The Linkage Problem for Group-labelled Graphs

by

Tony Chi Thong Huynh

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

Waterloo, Ontario, Canada, 2009

Author's Declaration

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

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

Abstract

This thesis aims to extend some of the results of the Graph Minors Project of Robertson and Seymour to “group-labelled graphs”. Let Γ be a group. A Γ-labelled graph is an oriented graph with its edges labelled from Γ, and is thus a generalization of a signed graph.

Our primary result is a generalization of the main result from Graph Minors XIII. For any finite abelian group Γ, and any fixed Γ-labelled graph H, we present a polynomial-time algorithm that determines if an input Γ-labelled graph G has an H-minor. The correctness of our algorithm relies on much of the machinery developed throughout the graph minors papers. We therefore hope it can serve as a reasonable introduction to the subject.

Remarkably, Robertson and Seymour also prove that for any sequence $G_1, G_2, \ldots$ of graphs, there exist indices $i < j$ such that G_i is isomorphic to a minor of G_j. Geelen, Gerards and Whittle recently announced a proof of the analogous result for Γ-labelled graphs, for Γ finite abelian. Together with the main result of this thesis, this implies that membership in *any* minor closed class of Γ-labelled graphs can be decided in polynomial-time. This also has some implications for well-quasi-ordering certain classes of matroids, which we discuss.

Acknowledgements

I would first like to thank my supervisor Jim Geelen for all his inspiration and encouragement. I feel extremely fortunate to have been his student. He provided the impetus as well as several key ideas to this thesis.

I would also like to acknowledge Guoli Ding, Bert Gerards, Thor Johnson, Stephan Kruetzer, Neil Robertson, and Paul Seymour for valuable contributions. In addition, I am very grateful to Bertrand Guenin, Naomi Nishimura, Bruce Richter, and Xingxing Yu for their careful reading and many helpful comments and suggestions.

Finally, thank you to all my family and friends for keeping me going all this time. Most importantly, I would like to thank my mother and father for their unwavering support and love.

for Mom and Dad.

Contents

Chapter 1

Introduction

1.1 History and Motivation

A graph H *is a minor of a graph* G if H can be obtained from a subgraph of G by contracting edges. We say that G has an H*-minor*, if H is isomorphic to a minor of G. This induces a natural ordering $\leq_m$ on the class of all graphs. That is, $H \leq_m G$ if and only if G has an H-minor. It turns out that many interesting graph properties are closed under this ordering, the canonical example being planarity.

The principal aim of this thesis is to extend some of the results of the Graph Minors Project of Robertson and Seymour to "group-labelled graphs".

The Graph Minors Project is widely considered to be the deepest and most important work in graph theory to date. Beginning in 1983, the project has spanned 23 papers. The ingenious methods used to construct and manipulate minors is a *tour de force* of prescient, creative, and disciplined reasoning.

The two main results of the project (as far as we are concerned) appear in Graph Minors XIII [40] and Graph Minors XX [42]. In Graph Minors XX, Wagner's Conjecture is positively settled. That is, Robertson and Seymour prove that (finite) graphs are well-quasi-ordered under taking minors.

Theorem 1.1.1 (Graph Minors Theorem)**.** *For any sequence $G_1, G_2, \dots$ of graphs, there exist indices $i < j$ such that G_i is isomorphic to a minor of G_j.*

Let $\mathcal{G}$ be a minor-closed class of graphs. A graph F is a *forbidden minor* for $\mathcal{G}$, if $F \notin \mathcal{G}$, but every proper minor of F is in $\mathcal{G}$. In this language, Theorem 1.1.1 asserts that every minor-closed class of graphs has a *finite* list of excluded minors. It can thus be viewed as a remarkable generalization of Kuratowski's Theorem.

Theorem 1.1.2 (Kuratowski's Theorem)**.** *A graph is planar if and only if it does not contain a K_5- or $K_{3,3}$-minor.*

Let us mention a few other examples of minor-closed families of graphs.

Example 1.1.3. Let Σ be a surface. Clearly, the class of graphs which embed in Σ is a minor-closed family.

Therefore, an important corollary of the Graph Minors Theorem is a generalized Kuratowski theorem for surfaces.

Corollary 1.1.4. *For any surface Σ, there is a finite set of graphs, $\mathcal{F}(\Sigma)$, such that a graph G embeds in Σ if and only if G does not contain an F-minor, for any $F \in \mathcal{F}(\Sigma)$.*

Example 1.1.5. A graph G is an *apex graph*, if there exists $v \in V(G)$ such that $G \backslash v$ is planar. It is easy to check that the class of apex graphs is minor-closed.

Example 1.1.6. A graph G is *knotless* if G has an embedding in $\mathbb{R}^3$ such that every cycle of G embeds as the *unknot*. That is, every cycle of G bounds a disk in $\mathbb{R}^3$. By performing edge contractions of G in $\mathbb{R}^3$, it is easy to see that the class of knotless graph is a minor-closed family.

The main result in Graph Minors XIII is an algorithmic counterpart to the Graph Minors Theorem. It asserts that for any fixed graph H, we can test if a graph has an H-minor in polynomial-time.

Theorem 1.1.7. *For any graph H, there is a polynomial-time algorithm which determines if an input graph G contains an H-minor.*

We remark that the running time of the algorithm is $O(|V(G)|^3)$, but the constants involved are enormous.

Together, Theorem 1.1.1 and Theorem 1.1.7 imply that there exists an algorithm to test membership in *any* minor-closed class of graphs in cubic-time. In particular, there is a cubic-time algorithm that tests if a graph is knotless.

Corollary 1.1.8. *There exists a cubic-time algorithm, which given any input graph G, correctly determines if G is knotless.*

This latest corollary aptly illustrates the combined utility of the Graph Minors Theorem and Theorem 1.1.7. Previously, there was *no* known algorithm (let alone a polynomial-time one) for testing knotlessness.

We now introduce group-labelled graphs, but we postpone definitions until the next section. A Γ-*labelled graph* is an oriented graph with its edges labelled from a group Γ. In the literature they are also known as *gain graphs* or *voltage graphs*. A $\mathbb{Z}_2$-labelled graph is called a *signed graph*.

Group-labelled graphs are a useful tool for constructing embeddings of graphs on surfaces. For example, they were utilized in the solution to Heawood's famous map-colouring problem by Ringel and Youngs [34]. Also, Zaslavsky [56, 57] showed that group-labelled graphs are connected to several interesting classes of matroids. We will discuss this further in Chapter 2.

In Section 1.3, we define a natural minor relation on the class of Γ-labelled graphs which reduces to the usual minor relation on graphs when Γ is trivial. With respect to this ordering, Geelen, Gerards and Whittle [17] recently announced that Theorem 1.1.1 does indeed extend to Γ-labelled graphs, for Γ finite abelian.

Theorem 1.1.9. *Let Γ be a finite abelian group. For any sequence $G_1, G_2, \ldots$ of Γ-labelled graphs, there exist indices $i < j$ such that G_i is isomorphic to a minor of G_j.*

The main result of this thesis is the extension of Theorem 1.1.7 to Γ-labelled graphs, for Γ finite abelian.

