distance at least $f_{5.2.1}(n^2, a)$ from the boundary vertices of G. So by Proposition 5.2.1 on universality, we only need to show there exists a graph (H, μ) which satisfies the claim so that the $\mathbb{F}_2^A$-weights of the edges sum to the same element in both graphs. Such a graph exists since $n \geq af_{5.2.1}(t, a) + 1$; we can add $f_{5.2.1}(t, a)$ batches of size a and then correct the sum of the weights with one additional edge. □

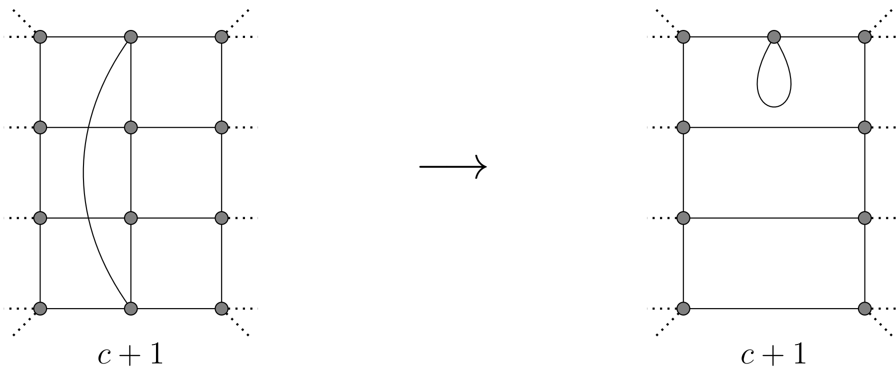

Figure 5.9: Combining two columns, before removing the final vertex.

Now let (H, μ) be the $n \times n$ grid-like graph from the claim; we may assume that it is $\mathbb{F}_2^V$-labelled. By the definition of n, the graph (H, μ) has a subgrid $\hat{H}$ of height $m+1$ and width $f_{5.3.2}(m, k)+1$ so that $V(\hat{H})$ has distance at least $f_{5.2.1}(t, a)$ from the boundary vertices of H. For each column c of faces of $\hat{H}$, there is a corresponding element of $\mathbb{F}_2^m$ which represents the μ_v-weights of the m faces on column c. Thus we obtain a word over $\mathbb{F}_2^m$ of length $f_{5.3.2}(m, k)$. By Lemma 5.3.2, this word has a decomposition $w_0, z_1, \ldots, z_k, w_{k+1}$ so that for each $i \in \{1, \ldots, k\}$, the subword z_i evaluates to zero. We will "combine" columns according to these subwords, thereby finding a complete immersion minor.

We will say how to remove each vertex of (H, μ) one at a time by specifying which incident edges to "combine". Suppose that we wish to "combine" columns c and $c+1$ of (H, μ). The vertices on column $c+1$ are incident to one face on column c and one face on column $c+1$; these are the vertices that we will remove. So, first of all, remove every vertex on column $c+1$ except for the vertex on the first row by "combining" the two incident horizontal edges as in Figure 5.9. By Claim 5.3.4.1, the final remaining vertex on column $c+1$ has a loop with μ_A-weight zero. Thus any way of removing this vertex results in an A-rich grid-like graph of the same richness and one smaller width. Moreover, since $f_{5.2.1}(t, a) \geq 1$, none of the faces of $\hat{H}$ contain boundary edges of H; so the weights of the faces of $\hat{H}$ sum.

Let us update the names of these graphs (H, μ) and $\hat{H}$ as we combine columns according to $w_0, z_1, \ldots, z_k, w_{k+1}$. We do not lose any richness as we "combine columns"; so in the end we still have richness at least $f_{5.2.1}(t, a)$. Furthermore, in the end, every face of $\hat{H}$ that is not on the first or last column has μ_v-weight zero. Let $\hat{H}'$ be the subgrid obtained from $\hat{H}$ by deleting every vertex on the first or last column. Then $\hat{H}'$ is an $(m+1) \times (k+1)$

subgrid of H so that $V(\hat{H}')$ has distance at least $f_{5.2.1}(t,a)$ from the boundary vertices of H. So, since $t=(m+1)(k+1)$, we can apply Proposition 5.2.1 on universality to an A-rich grid-like graph with underlying grid $\hat{H}'$.

So, since $k+1=(a+1)m2^a+1$, the labelled graph (H,μ) completely immerses a toroidal grid $(\widetilde{H},\widetilde{\mu})$ so that

(i) $(\widetilde{H},\widetilde{\mu})$ is A-rich with $m2^a$ batches,

(ii) each batch consists of outside-vertical edges on consecutive columns,

(iii) $\hat{H}'$ is the underlying subgrid of $\widetilde{H}$, and

(iv) each batch has the same set of $\widetilde{\mu}_A$-weights.

Furthermore, by applying Lemma 5.3.3 and shifting $\widetilde{\mu}_v$, we may assume that each edge of the underlying grid of $\hat{H}'$ has $\widetilde{\mu}_v$-weight zero. Now by the pigeonhole principle, there exist m batches which each have the same set of $\widetilde{\mu}_A\times\widetilde{\mu}_v$-weights. Then $(\widetilde{H},\widetilde{\mu})$ is an (A,v)-clean toroidal grid with richness m. This completes the proof of Proposition 5.3.4. □

5.4 Perturbing or growing

Informally, this section is dedicated to proving that if we have a "sufficiently large and rich" (A,v)-clean toroidal grid as a complete immersion minor, then either we can perturb v to A or we can add v to A.

In order to add v to A (that is, to find an $A\cup\{v\}$-rich grid-like graph as a complete immersion minor), we will use the following "rerouting lemma" on floodings. It says that under minor conditions, we can find a half-edge at which we can "reroute" the flooding in every possible way. See Section 4.3 for the relevant definitions; in particular, recall that an RES-graph is a "rooted Eulerian signed graph".

Lemma 5.4.1 (Rerouting Lemma). *If (G,γ,b) is a 4-edge-connected RES-graph with $\widetilde{\nu}(G,\gamma,b)\geq\ell$ and $|\delta(b)|/2>\ell+1$, then there exists a vertex $x\neq b$ and a half-edge h incident to x so that every transition at x which includes h is a transition of a flooding with at least ℓ non-zero circuits.*

Proof. The proof is a very reminiscent of the proof of Lemma 4.4.1, which, informally, said that "most of the relevant transitions maintain a basis of the flooding matroid". We say that a transition *works* if it is a transition of a flooding with at least ℓ non-zero circuits. So the goal is to choose a particular half-edge where all of the "relevant" transitions work.