Theorem 1.1.10. *For any finite abelian group Γ and any fixed Γ-labelled graph H, there is a polynomial-time algorithm which determines if an input Γ-labelled graph G contains an H-minor.*

The correctness of our algorithm relies on much of the machinery developed throughout the graph minors papers. Our aim is to present these results in as clear and unified a manner as possible. Indeed, it is an ancillary goal of ours that this thesis may serve as a suitable introduction to the subject.

We end our brief introduction by mentioning that Theorem 1.1.9 and Theorem 1.1.10 fit nicely into the matroid minors project of Geelen, Gerards, and Whittle. It turns out that group-labelled graphs are quite fundamental in understanding the structure of matroids representable over a fixed finite field $\mathbb{F}$. See [15] for a survey of this work.

1.2 Group-labelled Graphs

Let Γ be a group. A Γ-*labelled graph* is an oriented graph together with edge-labels from Γ. To be precise, a Γ-labelled graph G has a *vertex set* $V(G)$ and an *edge set* $E(G)$. Each $e \in E(G)$ is assigned a *head* in $V(G)$, a *tail* in $V(G)$ and a *group-label* in Γ. We denote these as $head_G(e), tail_G(e)$ and $\gamma_G(e)$ respectively. The head and tail of an edge are its *ends*. Let H and G be Γ-labelled graphs. We say that H is a *subgraph* of G if $V(H) \subseteq V(G)$, $E(H) \subseteq E(G)$, and each edge in H has the same head, tail, and group-value as it does in G. We say that G and H are *isomorphic*, if there is a bijection $f : V(G) \cup E(G) \to V(H) \cup E(H)$ such that for all $e \in E(G)$,

- $f(head_G(e)) = head_H(f(e))$,
- $f(tail_G(e)) = tail_H(f(e))$, and
- $\gamma_G(e) = \gamma_H(f(e))$.

We let $\vec{G}$ be the directed graph obtained from G by ignoring the group-labels, and $\widetilde{G}$ be the graph obtained from $\vec{G}$ by ignoring the directions of edges. Since Γ will almost always be abelian, we use additive notation for the group operation. A *walk* in G is a walk in $\widetilde{G}$. Let $e \in E(G)$ and $v \in V(G)$ be an end of e. Define

$$\gamma_G(e, v) := \begin{cases} \gamma_G(e) & \text{if } v = head_G(e) \ , \\ -\gamma_G(e) & \text{otherwise.} \end{cases}$$

Let $W = (v_0, e_1, v_1, e_2, v_2, e_3, \dots, e_k, v_k)$ be a walk in G. The *vertices* of W are $v_0, \dots, v_k$, and the *edges* of W are $e_1, \dots, e_k$. The *group-value*, or just *value*, of W is

$$\gamma_G(W) := \gamma_G(e_1, v_1) + \cdots + \gamma_G(e_k, v_k).$$

The *length of* W is $|W| := k$. W is *closed* if $v_0 = v_k$. W is a cycle if W is closed and $e_1, \dots, e_k, v_1, \dots, v_k$ are distinct. W is a *path* if $v_0, v_1, \dots, v_k$ are distinct. We abuse notation and call v_0 and v_k the *ends* of W, with

$tail_G(W) := v_0$ and $head_G(W) := v_k$. Two walks are *disjoint* if they do not share any vertices. A set of walks $\mathcal{W}$ is *disjoint* if any two members of $\mathcal{W}$ are disjoint.

Let P be a path and let a and b be vertices of P with a occurring before b. We let aP denote the maximal subpath of P starting from a. We let Pb denote the maximal subpath of P ending at b. Finally, we let aPb denote the subpath of P starting at a and ending at b.

Let $e \in E(G)$. We say that G' is obtained from G by *flipping* e if $G' = (G \backslash e) \cup f$, where $head_{G'}(f) = tail_G(e)$, $tail_{G'}(f) = head_G(e)$, and $\gamma_{G'}(f) = -\gamma_G(e)$. Note that flipping an edge does not change the group-value of any walk.

A $\mathbb{Z}_2$-labelled graph will be called a *signed-graph*. Note that the group-value of any path in a signed-graph only depends on the labels of the edges in the path and not on the orientation of those edges.

Next we define an equivalence relation on the class of Γ-labelled graphs. Let $v \in V(G)$ and $\alpha \in \Gamma$. Let G' be the Γ-labelled graph obtained from G by adding α to the label of every edge with head v and subtracting α from the label of every edge with tail v. Note that this operation does not change the group-value of any cycle. We say that G' is obtained from G by *shifting by α at v*. A Γ-labelled graph is *shifting-equivalent* to G if it can be obtained from G via any sequence of shifting operations.

We will need the following elementary lemma.

Lemma 1.2.1. *If G is a Γ-labelled graph and $\widetilde{H}$ is an acyclic subgraph of $\widetilde{G}$, then we can perform shifts so that all edges of H become zero-labelled.*

Proof. It suffices to prove the result when H is a tree. It is helpful to regard H as a rooted tree with root vertex r. We first shift at the neighbours of r in H so that all the edges in H incident to r are zero-labelled. We then regard the neighbours of r in H as new roots and proceed up the tree. □

Let Γ' be a subgroup of Γ. We say that a Γ-labelled graph G is Γ'-*balanced* if $\gamma_G(C) \in \Gamma'$, for all cycles C of G.

Lemma 1.2.2. *If a Γ-labelled graph is Γ'-balanced, then we may perform shifts so that every edge has its group-label in Γ'.*

Proof. According to Lemma 1.2.1, we can shift so that a spanning forest of G is zero-labelled. Since G is Γ'-balanced, it follows that each non-forest edge has its group-label in Γ'. □

Example 1.2.3. Let Γ be a finite group and let $n \in \mathbb{N}$. We define $K(\Gamma, n)$ to be the Γ-labelled graph with vertex set $[n]$ and edge set

$$\{(i, j, \gamma) : i, j \in [n], i \neq j, \gamma \in \Gamma\}.$$

The tail, head, and group-label of (i, j, γ) are i, j, and γ respectively. We call $K(\Gamma, n)$ a Γ-*labelled clique*. Note that $K(\Gamma, n)$ has $2|\Gamma|\binom{n}{2}$ edges.

1.3 Group-labelled Minors

Let G be a Γ-labelled graph and let $e \in E(G)$. The graph $G\backslash e$, is the subgraph of G with vertex set $V(G)$ and edge set $E(G)\backslash\{e\}$. We say that $G\backslash e$ is obtained from G by *deleting* e. If $\gamma_G(e) = 0$, we define the operation of contracting e as follows. Let e have ends u and v. The Γ-labelled graph G/e has edge set $E(G)\backslash\{e\}$ and vertex set $(V(G)\backslash\{u, v\}) \cup \{x_e\}$. For all $f \in E(G/e)$ we define $\gamma_{G/e}(f) := \gamma_G(f)$. Lastly, we set

$$head_{G/e}(f) := \begin{cases} head_G(f) & \text{if } head_G(f) \notin \{u, v\}\ , \\ x_e & \text{otherwise,} \end{cases}$$

and similarly for $tail_{G/e}(f)$. We say that G/e is obtained from G by *contracting* e.

Let H and G be Γ-labelled graphs. We say H *is a minor of* G, if H can be obtained from G via any sequence of the following operations:

- Shifting at a vertex,
- Deleting an edge,
- Contracting a zero-labelled edge,
- Deleting an isolated vertex.