Let $\mathcal{C}$ be an optimal flooding. Since $|\delta(b)|/2 > \ell + 1$, there are more than $\ell + 1$ circuits in $\mathcal{C}$ with more than one edge. Choose such a circuit $C \in \mathcal{C}$; if possible, choose C to be a zero circuit. Let $x \neq b$ be a vertex which is incident to a half-edge h in the first arc of C (thus, informally, x is the first vertex after b which is "hit" by C). Let h' denote the half-edge so that $\{h, h'\}$ is a transition of C. We consider each transition $\{h_1, h_2\} \neq \{h, h'\}$ of $\mathcal{C}$ at x one at a time. To prove the lemma, we need to show that both of the transitions $\{h, h_1\}$ and $\{h, h_2\}$ work. We break into two cases.

Case 1: The transition $\{h_1, h_2\}$ is a transition of C.

Then there are trails T_1, L, T_2 so that $C = (T_1, L, T_2)$ and exactly one half-edge from each of $\{h_1, h_2\}$ and $\{h, h'\}$ is contained in L, which is an x-circuit. Replacing C by (T_1, L^{-1}, T_2) does not change the number of non-zero circuits in the flooding. Furthermore, since G is 4-edge-connected, is it possible to "move L off of C" and "onto" a transition of $\mathcal{C}$ which does not include any of h, h', h_1, h_2. The number of non-zero circuits in this new flooding either stays the same or goes down by exactly 2 (because the parity of the number of non-zero circuits stays the same). Furthermore, it can only go down by 2 if C is non-zero, in which case there were at least $\ell + 2$ non-zero circuits in $\mathcal{C}$ to begin with. So in this case, both $\{h, h_1\}$ and $\{h, h_2\}$ work.

Case 2: The transition $\{h_1, h_2\}$ is a transition of a circuit $C' \in \mathcal{C} - \{C\}$.

Then there are trails T_1, T_2, R_1, R_2 so that $C = (T_1, T_2)$, $C' = (R_1, R_2)$, exactly one half-edge from $\{h, h'\}$ is in T_1, and exactly one half-edge from $\{h_1, h_2\}$ is in R_1. Again we have two new floodings to consider; we can replace C and C' by either (T_1, R_2) and (R_1, T_2), or by (T_1, R_1^{-1}) and (R_2^{-1}, T_2). Again, the number of non-zero trails either stays the same or goes down by exactly 2; the latter case can only occur when both C and C' are non-zero. So again, both $\{h, h_1\}$ and $\{h, h_2\}$ work.

In each of the cases, both $\{h, h_1\}$ and $\{h, h_2\}$ work, and so Lemma 5.4.1 holds. □

We are almost ready to state the main proposition about either adding v to A or perturbing v to A. However, since we have quite a bit of setup at this point, it is convenient to give another definition. (We give the graph of a circle-structure a different name than "G" in the following definition because it is convenient to use "G" for the tour graph.) So, an *expanded circle-structure* is a tuple $\mathcal{C} = (\widetilde{G}, C, B, A, v)$ so that

(i) $(\widetilde{G}, C)$ is a prime circle-structure,

(ii) B is a subset of C which induces a comparability grid,

(iii) $A \subseteq V(\widetilde{G}) - C$, and

(iv) $v \in V(\widetilde{G}) - C - A$.

We say that $\mathcal{C}$ *controls* a labelled graph (H, μ) if $V(H) \subseteq B$ and (H, μ) is a complete immersion minor of the labelled tour graph of $(\widetilde{G}, C)$.

Now we can state the main proposition, which will almost complete the proof of the Local Structure Theorem. The proof is fairly long and is broken up into two parts based on the flooding number of an associated RES-graph. However, we believe that it is easier to take care of both parts in one proposition due to the amount of setup required.

Proposition 5.4.2. *There is a function $f_{5.4.2} : \mathbb{Z}^+ \to \mathbb{Z}^+$ so that for any $m \in \mathbb{Z}^+$ and any expanded circle-structure $\mathcal{C} = (\widetilde{G}, C, B, A, v)$ which controls an (A, v)-clean toroidal grid of richness, height, and width at least $f_{5.4.2}(m)$, either*

(i) $\mathcal{C}$ controls an $A \cup \{v\}$-rich grid-like graph of richness, height, and width at least m, or

(ii) there is a circle-structure that is obtained from $(\widetilde{G}, C)$ by perturbing v to A so that the signature of v is $f_{5.4.2}(m)$-small relative to B.

Proof. Let $m \in \mathbb{Z}^+$, and define

$$\ell := 100(m-1)^2 + 5(m-1) + 2 \text{ and}$$
$$\ell' := 64(\ell-1)^2 + 4(\ell-1).$$

Finally, set $f_{5.4.2}(m)$ to be the maximum of $5m(\ell+2)$ and ℓ'.

Now suppose that the conditions of the proposition hold for an expanded circle-structure $\mathcal{C} = (\widetilde{G}, C, B, A, v)$. Let (H, μ) denote the relevant (A, v)-clean toroidal grid, and let N denote its subspace. So N is a subspace of $\mathbb{F}_2^A \times \mathbb{F}_2^v$ of dimension $|A|$, and its dimension does not decrease upon restricting to $\mathbb{F}_2^A$. Furthermore, let (G, λ) denote the labelled tour graph of $(\widetilde{G}, C)$. By performing the relevant shiftings in (G, λ), we may assume that (H, μ) is a shifting-free complete immersion minor of (G, λ) (see Section 3.4 for the definition).

The proof will be split into two cases based on the flooding number of an associated RES-graph. So, let $(\hat{G}, \hat{\lambda}, b)$ be the RES-graph that is obtained from G by labelling each edge in the quotient space $\lambda_A \times \lambda_v / N$ (where each edge e is labelled by the coset of

$\lambda_A(e) \times \lambda_v(e))$, and then identifying all vertices in $V(H)$ to a single new vertex b. Now we split into two cases; the first case is considerably harder than the second.

Case 1: The flooding number of $(\hat{G}, \hat{\lambda}, b)$ is at least ℓ.

In this case we will find, as a complete immersion minor of (G, λ), an $A \cup \{v\}$-rich grid-like graph of richness, height, and width at least m whose vertex-set is contained in $V(H)$. So we might as well assume that (G, λ) is a complete immersion minor-minimal graph so that

(i) (H, μ) is a shifting-free complete immersion minor of (G, λ), and

(ii) the flooding number of the associated RES-graph $(\hat{G}, \hat{\lambda}, b)$ is at least ℓ.

This is an abuse of notation, but we believe that the proof is easiest to read this way (without changing the names of the graphs).

We begin by proving that this new graph is 4-edge-connected using minimality.

Claim 5.4.2.1. *The graph G is 4-edge-connected.*

Proof. The graph G is certainly connected since it completely immerses H. So, going for a contradiction, suppose that G has a 2-edge-cut. By Lemma 5.3.1, the toroidal grid H is 4-edge-connected. So all of $V(H)$ is on one side of the cut; thus there exists $X \subseteq V(G) - V(H)$ so that $|\delta_G(X)| = 2$.

Now consider a trail T whose edge-set is precisely the set of edges with at least one end in X, and whose head and tail are not in X. Such a trail exists, and, moreover, every such trail has the same weight. Consider the labelled graph that is obtained from (G, λ) by deleting X and adding a new edge of weight $\lambda(T)$ whose ends are the ends of T. This graph is a complete immersion minor of (G, λ) and still contains (H, μ) as a shifting-free immersion minor. Moreover, the associated RES-graph still has the same flooding number as before; this contradicts the minimality of (G, λ). □

The next claim will use the Rerouting Lemma (stated as Lemma 5.4.1) to show that most of H is just a subgraph of G.

Claim 5.4.2.2. *All but at most $\ell + 1$ edges of H are also edges of G.*

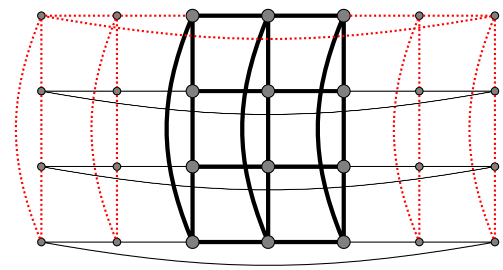

Figure 5.10: The subgraph of H induced by $V(\hat{H})$ (bold) and a depiction of how to find the complete immersion minor (the red dotted edges are the edge-set of a single trail).

Proof. Going for a contradiction, suppose otherwise. Then in any optimal flooding of the associated RES-graph, more than $\ell + 1$ circuits have more than one edge. Equivalently, $|\delta_{\hat{G}}(b)|/2 > \ell + 1$. Furthermore, $\hat{G}$ is 4-edge-connected since G is 4-edge-connected by Claim 5.4.2.1. So we can apply the Rerouting Lemma to find a vertex $x \neq b$ and a half-edge h incident to x so that, in $(\hat{G}, \hat{\lambda}, b)$, every transition at x which includes h is a transition of a flooding with at least ℓ non-zero circuits. However, splitting off one of these transitions in (G, λ) must maintain (H, μ) as a shifting-free immersion minor. This contradicts the minimality of (G, λ). □

Next we find another (A, v)-clean toroidal grid as a shifting-free complete immersion minor of (H, μ); almost all of its edges will also be edges of G.

Claim 5.4.2.3. *There is an (A, v)-clean toroidal grid $(\widetilde{H}, \widetilde{\mu})$ of richness, height, and width at least $5m$ so that*

(i) $(\widetilde{H}, \widetilde{\mu})$ is a shifting-free complete immersion minor of (H, μ), and

(ii) every edge of $\widetilde{H}$ that is not outside-horizontal is also an edge of G.

Proof. Recall from the definition of an (A, v)-clean toroidal grid that there are $5m(\ell + 2)$ disjoint subgrids of H which each have the same height as H and contain the ends of a batch of (H, μ). Then we can "break up H" into $\ell + 2$ disjoint grids, each of which contains $5m$ of those subgrids. By Claim 5.4.2.2, the vertex set of one of those grids, say $\hat{H}$, induces a subgraph in H which is also a subgraph of G. So we are just "missing" the outside-horizontal edges. We can find an appropriate collection edge-disjoint trails of H as depicted in Figure 5.10; thus we obtain the desired (A, v)-clean toroidal grid as a shifting-free complete immersion minor, with underlying grid $\hat{H}$.

□

Now let $(\widetilde{H}, \widetilde{\mu})$ be the (A, v)-clean toroidal grid from Claim 5.4.2.3. The subspace of $(\widetilde{H}, \widetilde{\mu})$ is still N. Consider the RES-graph which is obtained from G by labelling each edge in the quotient space $\lambda_A \times \lambda_v / N$ and identifying all vertices in $V(\widetilde{H})$ (instead of $V(H)$ as before) to a single new vertex, which is the root. If its flooding number is at least m, then (G, λ) completely immerses an $A \cup \{v\}$-rich grid-like graph of richness, height, and width at least m whose vertex-set contained in $V(\widetilde{H}) \subseteq V(H)$. (Each circuit in an optimal flooding "becomes" an outside-horizontal edge of $(\widetilde{H}, \widetilde{\mu})$, and the remaining edges of $(\widetilde{H}, \widetilde{\mu})$ stay the same. Then we can pair each batch of $(\widetilde{H}, \widetilde{\mu})$ with an outside-horizontal edge whose $\lambda_A \times \lambda_v$-weight is not in N.) So, since in this case we would be done, we may assume that the flooding number is less than m. We will obtain a contradiction.

The new RES-graph is is 4-edge-connected since G is 4-edge-connected by Claim 5.4.2.1. So we can apply Corollary 4.1.3 on floodings to obtain the following statement. (We are using the equivalent definition of shiftings in terms of cuts, and the fact that we only care about the weights of edges with both ends in X.) Thus, there exist $X' \subseteq X \subseteq V(G)$ so that X' contains $V(\widetilde{H})$ and

(i) there are at most $4(m-1)$ edges of G with exactly one end in X, and

(ii) there are at most $m-1$ edges of G with both ends in X which are either in $\delta(X')$ and have $\lambda_A \times \lambda_v$-weight in N, or are not in $\delta(X')$ and have $\lambda_A \times \lambda_v$-weight not in N.

Call an edge which satisfies either of these two conditions *N-bad*; so there are at most $5(m-1)$-many N-bad edges.

We will obtain a contradiction to the fact that the flooding number of $(\hat{G}, \hat{\lambda}, b)$ is at least ℓ. The key point is that almost all of $V(H)$ (which was identified to obtain b) is actually in X'.

Claim 5.4.2.4. *All but at most $25(m-1)^2$ vertices of H are in X'.*

Proof. Fix a collection of trails of (G, λ) which shows that it contains (H, μ) as a shifting-free complete immersion minor. Let Y be the set of all $v \in V(H)$ so that there is a path from v to $V(\widetilde{H})$ in the underlying grid of H whose corresponding trails, in G, do not use any N-bad edge. Thus there is a trail from v to $V(\widetilde{H})$ in (G, λ) which does not use any N-bad edge and has $\mathbb{F}_2^A \times \mathbb{F}_2^v$-weight in N. An edge of the trail has $\mathbb{F}_2^A \times \mathbb{F}_2^v$-weight not in N if and only if it is in $\delta(X')$. So in fact $Y \subseteq X'$.