We also say G *has* H *as a minor* in such a case. On the other hand, G *has an* H*-minor* if G has H' as a minor, where H' is isomorphic to H. We write $H \leq_m G$, if G has an H-minor.

Problem 1.3.1. Let H be a Γ-labelled graph. The H*-minor testing problem* is to determine if an input Γ-labelled graph G has an H-minor.

The next easy lemma provides a more global view of minors. We omit the proof.

Lemma 1.3.2. *Let H and G be Γ-labelled graphs with H a minor of G. Then there is a Γ-labelled graph G' (shifting-equivalent to G), and vertex-disjoint trees $\{T_v : v \in V(H)\}$ of G' such that*

- $\gamma_{G'}(e) = 0$ *for all* $v \in V(H)$ *and* $e \in E(T_v)$,
- $head_{G'}(e) \in V(T_{head_H(e)})$ *and* $tail_{G'}(e) \in V(T_{tail_H(e)})$ *for each* $e \in E(H)$, *and*
- $\gamma_{G'}(e) = \gamma_H(e)$, *for each* $e \in E(H)$.

Remark 1.3.3. Let H and G be Γ-labelled graphs, with H a minor of G. Lemma 1.3.2 implies that G is shifting equivalent to a graph G' such that H can be obtained from G' just by deleting edges, contracting zero-labelled edges, and deleting isolated vertices. Furthermore, the order in which we delete or contract edges in G' is irrelevant.

Definition 1.3.4. Let H and G be Γ-labelled graphs with H a minor of G. Let G' and $\{T_v : v \in V(H)\}$ be as given in the previous lemma. For $X \subseteq V(G)$, we say that we can *contract* H *onto* X if $|X| = |V(H)|$, and $X \cap V(T_v) \neq \emptyset$, for each $v \in V(H)$.

Let H and G be Γ-labelled graphs. We say that G has a *topological* H*-minor* if there is an injective map $f : V(H) \to V(G)$ and a family $\mathcal{P} := \{P_e : e \in E(H)\}$ of internally disjoint paths in G such that for each $e \in E(H)$ there exists a path P_e in G such that $head_G(P_e) = f(head_H(e))$, $tail_G(P_e) = f(tail_H(e))$, and $\gamma_G(P_e) = \gamma_H(e)$ for all $e \in E(H)$. We call the pair $(f, \mathcal{P})$ a *model of* H *in* G. If G has a topological H-minor, we write $H \leq_t G$. Note that if $H \leq_t G$, then $H \leq_t G'$, where G' is any Γ-labelled graph obtained from G by shifting at vertices of G that are not in $f(V(H))$.

Problem 1.3.5. Let H be a Γ-labelled graph. The *topological H-minor testing problem* is to determine if an input Γ-labelled graph G has a topological H-minor.

We omit the easy proof of the following lemma.

Lemma 1.3.6. *For any Γ-labelled graph H, there is a finite set $\mathcal{F}_H$ of Γ-labelled graphs, such that $H \leq_m G$ if and only if $F \leq_t G$ for some $F \in \mathcal{F}_H$.*

We remark that for a fixed Γ-labelled graph H, it is straightforward to construct $\mathcal{F}_H$.

1.4 Linkages

We now describe the fundamental problem we are interested in. Let G be a graph. A *pattern* Π in G is a collection of disjoint 2-sets of $V(G)$. Let $\Pi := \{\{s_i, t_i\} : i \in [k]\}$ be a pattern in G. A Π-linkage in G is a collection $\mathcal{P} := \{P_1, \dots, P_k\}$ of disjoint paths in G, such that for all $i \in [k]$, $ends(P_i) = \{s_i, t_i\}$. We refer to $\mathcal{P}$ as a *realization of* Π. The *size* of a pattern Π is simply $|\Pi|$. We call Π a k-pattern if it has size at most k.

Problem 1.4.1. The *k-linkage problem* is given a graph G and a k-pattern Π of G, to determine whether G has a Π-linkage.

Knuth (cf. Karp [24]) showed that if k is part of the input then the k-linkage problem is NP-complete.

We mention that for directed graphs, the natural corresponding problem is NP-complete, even if k is fixed. In fact, Even, Itai, and Shamir [12, 13] showed that the 2-linkage problem for directed graphs is NP-complete. Henceforth, we cease mentioning directed graphs.

A graph G is *k-linked* if *every* k-pattern in G has a realization. Evidently, if G is k-linked, then G is k-connected. On the other hand, Larman and Mani [27] and Jung [22] were the first to show that all graphs of sufficiently high connectivity are k-linked.

Theorem 1.4.2. *For each $k \in \mathbb{N}$, there exists $f(k) \in \mathbb{N}$ such that every $f(k)$-connected graph is k-linked.*

This function has since been substantially improved. Currently, the best bound is due to Thomas and Wollan [53].

Theorem 1.4.3. *If G is a $2k$-connected graph with at least $5k|V(G)|$ edges, then G is k-linked.*

In particular, this implies that every $10k$-connected graph is k-linked.

The first value of k for which the k-linkage problem is interesting is $k = 2$. In fact, there is a beautiful characterization of the 2-linkage problem for graphs, that we would be remiss not to mention. It asserts that in 4-connected graphs, the only obstruction to a 2-linkage problem is topological.

Theorem 1.4.4 (Seymour [48], Shiloach [49], Thomassen [54])**.** *Let G be a 4-connected graph and let $\Pi := \{\{s_1, t_1\}, \{s_2, t_2\}\}$ be a 2-pattern in G. Then G does not have a Π-linkage, if and only if G has an embedding in the plane such that s_1, s_2, t_1, t_2 all appear on the boundary of the outer face of G (in that clockwise order).*

Using any planarity testing algorithm as a subroutine, for example Boyer and Myrvold [3], it is easy to obtain a polynomial-time algorithm that decides the 2-linkage problem.

Amazingly, Robertson and Seymour generalize this result to *any* fixed k. The main result of Graph Minors XIII [40] is that there is a polynomial-time algorithm that solves the k-linkage problem.

Theorem 1.4.5. *Fix $k \in \mathbb{N}$. For any graph G and any k-pattern Π in G, there is a polynomial-time algorithm that determines if G has a Π-linkage.*

As previously noted, the proof of Theorem 1.4.5 relies on much of the theory developed throughout the graph minors papers. Indeed, the correctness of the algorithm hinges upon a lemma whose proof is deferred until Graph Minors XXI [43] and Graph Minors XXII [44].

We will shortly show that Theorem 1.4.5 also yields an algorithm to test for minors. However, it is possible to attack the minor-testing problem directly. In fact, Graph Minors XIII [40] solves a generalization of *both* the k-linkage problem and the minor-testing problem, that runs in $O(|V(G)|^3)$-time. Kawarabayashi, Li, and Reed [25] recently announced an improved $O(|V(G)| \log |V(G)|)$-time algorithm for minor-testing.

Remark 1.4.6. What we are calling the *k-linkage problem* for graphs is usually called the *k-disjoint paths problem*. Typically, the k-linkage problem refers to the topological H-minor testing problem, where H has at most

k edges, and maximum degree 2. The two problems are equivalent, but we prefer to use the term *linkage* since the generalization to group-labelled graphs is less verbose.