We complete the proof by showing that $|V(H) - Y| \leq 25(m-1)^2$. Observe that H and $\widetilde{H}$ have height and width at least $5m > 5(m-1)$, and the edge-boundary of Y within the underlying grid of H has size at most $5(m-1)$ (since there are at most that many N-bad edges). So by Lemma 5.1.2 about cuts in a grid, either Y or $V(H) - Y$ has size at most $25(m-1)^2$. So the claim follows since $|Y| \geq |V(\widetilde{H})| \geq 25m^2 > 25(m-1)^2$. $\square$

Now consider a collection of trails of (G, λ) which corresponds to an optimal flooding of $(\hat{G}, \hat{\lambda}, b)$. By Claim 5.4.2.4, at most $100(m-1)^2$ of the trails have an end outside of X'. Each of the other trails whose $\lambda_A \times \lambda_v$-weight is not in N must use an N-bad edge. So

$$100(m-1)^2 + 5(m-1) + 2 = \ell \leq \widetilde{\nu}(\hat{G}, \hat{\lambda}, b) \leq 100(m-1)^2 + 5(m-1),$$

a contradiction. This completes the first case. Now we move onto the second and final case.

Case 2: The flooding number of $(\hat{G}, \hat{\lambda}, b)$ is less than ℓ.

In the last case we were working with a certain "minimal" graph (G, λ). Recall that now (G, λ) is just the labelled tour graph of the circle-structure $(\widetilde{G}, C)$, shifted so that it contains (H, μ) as a shifting-free complete immersion minor. So, by Theorem 3.2.3 of Bouchet [5, 8] about prime circle graphs, the graph G is internally 6-edge-connected, and therefore 4-edge-connected. So we can apply Corollary 4.1.3 on floodings.

First, however, consider what it means to shift the quotient labelling that is obtained from (G, λ). Recall that the subspace N of (H, μ) has dimension $|A|$, and that its dimension does not change upon restricting from $\mathbb{F}_2^A \times \mathbb{F}_2^v$ to $\mathbb{F}_2^A$. So the vector which is zero everywhere except for the v-entry is not in N. Thus, equivalently, we can just shift λ_v. So, by Corollary 4.1.3 applied to $(\hat{G}, \hat{\lambda}, b)$, which has flooding number at most $\ell - 1$, there exist a shifting λ'_v of λ_v and a set $X \subseteq V(G)$ which contains $V(H)$ so that

(i) there are at most $4(\ell - 1)$ edges of G with exactly one end in X, and

(ii) there are at most $\ell - 1$ edges of G with both ends in X that have $\lambda_A \times \lambda'_v$-weight not in N.

First we show that almost all of B (recall that this is the set inducing a comparability grid in $\widetilde{G}$) is in X.

Claim 5.4.2.5. *All but at most $16(\ell-1)^2$ vertices of B are in X.*

Proof. Let $n \in \mathbb{Z}^+$ be such that B induces the $n \times n$ comparability grid in $\widetilde{G}$. Since $V(H) \subseteq B$ and H has height and width at least ℓ', we have that $n \geq \sqrt{|V(H)|} \geq \ell'$. So in particular $n \geq 5$ and by Lemma 5.1.1, any tour graph of $\widetilde{G}[B]$ is an $n \times n$ grid-like graph. One of these tour graphs is a complete immersion minor of G. So by Lemma 5.1.2 about cuts in a grid and since $|\delta_G(X)| \leq 4(\ell - 1) < \ell' \leq n$, either $B \cap X$ or $B - X$ has size at most $16(\ell - 1)^2$. So since

$$|B \cap X| \geq |V(H)| \geq (\ell')^2 > 16(\ell - 1)^2,$$

we must have that $|B - X| \leq 16(\ell - 1)^2$, as desired. □

Finally, we show how to perturb v to A.

Claim 5.4.2.6. *There is a circle-structure obtained from $(\widetilde{G}, C)$ by perturbing v to A so that each edge with $\lambda_A \times \lambda'_v$-weight in N has $\mathbb{F}_2^v$-weight zero (in the perturbation).*

Proof. First of all, we may assume that A is non-empty, since otherwise N has dimension 0 and the claim trivially holds. Now recall from Lemma 3.3.1 that for each $a \in A$, if we add a to v then the new signature of v will be $\lambda'_v + \lambda_a$. So we wish to find a vector $x \in \mathbb{F}_2^A$ (that indicates which vertices to add to v) so that the following holds.

(*) For each $\alpha \in \mathbb{F}_2^A$ and $\alpha_v \in \mathbb{F}_2^v$ so that $\alpha \times \alpha_v \in N$, we have $\alpha \cdot x^\intercal = \alpha_v$.

Then the new signature of each edge with $\lambda_A \times \lambda'_v$-weight $\alpha \times \alpha_v \in N$ will be $\alpha_v + \alpha \cdot x^\intercal = 0$, as desired.

We can view (*) as a system of linear inequalities of the form $Mx^\intercal = b$. By the definition of an (A, v)-clean toroidal grid, both M and the matrix $\begin{bmatrix} M \mid b \end{bmatrix}$ have rank $|A|$. It follows that $Mx^\intercal = b$ has a solution, which completes the proof of the claim. □

Now we perturb v to A as in Claim 5.4.2.6 and consider the set of vertices $X \cup B$. Since there are at most $16(\ell - 1)^2$ vertices of B that are not in X by Claim 5.4.2.5, there are

(i) at most $64(\ell - 1)^2 + 4(\ell - 1)$ edges of G with exactly one end in $X \cup B$, and

(ii) at most $64(\ell - 1)^2 + (\ell - 1)$ edges of G whose $\mathbb{F}_2^v$-weight is non-zero (in the perturbation).

Thus, in the perturbation, the signature of v is $f_{5.4.2}(m)$-small relative to B. This completes the proof of Proposition 5.4.2. □

5.5 Completing the proof

In this section we finish the proof of the Local Structure Theorem, which is restated below for convenience.

The Local Structure Theorem (Geelen, McCarty, Wollan)**.** *For any $t \in \mathbb{N}$, there exists $\ell_t \in \mathbb{N}$ so that if (G, C) is a prime circle-structure whose graph does not have all $(t+1)$-vertex graphs as vertex-minors, up to isomorphism, and if B is a subset of C that induces a comparability grid, then there is a rank-t perturbation of (G, C) so that each signature is ℓ_t-small relative to B.*

Proof. We write $f_{5.2.2}$ for the function from Lemma 5.2.2 and $f_{5.3.4}$ and $f_{5.4.2}$ for the functions from Propositions 5.3.4 and 5.4.2, respectively.