Let us turn our attention to group-labelled graphs. Let G be a Γ-labelled graph. We will use much of the same terminology and notation that we introduced for graphs. For example, a *pattern* Π in G is any set of triples of the form (x, y, γ), where x and y are distinct vertices of G, $\gamma \in \Gamma$, and no vertex of G appears in more than one triple of Π.

Let $\Pi = \{(s_i, t_i, \gamma_i) : i \in [k]\}$ be a pattern. The *vertex set* of Π is the set $V(\Pi) := \{s_1, t_1, \ldots, s_k, t_k\}$. The *size* of Π is k. A Π-linkage in G is a set $\mathcal{P} := \{P_i : i \in [k]\}$ of disjoint paths in G such that for all $i \in [k]$, $tail_G(P_i) = s_i$, $head_G(P_i) = t_i$ and $\gamma_G(P_i) = \gamma_i$.

Again, we will call $\mathcal{P}$ a *realization* of Π. Conversely, if $\mathcal{P}$ is a realization of Π, we define the *pattern* of $\mathcal{P}$ to be $\Pi(\mathcal{P}) := \Pi$. The *vertices* and *edges* of $\mathcal{P}$ are defined in the obvious way. Namely,

$$V(\mathcal{P}) := \{v \in V(G) : \ v \in V(P) \text{ for some } P \in \mathcal{P}\},$$

$$E(\mathcal{P}) := \{e \in E(G) : \ e \in E(P) \text{ for some } P \in \mathcal{P}\}.$$

The subgraph $\bigcup \mathcal{P}$ of G will often just be denoted by $\mathcal{P}$, if no confusion is likely to arise. If a pattern Π is of size at most k, we call Π a *k-pattern*.

Let G be a Γ-labelled graph, Π be a pattern in G and $\alpha \in \Gamma$. Let G' be the Γ-labelled graph obtained from G by shifting by α at x. Clearly, if $x \notin V(\Pi)$, then G has a Π-linkage if and only if G' does. Otherwise, define

$$\Pi' := \begin{cases} (\Pi \backslash \{(x, y, \gamma)\}) \cup \{(x, y, \gamma - \alpha)\}, & \text{if } (x, y, \gamma) \in \Pi \\ (\Pi \backslash \{(y, x, \gamma)\}) \cup \{(y, x, \gamma + \alpha)\}, & \text{if } (y, x, \gamma) \in \Pi. \end{cases}$$

Note that G has a Π-linkage if and only if G' has a Π'-linkage.

We will make a few more rudimentary observations. If $e \in E(G)$, and $G \backslash e$ has a Π-linkage, then obviously G has a Π-linkage. On the other hand, consider G/e. If e is zero-labelled and at most one end of e is in $V(\Pi)$, then we can naturally regard Π as a pattern in G/e. With this convention, if G/e has a Π-linkage, then G has a Π-linkage. If both ends of e are in $V(\Pi)$, then by convention G/e does not have a Π-linkage. Similarly, if we delete a vertex in $V(\Pi)$ then the resulting graph does not have a Π-linkage. Therefore, if Π is a pattern, then the class of Γ-labelled graphs that do *not* have a Π-linkage is minor-closed.

We are interested in the following algorithmic problem. Fix a group Γ, and $k \in \mathbb{N}$.

Problem 1.4.7. The (Γ, k)-*linkage problem* is given a Γ-labelled graph G and a k-pattern Π of G, to determine if G has a Π-linkage.

This of course is the natural generalization of Problem 1.4.1. To avoid complex complexity issues we assume that Γ is given to us via its multiplication table. Let us consider some examples.

The simplest example is of course the $(\Gamma, 1)$-linkage problem. That is, let G be a Γ-labelled graph and let $s, t \in V(G)$, $\gamma \in \Gamma$. Does there exist a path P in G from s to t with $\gamma(P) = \gamma$? For graphs, this problem is trivial, but for group-labelled graphs we will show that it is deceptively difficult. For example, as a special case it includes the 2-linkage problem for graphs. To see this let G be a graph and let $\Pi := \{\{s_1, t_1\}, \{s_2, t_2\}\}$ be a 2-pattern in G. Let G' be any $\mathbb{Z}_3$-labelled graph with $\widetilde{G'} = G$ and such that $\gamma_{G'}(e) = 0$ for all $e \in E(G')$. Add a new edge f to G' with $tail_{G'}(f) = t_1$, $head_{G'}(f) = s_2$, and $\gamma_{G'}(f) = 1 \in \mathbb{Z}_3$. Finally, consider the 1-pattern $\Pi' := \{(s_1, t_2, 1)\}$ in $G' \cup \{f\}$. Clearly, G has a Π-linkage if and only if $G' \cup \{f\}$ has a Π'-linkage.

Bert Gerards observed that we can generalize the previous example as follows. If we let $\Gamma := \mathcal{S}_k$, the symmetric group on $[k]$, then it is easy to show that the $(\Gamma, 1)$-linkage problem in $\mathcal{S}_k$-labelled graphs contains the k-linkage problem in graphs. To see this let G be a graph and $\Pi := \{\{s_i, t_i\}, i \in [k]\ \}$ be a k-pattern in G. We let id denote the identity of $\mathcal{S}_k$, and use the convention that group elements are multiplied from left to right. Now, let G' be any $\mathcal{S}_k$-labelled graph such that $\widetilde{G'} = G$ and $\gamma_{G'}(e) = id$ for all $e \in E(G')$. For each $i \in [k-1]$, add a new edge f_i to G' such that $tail_{G'}(f_i) = t_i$, $head_{G'}(f_i) = s_{i+1}$ and $\gamma_{G'}(f_i) = (i\ i+1) \in \mathcal{S}_k$. Let $\gamma := (12 \ldots k) \in \mathcal{S}_k$, and consider the 1-pattern $\Pi' := \{(s_1, t_k, \gamma)\}$ in $G' \cup \{f_i : i \in [k]\}$. Clearly, G has a Π-linkage if and only if $G' \cup \{f_i : i \in [k]\}$ has a Π'-linkage. Of course, this example crucially exploits the fact that $\mathcal{S}_k$ is non-abelian.

The Π-linkage problem in signed graphs also includes the problem of finding disjoint paths in graphs with specified parities. Let G be a graph and let $\Pi := \{\{s_i, t_i\} : i \in [k]\}$ be a pattern in G. Further, let $\{p_i \in \mathbb{Z}_2 : i \in [k]\}$, be a specified set of parities. We may ask if there is a realization $\mathcal{P} := \{P_1, \ldots, P_k\}$ of Π in G, with the additional property that $|P_i| \equiv p_i$ (mod 2), for each $i \in [k]$. Let G' be a $\mathbb{Z}_2$-labelled graph such that

$\widetilde{G'} = G$ and $\gamma_{G'}(e) = 1 \in \mathbb{Z}_2$ for all $e \in E(G)$. Let

$$\Pi' := \{(s_i, t_i, p_i) : i \in [k]\},$$

be the pattern in G' induced by Π and $\{p_i \in \mathbb{Z}_2 : i \in [k]\}$. Clearly, the required paths exist in G if and only if G' has a Π'-linkage.

The main result of this thesis is that for any fixed $k \in \mathbb{N}$ and any finite abelian group Γ, there is a polynomial-time algorithm that decides the (Γ, k)-linkage problem.