Let $t \in \mathbb{N}$; we may assume that $t \geq 1$. We would like to grow the size of a set A for which we have an A-rich grid-like graph as a complete immersion minor. We are happy to give up some richness, height, and width in order to make A larger. In light of Lemma 5.2.2 about finding vertex-minors, we know where we would like to end up; so we begin by defining $m_{t+1} \coloneqq f_{5.2.2}(t+1)$.

Now we define m_i for each $i \in \{0, 1, \ldots, t\}$ by induction on $t - i$. So suppose that $i \leq t$ and we have already defined m_{i+1}. Then define

$$m_i \coloneqq f_{5.3.4}(f_{5.4.2}(m_{i+1}), i).$$

This finishes the inductive definitions. Finally, let ℓ_t be the maximum of 64, $4(m_0 - 1)^2$, and $\max_{i \in \{0,\ldots,t\}} f_{5.4.2}(m_{i+1})$.

So, let (G, C) be a prime circle-structure so that G forbids a $(t+1)$-vertex graph as a vertex-minor, up to isomorphism, and let $B \subseteq C$ be a set which induces an $n \times n$ comparability grid. If $n < 5$ or $n < m_0$, then each signature of the labelled tour graph is already ℓ_t-small relative to B since $4|B| = 4n^2 \leq \ell_t$. So we may assume that $n \geq 5$ and $n \geq m_0$. Thus by Lemma 5.1.1, any tour graph of the $n \times n$ comparability grid is an $n \times n$ grid-like graph. So the labelled tour graph completely immerses an $\emptyset$-rich grid-like graph that has vertex-set B, is arbitrarily rich, and has height and width at least m_0.

Now choose the largest integer $i \in \{0, 1, \ldots, t+1\}$ so that the labelled tour graph completely immerses an A-rich grid-like graph (H, μ) so that A is a subset of $V(G) - C$ of size i, $V(H) \subseteq B$, and (H, μ) has richness, height, and width at least m_i. Such an integer exists since we could take $i = 0$. Furthermore, $i \leq t$ since otherwise Lemma 5.2.2 would imply that G has every graph with vertex-set A as a vertex-minor.

Now consider a vertex $v \in V(G) - C - A$. By Proposition 5.3.4, the graph (H, μ) completely immerses an (A, v)-clean toroidal grid of richness, height, and width at least $f_{5.4.2}(m_{i+1})$. Then (G, C, B, A, v) controls that (A, v)-clean toroidal grid. Then by the maximality of i, we must have the second outcome of Proposition 5.4.2; there is a circle-structure that is obtained from (G, C) by perturbing v to A so that the signature of v is $f_{5.4.2}(m_{i+1})$-small relative to B. Finally, we can also perturb vertices in A (to A) to make their signatures 0-small relative to B. Thus there is a rank-$|A|$ perturbation of (G, C) so that each signature is ℓ_t-small relative to B. This completes the proof of the Local Structure Theorem. □

5.6 Going beyond the local structure

To conclude, let us return to the Structural Conjecture of Geelen from Section 1.5. Consider a graph G of huge rank-width in a proper vertex-minor-closed class. By the Grid Theorem for Vertex-Minors, we can assume (after possibly performing some local complementations) that G contains a large comparability grid as an induced subgraph. Unfortunately, comparability grids are not prime. However, by looking within the comparability grid, it is not hard to reduce to the case that there exist $B \subseteq C \subseteq V(G)$ so that C induces a prime circle graph and B induces a huge comparability grid. We can then apply the Local Structure Theorem to the circle-structure (G, C) and the set B.

In this manner we obtain a low-rank perturbation $(\widetilde{G}, C)$ of (G, C) which makes each signature small relative to B. In an ideal world we could find the "global" structure from the "local" structure using the following approach. First perform the perturbation to obtain $\widetilde{G}$ from G, then grow C, and then repeat this process until $C = V(G)$. We *grow* C by finding a set $\widetilde{C}$ which properly contains C so that $\widetilde{G}[\widetilde{C}]$ is a prime circle graph. There are two main problems with this approach. First, we might get stuck at any one step and not be able to grow C. Second, the perturbations from many different steps might not combine into a single low-rank perturbation of G. (In that case we would want to find a vertex-minor.)

We suspect that the first problem can be avoided for sufficiently rank-connected graphs. (Recall that the Weak Structural Conjecture says that such graphs are low-rank perturbations of circle graphs.) So let us focus on the second problem. We need a better way of keeping track of perturbations. The following definition is similar to the definition of a t-perturbation of a graph, but it behaves much better under local complementation. Two graphs G and $\widetilde{G}$ with the same vertex set are *t-similar* if there exists a graph $\hat{G}$ that has

t more vertices than G (or, equivalently, than $\widetilde{G}$) so that $\hat{G}$ contains both G and $\widetilde{G}$ as vertex-minors. This definition is now invariant under locally complementing in G and $\widetilde{G}$.

Moreover, the proof of the Local Structure Theorem actually gives a circle-structure whose graph $\widetilde{G}$ is t-similar to the original graph G. We believe that we can keep track of the perturbation much better by working in $\hat{G}$ instead of $\widetilde{G}$. With the appropriate conditions, we believe that the set of t new vertices of $\hat{G}$ can be treated like the original set A from the proof presented here. We are optimistic that this approach can eventually lead to a proof of the Structural Conjecture.

References

[1] M. Bonamy and M. Pilipczuk, *Graphs of bounded cliquewidth are polynomially χ-bounded*, Advances in Combinatorics **https://doi.org/10.19086/aic.13668** (2020).

[2] É. Bonnet, C. Geniet, E. J. Kim, S. Thomassé, and R. Watrigant, *Twin-width III: max independent set and coloring*, CoRR **abs/2007.14161** (2020).

[3] É. Bonnet, E. J. Kim, S. Thomassé, and R. Watrigant, *Twin-width I: tractable FO model checking*, 2020 IEEE 61st Annual Symposium on Foundations of Computer Science—FOCS 2020, IEEE Computer Soc., Los Alamitos, CA, 2020, pp. 601–612.

[4] A. Bouchet, *Isotropic systems*, European J. Combin. **8** (1987), no. 3, 231–244.

[5] A. Bouchet, *Reducing prime graphs and recognizing circle graphs*, Combinatorica **7** (1987), no. 3, 243–254.

[6] ______, *Reducing prime graphs and recognizing circle graphs*, Combinatorica **7** (1987), no. 3, 243–254.

[7] A. Bouchet, *Graphic presentations of isotropic systems*, J. Combin. Theory Ser. B **45** (1988), no. 1, 58–76.

[8] A. Bouchet, *Connectivity of isotropic systems*, Combinatorial Mathematics: Proceedings of the Third International Conference (New York, 1985), Ann. New York Acad. Sci., vol. 555, New York Acad. Sci., New York, 1989, pp. 81–93.

[9] A. Bouchet, *An efficient algorithm to recognize locally equivalent graphs*, Combinatorica **11** (1991), no. 4, 315–329.

[10] A. Bouchet, *Circle graph obstructions*, J. Combin. Theory Ser. B **60** (1994), no. 1, 107–144.

[11] A. Bouchet, *Multimatroids. III. Tightness and fundamental graphs*, vol. 22, 2001, Combinatorial geometries (Luminy, 1999), pp. 657–677.

[12] S. Bravyi and R. Raussendorf, *Measurement-based quantum computation with the toric code states*, Phys. Rev. A **76** (2007), 022304.

[13] R. Brignall and D. Cocks, *Uncountably many minimal hereditary classes of graphs of unbounded clique-width*, arXiv:2104.00412 (2021).

[14] R. Brijder and H. J. Hoogeboom, *The algebra of gene assembly in ciliates*, pp. 289–307, Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.

[15] R. Brijder and L. Traldi, *Isotropic matroids I: Multimatroids and neighborhoods*, Electron. J. Combin. **23** (2016), no. 4, Paper 4.1, 41.

[16] R. Brijder and L. Traldi, *Isotropic matroids II: Circle graphs*, Electron. J. Combin. **23** (2016), no. 4, Paper 4.2, 38.

[17] R. Brijder and L. Traldi, *A characterization of circle graphs in terms of multimatroid representations*, Electron. J. Combin. **27** (2020), no. 1, Paper No. 1.25, 35.

[18] R. Brijder and L. Traldi, *A characterization of circle graphs in terms of total unimodularity*, 2020.

[19] J. R. Bunch, L. Kaufman, and Be. N. Parlett, *Hanbook Series Linear Algebra: Decomposition of a symmetric matrix*, Numer. Math. **27** (1976), no. 1, 95–109.

[20] M. Chudnovsky, *The Erdös-Hajnal conjecture—a survey*, J. Graph Theory **75** (2014), no. 2, 178–190.

[21] M. Chudnovsky, Z. Dvořák, T. Klimošová, and P. Seymour, *Immersion in four-edge-connected graphs*, J. Combin. Theory Ser. B **116** (2016), 208–218.

[22] M. Chudnovsky, J. Geelen, B. Gerards, L. Goddyn, M. Lohman, and P. Seymour, *Packing non-zero A-paths in group-labelled graphs*, Combinatorica **26** (2006), no. 5, 521–532.

[23] M. Chudnovsky and S. Oum, *Vertex-minors and the Erdös-Hajnal conjecture*, Discrete Math. **341** (2018), no. 12, 3498–3499.

[24] R. Churchley, *Odd disjoint trails and totally odd graph immersions*, PhD thesis, Simon Fraser University (2017).

[25] A. Collins, J. Foniok, N. Korpelainen, V. Lozin, and V. Zamaraev, *Infinitely many minimal classes of graphs of unbounded clique-width*, Discrete Applied Mathematics (2017).

[26] B. Courcelle, *The monadic second-order logic of graphs. VIII. Orientations*, Ann. Pure Appl. Logic **72** (1995), no. 2, 103–143.

[27] B. Courcelle, J. Makowsky, and U. Rotics, *Linear time solvable optimization problems on graphs of bounded clique-width*, Theory Comput. Syst. **33** (2000), no. 2, 125–150.

[28] B. Courcelle and S. Oum, *Vertex-minors, monadic second-order logic, and a conjecture by Seese*, J. Combin. Theory Ser. B **97** (2007), no. 1, 91–126.

[29] W. H. Cunningham, *Decomposition of directed graphs*, SIAM J. Algebraic Discrete Methods **3** (1982), no. 2, 214–228.

[30] A. Dahlberg and S. Wehner, *Transforming graph states using single-qubit operations*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **376** (2018).

[31] J. Davies, *Vertex-minor-closed classes are χ-bounded*, arXiv:2008.05069 (2020).

[32] J. Davies, *Improved bounds for colouring circle graphs*, arXiv:2107.03585 (2021).

[33] J. Davies and R. McCarty, *Circle graphs are quadratically χ-bounded*, Bulletin of the London Mathematical Society **53** (2021), no. 3, 673–679.

[34] H. de Fraysseix, *Local complementation and interlacement graphs*, Discrete Math. **33** (1981), no. 1, 29–35.

[35] E.D. Demaine, M.T. Hajiaghayi, and K. Kawarabayashi, *Algorithmic graph minor theory: Decomposition, approximation, and coloring*, 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05), 2005, pp. 637–646.

[36] M. Van den Nest, J. Dehaene, and B. De Moor, *Graphical description of the action of local clifford transformations on graph states*, Phys. Rev. A **69** (2004), 022316.

[37] M. DeVos, G. Ding, B. Oporowski, D. P. Sanders, B. Reed, P. Seymour, and D. Vertigan, *Excluding any graph as a minor allows a low tree-width 2-coloring*, J. Combin. Theory Ser. B **91** (2004), no. 1, 25–41.

[38] M. DeVos, J. McDonald, B. Mohar, and D. Scheide, *A note on forbidding clique immersions*, Electron. J. Combin. **20** (2013), no. 3, Paper 55, 5.

[39] R. Diestel, *Graph theory*, fifth ed., Graduate Texts in Mathematics, vol. 173, Springer, Berlin, 2017.

[40] G. Ding, B. Oporowski, J. Oxley, and D. Vertigan, *Unavoidable minors of large 3-connected matroids*, J. Combin. Theory Ser. B **71** (1997), no. 2, 244–293.

[41] G. Durán, L. Grippo, and M. Safe, *Structural results on circular-arc graphs and circle graphs: a survey and the main open problems*, Discrete Appl. Math. **164** (2014), no. part 2, 427–443.

[42] Z. Dvořák and D. Král', *Algorithms for classes of graphs with bounded expansion*, Graph-Theoretic Concepts in Computer Science (Berlin, Heidelberg) (C. Paul and M. Habib, eds.), Springer Berlin Heidelberg, 2010, pp. 17–32.

[43] Z. Dvořák and D. Král', *Classes of graphs with small rank decompositions are χ-bounded*, European J. Combin. **33** (2012), no. 4, 679–683.

[44] Z. Dvořák, *A stronger structure theorem for excluded topological minors*, arXiv:1209.0129 (2012).