Theorem 1.4.8. *Fix $k \in \mathbb{N}$ and Γ a finite abelian group. If Π is a k-pattern of a Γ-labelled graph G, then there is a polynomial-time algorithm that determines if G has a Π-linkage.*

As promised, we now show that Theorem 1.4.8 also yields polynomial-time algorithms for both the topological minor-testing problem and the minor-testing problem for Γ-labelled graphs.

Theorem 1.4.9. *For any finite abelian group Γ, and any fixed Γ-labelled graph H, there is a polynomial-time algorithm that determines if G has a topological H-minor, for any input Γ-labelled graph G.*

Proof. We first define the operation of duplicating vertices. Let G be a Γ-labelled graph, let $v \in V(G)$, and let $E(v)$ be the edges of G incident to v. Let v' be a copy of v and $E'(v) := \{e' : e \in E(v)\}$ be a copy of $E(v)$. Let G' be the Γ-labelled graph with vertex set $V(G) \cup \{v'\}$, and edge set $E(G) \cup E'(v)$, such that $G' \backslash v' = G$ and $G' \backslash v$ is isomorphic to G in the natural way. Namely, the function

$f : V(G) \cup E(G) \to V(G' \backslash \{v\}) \cup E(G' \backslash v)$, such that

$$f(\alpha) := \begin{cases} a', & \text{if } a \in \{v\} \cup E(v) \\ a, & \text{otherwise} \end{cases}$$

is an isomorphism from G to $G' \backslash v$. We say that G' is obtained from G by *duplicating* v.

Now let H be a fixed Γ-labelled graph, and G be an input Γ-labelled graph. Let $f : V(H) \to V(G)$ be an injection and consider $x \in V(H)$. If $deg_H(x) = n$, then we duplicate $f(x)$ ($n-1$ times) in G. Denote the copies of $x := x_1$ as $x_2, \ldots, x_n$. Repeat this for all vertices of H and let G' be the resulting graph. Now for each $e = (u, v, \gamma) \in E(H)$ it is easy to choose indices $i \leq deg_H(u)$ and $j \leq deg_H(v)$ such that the pattern

$$\Pi' := \{(u_i, v_j, \gamma) : (u, v, \gamma) \in E(H)\}$$

has size exactly $E(H)$. Moreover, it is easy to see that there is a model $(f', \mathcal{P})$ of H in G with $f' = f$ if and only if G' has a Π'-linkage. Enumerating over all possible choices of f, and then applying Theorem 1.4.8 gives the desired result. □

Theorem 1.4.8 also proves Theorem 1.1.10, which was our main result from the Introduction.

Theorem 1.1.10. *Let Γ be a finite abelian group, and let H be a fixed Γ-labelled graph. There is a polynomial-time algorithm that tests if $H \leq_m G$ for any input Γ-labelled graph G.*

Proof. Construct the set $\mathcal{F}_H$ given in Lemma 1.3.6 and then apply Theorem 1.4.9. □

We thus focus all our efforts in proving Theorem 1.4.8.

Remark 1.4.10. Let H be a Γ-labelled graph, with Γ finite abelian. We remark that in the proof of Theorem 1.1.10 as a corollary to Theorem 1.4.8, the complexity for H-minor testing is $O(|V(G)|^\alpha)$, where α depends on H. In a subsequent paper, we will show how to generalize the techniques in this thesis to directly obtain an algorithm for H-minor testing, that runs in $O(|V(G)|^\beta)$-time, where β does not depend on H.

Remark 1.4.11. Let G be a Γ-labelled graph, Π be a pattern in G, and G' be a Γ-labelled graph obtained from G by flipping an edge. Note that G has a Π-linkage if and only if G' does. So, henceforth we will not distinguish Γ-labelled graphs that are equivalent up to flipping edges.

1.5 Overview of the Algorithm

In this section we give an informal sketch of the algorithm, to better motivate the reader for some of the technical results that follow. Definitions of unknown terms will be given later.

Let Γ be a finite abelian group, G be a Γ-labelled graph, and Π be a k-pattern in G. We wish to determine whether G has a Π-linkage.

We begin by testing if G has small *branch-width*. If so, then we can solve the problem directly via a theorem from logic, and Chapter 3 describes how to do so.

In the case that G has huge branch-width, the interesting idea is that we do not try to solve the problem directly. Rather, we find a vertex whose deletion does not affect the output. That is, we say that a vertex $v \in V(G)$ is *redundant for* Π provided that G has a Π-linkage if and only if $G \backslash v$ has a Π-linkage. The algorithm finds a redundant vertex, deletes it, and then recurses. Eventually, we reduce to the small branch-width case, where we can solve the problem directly.

Much of our work is therefore dedicated to finding redundant vertices and certifying that they are indeed redundant. This will require various results from graph structure theory. To begin with, since we are in the case that G has huge branch-width, the Grid Theorem (Theorem 5.1.2), implies that the underlying graph $\widetilde{G}$ has a large grid-minor. We can find such a grid J efficiently, and are interested in how the rest of G attaches to J.

We attempt to use the large grid-minor J to find a big clique-minor K in $\widetilde{G}$. We show that we can use such a K to construct a $K(\{0\}, m)$-minor in G, where m is still big. We then try to use this $K(\{0\}, m)$-minor to build a $K(\Gamma', m')$-minor, where Γ' is a subgroup of Γ properly containing $\{0\}$, and m' is still big. We then recurse. If we are lucky, we are able to find a $K(\Gamma, n)$-minor, where n is still big. Big $K(\Gamma, n)$-minors play the same role for Γ-labelled graphs as big clique-minors do for graphs. That is, for Γ-labelled graphs, it is relatively straightforward to find a redundant vertex within a big Γ-labelled clique-minor. The details are given in Chapter 6. If we cannot find a big clique-minor labelled over the full group Γ, then we use a structure theorem for Γ-labelled graphs to find a redundant vertex. This is also handled in Chapter 6.

The remaining case is if our large grid-minor J does not *control* a big clique-minor K in $\widetilde{G}$. In this instance, we use the Graph Minors

Structure Theorem (Theorem 5.3.1), which asserts that $\widetilde{G}$ *essentially embeds* in a surface. Chapter 8 describes how to find a redundant vertex when G is truly embedded in a surface. Chapter 9 sorts out the technical difficulties associated with the essential embedding, namely the *vortices*. The idea is to remove the vortices at the cost of introducing a few more linkage vertices. We can then apply the results from Chapter 8 to find a suitable redundant vertex.

Chapter 2

Matroids

Before delving into the details of our algorithm, we take a brief foray into matroid theory. The main goal is to show that group-labelled graphs encode two natural classes of matroids. For further connections between group-labelled graphs and matroids, see Zaslavsky [56, 57]. Also, many of our later results will be phrased in terms of matroids, so this chapter is quite pertinent. However, our treatment of matroids is rather terse, focusing mainly on their relationship to group-labelled graphs. For a more thorough introduction to matroid theory, please read Oxley's excellent introductory text [30].