[45] Z. Dvořák and T. Klimošová, *Strong immersions and maximum degree*, SIAM J. Discrete Math. **28** (2014), no. 1, 177–187.

[46] Z. Dvořák and P. Wollan, *A structure theorem for strong immersions*, J. Graph Theory **83** (2016), no. 2, 152–163.

[47] J. Edmonds, *Submodular functions, matroids, and certain polyhedra*, pp. 11–26, Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.

[48] A. Ehrenfeucht, T. Harju, I. Petre, D. M. Prescott, and G. Rozenberg, *Formal systems for gene assembly in ciliates*, Theoretical Computer Science **292** (2003), no. 1, 199–219, Selected Papers in honor of Jean Berstel.

[49] P. Erdős and L. Pósa, *On independent circuits contained in a graph*, Canadian J. Math. **17** (1965), 347–352.

[50] P. Erdös, *Graph theory and probability*, Canadian Journal of Mathematics **11** (1959), 34–38.

[51] L. Esperet, *Graph colorings, flows and perfect matchings*, Habilitation thesis, Université Grenoble Alpe (2017), 24.

[52] C. P. Gabor, K. J. Supowit, and W. L. Hsu, *Recognizing circle graphs in polynomial time*, J. Assoc. Comput. Mach. **36** (1989), no. 3, 435–473.

[53] J. Gajarský, P. Hliněný, J. Obdržálek, D. Lokshtanov, and M. S. Ramanujan, *A new perspective on FO model checking of dense graph classes*, ACM Trans. Comput. Log. **21** (2020), no. 4, Art. 28, 23.

[54] J. Gajarský, S. Kreutzer, J. Nešetřil, P. Ossona de Mendez, M. Pilipczuk, S. Siebertz, and S. Toruńczyk, *First-order interpretations of bounded expansion classes*, ACM Trans. Comput. Log. **21** (2020), no. 4, Art. 29, 41.

[55] J. Gajarský, M. Pilipczuk, and S. Toruńczyk, *Stable graphs of bounded twin-width*, arXiv:2107.03711 (2021).

[56] M. R. Garey, D. S. Johnson, G. L. Miller, and C. H. Papadimitriou, *The complexity of coloring circular arcs and chords*, SIAM J. Algebraic Discrete Methods **1** (1980), no. 2, 216–227.

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

[58] B. Geelen, J.and Gerards and G. Whittle, *Solving Rota's conjecture*, Notices Amer. Math. Soc. **61** (2014), no. 7, 736–743.

[59] J. Geelen, personal communication, 2018.

[60] J. Geelen, B. Gerards, and G. Whittle, *Excluding a planar graph from* GF(q)-*representable matroids*, J. Combin. Theory Ser. B **97** (2007), no. 6, 971–998.

[61] J. Geelen, B. Gerards, and G. Whittle, *Tangles, tree-decompositions and grids in matroids*, J. Combin. Theory Ser. B **99** (2009), no. 4, 657–667.

[62] J. Geelen, B. Gerards, and G. Whittle, *The highly connected matroids in minor-closed classes*, Ann. Comb. **19** (2015), no. 1, 107–123.

[63] J. Geelen, J. P. S. Kung, and G. Whittle, *Growth rates of minor-closed classes of matroids*, J. Combin. Theory Ser. B **99** (2009), no. 2, 420–427.

[64] J. Geelen, O. Kwon, R. McCarty, and P. Wollan, *The grid theorem for vertex-minors*, J. Combin. Theory Ser. B (2020).

[65] J. Geelen and E. Lee, *Naji's characterization of circle graphs*, arXiv:1807.10988 (2019).

[66] J. Geelen and S. Oum, *Circle graph obstructions under pivoting*, J. Graph Theory **61** (2009), no. 1, 1–11.

[67] A. M. H. Gerards, *On Tutte's characterization of graphic matroids—a graphic proof*, J. Graph Theory **20** (1995), no. 3, 351–359.

[68] M. Grohe and D. Marx, *Structure theorem and isomorphism test for graphs with excluded topological subgraphs*, SIAM J. Comput. **44** (2015), no. 1, 114–159.

[69] A. Gyárfás, *On the chromatic number of multiple interval graphs and overlap graphs*, Discrete Math. **55** (1985), no. 2, 161–166.

[70] F. Harary, *On the notion of balance of a signed graph*, Michigan Math. J. **2** (1953/54), 143–146 (1955).

[71] M. Hein, W. Dür, J. Eisert, R. Raussendorf, M. Van den Nest, and H. J. Briegel, *Entanglement in graph states and its applications*, **162** (2006).

[72] P. Hliněný, F. Pokrývka, and B. Roy, *FO model checking of geometric graphs*, 12th International Symposium on Parameterized and Exact Computation, LIPIcs. Leibniz Int. Proc. Inform., vol. 89, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2018, pp. Art. No. 19, 12.

[73] J. Jeong, Kim E. J., and S. Oum, *Finding branch-decompositions of matroids, hypergraphs, and more*, arXiv:1711.01381 (2017).

[74] C.D.T. Johnson, *Eulerian digraph immersion*, Ph.D. thesis, Princeton University (2002).

[75] R. Jozsa, *An introduction to measurement based quantum computation*, Quantum Information Processing **199** (2005).

[76] N. Kakimura, K. Kawarabayashi, and Y. Kobayashi, *Packing edge-disjoint odd Eulerian subgraphs through prescribed vertices in 4-edge-connected graphs*, SIAM J. Discrete Math. **31** (2017), no. 2, 766–782.

[77] K. Kawarabayashi and Y. Kobayashi, *Edge-disjoint odd cycles in 4-edge-connected graphs*, J. Combin. Theory Ser. B **119** (2016), 12–27.

[78] K. Kawarabayashi, Y. Kobayashi, and B. Reed, *The disjoint paths problem in quadratic time*, J. Combin. Theory Ser. B **102** (2012), no. 2, 424–435.

[79] R. Kim, O. Kwon, S. Oum, and V. Sivaraman, *Classes of graphs with no long cycle as a vertex-minor are polynomially χ-bounded*, J. Combin. Theory Ser. B **140** (2020), 372 – 386.

[80] A. Kostochka, *Upper bounds on the chromatic number of graphs*, Trudy Inst. Mat **10** (1988), 204–226.

[81] A. Kostochka, *Coloring intersection graphs of geometric figures with a given clique number*, Contemp. Math. **342** (2004), 127–138.

[82] A. Kostochka and J. Kratochvíl, *Covering and coloring polygon-circle graphs*, Discrete Math. **163** (1997), no. 1-3, 299–305.