2.1 Basics

A *matroid* M consists of a finite ground set $E(M)$ and a *rank function* $r_M : 2^{E(M)} \to \mathbb{Z}$ satisfying

(R0) $0 \leq r_M(X) \leq |X|$, for all $X \subseteq E(M)$

(R1) $r_M(X) \leq r_M(Y)$, for all $X \subseteq Y \subseteq E(M)$

(R2) $r_M(X) + r_M(Y) \geq r_M(X \cap Y) + r_M(X \cup Y)$, for all $X, Y \subseteq E(M)$.

We now give some examples of matroids.

Example 2.1.1. Let $\mathbb{F}$ be a field, let R and E be finite sets, and let $A \in \mathbb{F}^{R \times E}$. For $X \subseteq E$ define $r(X)$ to be the rank of the submatrix of A consisting of the columns indexed by X. It is easy to verify that r is the rank function of a matroid on E. This matroid, denoted $M_{\mathbb{F}}(A)$, is called the *column matroid* of A.

A matroid M is $\mathbb{F}$*-representable* if $M = M_{\mathbb{F}}(A)$ for some A. We say that M is *binary* if it is representable over the binary field $\mathbb{F}_2$, *ternary* if it representable over $\mathbb{F}_3$, and *regular* if it is representable over every field.

Let M be a matroid. A set $X \subseteq E(M)$ is *independent* if $r_M(X) = |X|$, and is *dependent* otherwise. *Bases* are maximal independent sets. *Circuits* are minimal dependent sets. The *closure* of X is the set

$$cl_M(X) := \{x \in E(M) : r_M(X \cup \{x\}) = r_M(X)\}.$$

A *flat* is a set which is equal to its closure.

Example 2.1.2. Let $G = (V, E)$ be a graph. It is easy to check that the collection of edge sets of cycles in G define the circuits of a matroid on E. This matroid, denoted $M(G)$, is the *cycle matroid* of G. A matroid M is *graphic* if $M = M(G)$ for some graph G.

We remark that it is not too difficult to show that all graphic matroids are regular.

Example 2.1.3. Let k and n be non-negative integers with $k \leq n$. We let $\mathcal{U}_{k,n}$ be the matroid whose ground set is $[n]$ and whose independent sets are all subsets of $[n]$ of size at most k. Such a matroid is called a *uniform matroid*.

2.2 Matroid Minors

Let M be a matroid. If D and C are disjoint subsets of $E(M)$, then we define a function $r_{M\backslash D/C}$ on $E(M)\backslash(D \cup C)$ such that

$$r_{M\backslash D/C}(X) := r_M(X \cup C) - r_M(C)$$

for $X \subseteq E(M)\backslash(D \cup C)$.

It is easy to check that $r_{M\backslash D/C}$ is the rank function of a matroid on $E(M)\backslash(C \cup D)$. We denote this matroid as $M\backslash D/C$, and say that $M\backslash D/C$ is a *minor of* M obtained by *deleting* D and *contracting* C. Let M and N be matroids. We say that M has an *N-minor* if N is isomorphic to a minor of M. We write $N \leq_m M$ if M has an N-minor.

A class of matroids $\mathcal{M}$ is *minor-closed* if $N \in \mathcal{M}$ whenever $M \in \mathcal{M}$ and $N \leq_m M$. Naturally, this definition of minors agrees with the usual minor relation on graphs.

Lemma 2.2.1. *Let G be a graph and let $e \in E(G)$. Then $M(G)\backslash e = M(G\backslash e)$ and $M(G)/e = M(G/e)$.*

It follows that the class of graphic matroids is a minor-closed family. For any (possibly infinite) field $\mathbb{F}$, the class of $\mathbb{F}$-representable matroids is also a minor-closed family.

Lemma 2.2.2. *If M is an $\mathbb{F}$-representable matroid and $e \in E(M)$, then both $M\backslash e$ and M/e are $\mathbb{F}$-representable.*

Proof. Suppose $M = M_{\mathbb{F}}(A)$. Evidently, $M\backslash e = M_{\mathbb{F}}(A\backslash e)$, where A is the matrix obtained from A by deleting its eth column. Now if e is the zero column, then $M/e = M_{\mathbb{F}}(A\backslash e)$. If e is not the zero column, then $M/e = M_{\mathbb{F}}(A/e)$, where A/e is the matrix obtained by performing row operations on A until the eth column becomes $[1, 0, \ldots, 0]^T$, and then deleting the first row and eth column. $\square$

Let $\mathcal{M}$ be a minor-closed class of matroids. A matroid N is an *excluded-minor for $\mathcal{M}$* if $N \notin \mathcal{M}$, but every proper minor of N is in $\mathcal{M}$. We end this section by mentioning a beautiful result of Tutte [55] that connects representability and minors.

Theorem 2.2.3. *A matroid is binary if and only if it does not have a $\mathcal{U}_{2,4}$-minor.*

That is, up to isomorphism, $\mathcal{U}_{2,4}$ is the only excluded-minor for the class of binary matroids.

2.3 Dowling Matroids

In this section we define an interesting class of matroids first introduced by Dowling [8]. We begin by providing some motivation.

Let M be a matroid. Recall that a flat of M is a subset of $E(M)$ which is equal to its closure. The set of flats of a matroid, ordered under inclusion, turns out to be a special type of lattice, called a *geometric lattice*. On the other hand, every geometric lattice is the lattice of flats of some (simple) matroid. So, lattices are another way to view matroids. See [50] for the appropriate definitions and proofs.

An example of a geometric lattice is the set of partitions of $[n]$, ordered by refinement, which we denote by P_n. For each finite group Γ and $n \in \mathbb{N}$, Dowling defines a geometric lattice $Q_n(\Gamma)$ of rank n which shares many properties with P_{n+1}. In fact, $Q_n(\Gamma) = P_{n+1}$ when Γ is trivial. Here we are only interested in the special case when Γ is the multiplicative group of a finite field $\mathbb{F}$, in which case $Q_n(\Gamma)$ is $\mathbb{F}$-representable.

Let $\mathbb{F}$ be a finite field. A matroid M is a *Dowling matroid over* $\mathbb{F}$ if $M := M_{\mathbb{F}}(A)$, for some A, where A has at most two non-zero entries per column.

Let $\mathbb{F}^*$ be the multiplicative group of $\mathbb{F}$ and let G be a $\mathbb{F}^*$-labelled graph. We will show that there is a natural Dowling matroid associated with G. Let $A \in \mathbb{F}^{V(G)\times E(G)}$ be defined as follows. If e is a non-loop edge of G with $head_G(e) = u$, $tail_G(e) = v$, and $\gamma_G(e) = \gamma$, then the eth column of A has precisely two non-zero entries $a_{v,e} = 1$ and $a_{u,e} = -\gamma$. If e is a loop with $head_G(e) = tail_G(e) = v$ and $\gamma_G(e) = \gamma$, then the eth column of A has exactly one non-zero entry $a_{v,e} = 1 - \gamma$. We call A the *$\mathbb{F}$-incidence matrix of G*. The *Dowling matroid of G* is $D_{\mathbb{F}}(G) := M_{\mathbb{F}}(A)$.

Lemma 2.3.1. *If G and G' are shifting equivalent $\mathbb{F}^*$-labelled graphs, then $D_{\mathbb{F}}(G) = D_{\mathbb{F}}(G')$.*

Proof. Let G' be obtained from G by shifting by $\alpha \in \mathbb{F}^*$ at $v \in V(G)$. It suffices to show that $D_{\mathbb{F}}(G) = D_{\mathbb{F}}(G')$. Let A and A' be the $\mathbb{F}$-incidence matrices of G and G' respectively. Let T be the set of edges of G with tail v. We define B_1 to be the matrix obtained from A by multiplying the vth row of A by α. We let B_2 be the matrix obtained from B_1 by multiplying each column in T by α^{-1}. We have

$$D_{\mathbb{F}}(G) = M_{\mathbb{F}}(A) = M_{\mathbb{F}}(B_1) = M_{\mathbb{F}}(B_2) = M_{\mathbb{F}}(A') = D_{\mathbb{F}}(G'),$$

as required. □

Lemma 2.3.2. *If H and G are $\mathbb{F}^*$-labelled graphs with H a minor of G, then $D_{\mathbb{F}}(H)$ is a minor of $D_{\mathbb{F}}(G)$.*

Proof. Let G be an $\mathbb{F}^*$-labelled graph, with $\mathbb{F}$-incidence matrix A. If $e \in E(G)$, then evidently $G\backslash e$ has $\mathbb{F}$-incidence matrix $A\backslash e$, where $A\backslash e$ is the matrix obtained from A by deleting the eth column. Thus, for all $e \in E(G)$, $D_{\mathbb{F}}(G\backslash e) = D_{\mathbb{F}}(G)\backslash e$. If $e \in E(G)$ is a loop with $\gamma_G(e) = 1$, then observe that the eth column of A is a zero-column. Thus, G/e also has $\mathbb{F}$-incidence matrix $A\backslash e$, when e is a 1-labelled loop. Finally, if e is a non-loop edge with $head_G(e) = u$, $tail_G(e) = v$, and $\gamma_G(e) = 1$, then note that the eth column of A has exactly two non-zero entries $a_{v,e} = 1$ and $a_{u,e} = -1$. Let A/e be the matrix obtained from A by adding the vth row of A to the uth row of A, and then deleting the eth column and vth row. It is straightforward to verify that G/e has $\mathbb{F}$-incidence matrix A/e. Thus, if e is a 1-labelled edge then $D_{\mathbb{F}}(G/e) = D_{\mathbb{F}}(G)/e$. $\square$

2.4 Lifting Graphic Matroids

Fix $m \in \mathbb{N}$. A matroid M is an *m-lift* of a graphic matroid if

$$M = M_{\mathbb{F}}(\begin{bmatrix} A \\ B \end{bmatrix}),$$

for some field $\mathbb{F}$, where B is the signed incidence matrix of a graph, and A has m rows. In this case we say M is an m-lift of the graphic matroid $M_{\mathbb{F}}(B)$.

Let G be a $\mathbb{F}_2^m$-labelled graph. We now exhibit a binary matroid $L(G)$ associated with G such that $L(G)$ is an m-lift of $M(\widetilde{G})$. Let $B \in \mathbb{F}_2^{V(G)\times E(G)}$ be the incidence matrix of $\widetilde{G}$. Let $A \in \mathbb{F}_2^{m\times E(G)}$ be the matrix whose eth column is the label of e in G. Define

$$L(G) = M_{\mathbb{F}_2}(\begin{bmatrix} A \\ B \end{bmatrix}).$$

We say that $\begin{bmatrix} A \\ B \end{bmatrix}$ is the *$\mathbb{F}_2^m$-incidence matrix of G*, and that $L(G)$ is the *lift matroid of G*.

Lemma 2.4.1. *If G and G' are shifting equivalent $\mathbb{F}_2^m$-labelled graphs, then $L(G) = L(G')$.*

Proof. Suppose that G' is obtained from G by shifting by $\delta \in \mathbb{F}_2^m$ at $v \in V(G)$. By symmetry it will suffice to consider $\delta = (1, 0, \ldots, 0)$. Let C and C' be the $\mathbb{F}_2^m$-incidence matrices of G and G' respectively. Note that C' is obtained from C by adding the vth row of C to the first row of C. $\square$

Lemma 2.4.2. *If H and G are $\mathbb{F}_2^m$-labelled graphs with H a minor of G, then $L(H)$ is a minor of $L(G)$.*

Proof. Let G be a $\mathbb{F}_2^m$-labelled graph with $\mathbb{F}_2^m$-incidence matrix C. Evidently, $L(G\backslash e) = L(G)\backslash e$, for any $e \in E(G)$. If $e \in E(G)$ is a zero-labelled loop, then it is also clear that $L(G/e) = L(G\backslash e)$. Finally if e is a non-loop edge of G with $head_G(e) = u$, $tail_G(e) = v$, and $\gamma_G(e) = 0$, then let C' be the matrix obtained from C by adding vth row to the uth row and then deleting the eth column and vth row. Clearly, C' is the $\mathbb{F}_2^m$-incidence matrix of G/e. Thus, $L(G/e) = L(G)/e$, as required. $\square$

2.5 Well-Quasi-Ordering and Rota's Conjecture

Let $\leq$ be a relation on a set X. We say that $(X, \leq)$ is a *quasi-ordering* if it is reflexive and transitive. For example, the minor-relation on the class of Γ-labelled graphs is clearly a quasi-ordering, as is the minor relation on matroids. An *antichain* is a set of pairwise incomparable elements of X. A quasi-ordering $(X, \leq)$ is a *well-quasi-ordering* if it contains no infinite strictly decreasing chain $x_0 > x_1 > \ldots$, and no infinite antichain. Let $\mathcal{G}$ be the class of all finite graphs, and let $\leq_m$ be the minor relation on $\mathcal{G}$. In this language, the main result of Graph Minors XX [42] asserts that finite graphs are well-quasi-ordered under taking minors.

Theorem 2.5.1. *$(\mathcal{G}, \leq_m)$ is a well-quasi-ordering.*

As previously alluded to, this has been generalized to Γ-labelled graphs, for Γ finite abelian by Geelen, Gerards, and Whittle [17].

Theorem 2.5.2. *Let Γ be a finite abelian group, $\mathcal{G}_\Gamma$ be the class of all Γ-labelled graphs, and $\leq_m$ be the minor relation on $\mathcal{G}_\Gamma$. Then $(\mathcal{G}_\Gamma, \leq_m)$ is a well-quasi-ordering.*

Theorem 2.5.2 yields some nice corollaries. For example, by combining Theorem 2.5.2 with the main result of this thesis, there exists (but we do not know it) an efficient test for membership in *any* minor-closed class of Γ-labelled graphs, for Γ finite abelian.

Corollary 2.5.3. *Let Γ be a finite abelian group, and let $\mathcal{C}$ be any minor-closed class of Γ-labelled graphs. There is a polynomial-time algorithm that, given a Γ-labelled graph G as input, decides if $G \in \mathcal{C}$.*

Proof. Let $\mathcal{C}$ be a minor-closed class of Γ-labelled graphs. By Theorem 2.5.2, $\mathcal{C}$ has a finite set $\mathcal{F}$ of excluded minors. By Theorem 1.1.10, for each $F \in \mathcal{F}$ there is a polynomial-time algorithm to test if G has an F-minor. This clearly yields a polynomial-time algorithm to test if $G \in \mathcal{C}$, namely just test if $F \leq_m G$ for each $F \in \mathcal{F}$. □

Theorem 2.5.2 also has consequences for well-quasi-ordering matroids.