[83] A. Kotzig, *Eulerian lines in finite* 4*-valent graphs and their transformations*, Theory of Graphs (Proc. Colloq., Tihany, 1966), Academic Press, New York, 1968, pp. 219–230.

[84] O. Kwon and S. Oum, *Graphs of small rank-width are pivot-minors of graphs of small tree-width*, Discrete Appl. Math. **168** (2014), 108–118.

[85] ______, *Scattered classes of graphs*, SIAM J. Discrete Math. **34** (2020), no. 1, 972–999.

[86] E. Lee, *Circle graph obstructions*, Master's thesis, University of Waterloo (2017), http://hdl.handle.net/10012/12309.

[87] R. J. Lipton and R. E. Tarjan, *Applications of a planar separator theorem*, SIAM J. Comput. **9** (1980), no. 3, 615–627.

[88] L. Lovász, *Combinatorial problems and exercises*, North-Holland Publishing Co., Amsterdam-New York, 1979.

[89] W. Mader, *Homomorphieeigenschaften und mittlere Kantendichte von Graphen*, Math. Ann. **174** (1967), 265–268.

[90] W. Mader, *A reduction method for edge-connectivity in graphs*, Ann. Discrete Math. **3** (1978), 145–164.

[91] E. Máčajová and M. Škoviera, *Nowhere-zero flows on signed Eulerian graphs*, SIAM J. Discrete Math. **31** (2017), no. 3, 1937–1952.

[92] E. Máčajová and M. Škoviera, *Odd decompositions of Eulerian graphs*, SIAM J. Discrete Math. **31** (2017), no. 3, 1923–1936.

[93] F. Mazoit and S. Thomassé, *Branchwidth of graphic matroids*, Surveys in combinatorics 2007, London Math. Soc. Lecture Note Ser., vol. 346, Cambridge Univ. Press, Cambridge, 2007, pp. 275–286.

[94] W. Naji, *Reconnaissance des graphes de cordes*, Discrete Math. **54** (1985), no. 3, 329–337.

[95] C. St. J. A. Nash-Williams, *On well-quasi-ordering trees*, Theory of Graphs and its Applications (Proc. Sympos. Smolenice, 1963), Publ. House Czechoslovak Acad. Sci., Prague, 1964, pp. 83–84.

[96] J. Nešetřil and P. Ossona de Mendez, *Sparsity*, Algorithms and Combinatorics, vol. 28, Springer, Heidelberg, 2012, Graphs, structures, and algorithms.

[97] J. Nešetřil and P. Ossona de Mendez, *Structural sparsity*, Uspekhi Mat. Nauk **71** (2016), no. 1(427), 85–116.

[98] H.-T. Nguyen and S. Oum, *The average cut-rank of graphs*, European J. Combin. **90** (2020), 103183, 22.

[99] S. Oum, *Rank-width and vertex-minors*, J. Combin. Theory Ser. B **95** (2005), no. 1, 79–100.

[100] ______, *Rank-width and well-quasi-ordering*, SIAM J. Discrete Math. **22** (2008), no. 2, 666–682.

[101] S. Oum, *Excluding a bipartite circle graph from line graphs*, J. Graph Theory **60** (2009), no. 3, 183–203.

[102] S. Oum, *Rank-width: algorithmic and structural results*, Discrete Appl. Math. **231** (2017), 15–24.

[103] S. Oum, personal communication, 2021.

[104] S. Oum and P. Seymour, *Approximating clique-width and branch-width*, J. Combin. Theory Ser. B **96** (2006), no. 4, 514–528.

[105] J. Oxley, *Matroid theory*, second ed., Oxford Graduate Texts in Mathematics, vol. 21, Oxford University Press, Oxford, 2011.

[106] R. Raussendorf and H. J. Briegel, *A one-way quantum computer*, Phys. Rev. Lett. **86** (2001), 5188–5191.

[107] B. Reed, *Mangoes and blueberries*, Combinatorica **19** (1999), no. 2, 267–296.

[108] N. Robertson and P. Seymour, *Graph minors. V. Excluding a planar graph*, J. Combin. Theory Ser. B **41** (1986), no. 1, 92–114.

[109] N. Robertson and P. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Combin. Theory Ser. B **52** (1991), no. 2, 153–190.

[110] N. Robertson and P. Seymour, *Graph minors. XIII. The disjoint paths problem*, J. Combin. Theory Ser. B **63** (1995), no. 1, 65–110.

[111] N. Robertson and P. Seymour, *Graph minors. XVI. Excluding a non-planar graph*, J. Combin. Theory Ser. B **89** (2003), no. 1, 43–76.

[112] N. Robertson and P. Seymour, *Graph minors. XX. Wagner's conjecture*, J. Combin. Theory Ser. B **92** (2004), no. 2, 325–357.

[113] N. Robertson and P. Seymour, *Graph minors XXIII. Nash-Williams' immersion conjecture*, J. Combin. Theory Ser. B **100** (2010), no. 2, 181–205.

[114] N. Robertson, P. Seymour, and R. Thomas, *Quickly excluding a planar graph*, J. Combin. Theory Ser. B **62** (1994), no. 2, 323–348.

[115] A. Scott and P. Seymour, *A survey of χ-boundedness*, J. Graph Theory **95** (2020), no. 3, 473–504.

[116] L. Traldi, *Notes on a theorem of Naji*, Discrete Math. **340** (2017), no. 1, 3217–3234.

[117] W. T. Tutte, *Lectures on matroids*, J. Res. Nat. Bur. Standards Sect. B **69B** (1965), 1–47.

[118] W. T. Tutte, *Menger's theorem for matroids*, J. Res. Nat. Bur. Standards Sect. B **69B** (1965), 49–53.

[119] W. T. Tutte, *Connectivity in matroids*, Canadian J. Math. **18** (1966), 1301–1324.

[120] M. Van den Nest, J. Dehaene, and B. De Moor, *Graphical description of the action of local clifford transformations on graph states*, Phys. Rev. A **69** (2004), 022316.

[121] M. Van den Nest, W. Dür, G. Vidal, and H. J. Briegel, *Classical simulation versus universality in measurement-based quantum computation*, Phys. Rev. A **75** (2007), 012337.

[122] H. Whitney, *Non-separable and planar graphs*, Trans. Amer. Math. Soc. **34** (1932), no. 2, 339–362.

[123] H. Whitney, *2-Isomorphic Graphs*, Amer. J. Math. **55** (1933), no. 1-4, 245–254.

[124] P. Wollan, *The structure of graphs not admitting a fixed immersion*, J. Combin. Theory Ser. B **110** (2015), 47–66.

[125] D. Zuckerman, *Linear degree extractors and the inapproximability of max clique and chromatic number*, Theory Comput. **3** (2007), 103–128.