Corollary 2.5.4. *For any finite field $\mathbb{F}$, the class of Dowling matroids over $\mathbb{F}$ is well-quasi-ordered under the minor relation.*

Proof. Immediate from Theorem 2.5.2 and Lemma 2.3.2. □

Corollary 2.5.5. *For any $m \in \mathbb{N}$, the class of binary matroids that are an m-lift of a graphic matroid is well-quasi-ordered under the minor relation.*

Proof. Immediate from Theorem 2.5.2 and Lemma 2.4.2. □

We finish this section by stating two outstanding problems in matroid theory.

Conjecture 2.5.6 (Well-quasi-ordering Conjecture)**.** *For any finite field $\mathbb{F}$ and any sequence $M_1, M_2, \ldots$ of $\mathbb{F}$-representable matroids, there exist indices $i < j$ such that M_i is isomorphic to a minor of M_j.*

The second conjecture was made by Rota [45], and is a vast generalization of Theorem 2.2.3.

Conjecture 2.5.7 (Rota's Conjecture)**.** *For any finite field $\mathbb{F}$, there are, up to isomorphism, only a finite number of excluded-minors for the class of $\mathbb{F}$-representable matroids.*

2.6 Matroid Intersection

We end our chapter on matroids with the Matroid Intersection Theorem, which is a beautiful classical result due to Edmonds [10]. It easily implies several min-max combinatorial relations, including König's Theorem, for example.

Let M_1 and M_2 be two matroids with the same ground set E. A subset X of E is a *common independent set* of M_1 and M_2 if X is independent in M_1 and also in M_2. Let X be a common independent subset of M_1 and M_2 and let $A \subseteq E$. Observe that

$$|X| = |X \cap A| + |X \cap (E \backslash A)| \leq r_1(A) + r_2(E \backslash A).$$

Therefore, the maximum size of a common independent set of M_1 and M_2 is at most $\min_{A \subseteq E}\{r_1(A) + r_2(E \backslash A)\}$. The Matroid Intersection Theorem asserts that equality is always attained.

Theorem 2.6.1 (Matroid Intersection Theorem)**.** *Let M_1 and M_2 be two matroids with the same ground set E. The maximum size of a common independent set of M_1 and M_2 is*

$$\min_{A \subseteq E}\{r_1(A) + r_2(E \backslash A)\}.$$

Proof. See [10] or [30]. □

We will require this theorem at a later juncture.

Chapter 3

Branch-width and Logic

The prime objective of this chapter is to solve the Π-linkage problem over classes of Γ-labelled graphs of bounded "branch-width". Our approach is to encode the Π-linkage problem as a model-checking problem in a certain logic called "monadic second-order logic". We thank Stephan Kreutzer for showing us how to do so. We remark that it is also possible to solve such instances by standard techniques from dynamic programming. However, we choose the logic approach since many difficult graph problems (including NP-hard problems) can be encoded in this way. It is therefore preferable to handle all such problems in a unified manner, via Courcelle's Theorem [5].

3.1 Branch-width

Branch-width is a measure of how tree-like a graph is. We choose to work with branch-width (instead of tree-width), since branch-width can be defined in a more general framework which includes both graphs and matroids as special cases.

Let E be a finite set. A *connectivity function on* E is a function $\lambda : 2^E \to \mathbb{Z}$ satisfying

- $\lambda(X) = \lambda(E - X)$, for all $X \subseteq E$. *(Symmetry)*
- $\lambda(X) + \lambda(Y) \geq \lambda(X \cup Y) + \lambda(X \cap Y)$, for all $X, Y \subseteq E$. *(Submodularity)*

A *connectivity system* is a pair $K = (E, \lambda)$, where λ is a connectivity function on E. We now describe the two connectivity functions that we are principally interested in.

Example 3.1.1. Let M be a matroid with ground set E and rank function r. Define $\lambda_M : 2^E \to \mathbb{Z}$ via

$$\lambda_M(X) = r(X) + r(E - X) - r(E) + 1.$$

It is readily checked that $K_M := (E, \lambda_M)$ is a connectivity system.

Example 3.1.2. Let G be a graph. We define a connectivity function λ_G on $V(G) \cup E(G)$ as follows. Let $A \subseteq V(G) \cup E(G)$. We abuse notation and let A also denote the minimal subgraph of G whose edge set is $A \cap E(G)$ and whose vertex set is $A \cap V(G)$ together with the ends of edges in $A \cap E(G)$. We define $\Lambda_G(A)$ to be the set of vertices in both A and $(V(G) \cup E(G)) \backslash A$ (regarded as subgraphs of G). Finally, we define $\lambda_G(A) := |\Lambda_G(A)|$. It is easy to check that $K_G := (V(G) \cup E(G), \lambda_G)$ is also a connectivity system.

Let $K = (E, \lambda)$ be a connectivity system. A tree is *cubic* if each of its vertices has degree 3 or 1. A *branch-decomposition* of K is a pair (T, f), where T is a cubic tree, with set of leaves $\mathcal{L}$, and f is an injective map from E to $\mathcal{L}$. Let e be an edge of T. Let X be one of the two components of $T - e$. We define the *width* of e, denoted $w(e)$, to be $\lambda(f(X \cap \mathcal{L}))$. Note that $w(e)$ is well defined as λ is symmetric. The *width* of (T, f), denoted $w(T, f)$, is the maximum width of its edges. The *branch-width* of K, denoted $bw(K)$, is the minimum width of all its branch-decompositions.

The branch-width of a matroid or graph is the branch width of their associated connectivity system. The branch-width of a group-labelled graph G is the branch-width of $\widetilde{G}$.

For completeness, we now include a definition of tree-width. Let G be a graph. A *tree-decomposition* of G is a pair (T, W), where T is a tree, and $W := \{W_t : t \in V(T)\}$ is a family of subgraphs of G satisfying

- $\bigcup_{t \in V(T)} W_t = G$, and
- if $t_1, t_2, t_3 \in V(T)$ and t_2 lies on the path of T between t_1 and t_3, then $W_{t_1} \cap W_{t_3} \subseteq W_{t_2}$.

The *width* of (T, W) is $\max\{|V(W_t)| - 1 : t \in V(T)\}$, and the *tree-width* of G, denoted $tw(G)$, is the minimum width of all its tree-decompositions. The *path-width* of G is defined similarly, except that we insist that the tree T is a path.

We remark that branch-width and tree-width are within a constant factor of each other.

Theorem 3.1.3. *For any graph G,*

$$bw(G) \leq tw(G) \leq \frac{3}{2}bw(G).$$

Proof. See Graph Minors X [37, Theorem 5.1] or Richter [33, Lemma 2.7]. □

Hence, a class of graphs has bounded branch-width if and only if it has bounded tree-width.

Bodlaender and Thilikos [2] proved that for any constant ω we can test if a graph has branch with at most w in linear-time.

Theorem 3.1.4. *For any fixed $\omega \in \mathbb{N}$, there is a linear-time algorithm that checks if a graph has branch-width at most ω, and if so, outputs a branch-decomposition of minimum width.*

3.2 Relational Structures

The next few sections are self-contained but are not intended as a comprehensive introduction to mathematical logic. For supplementary details, please see [9].

An *r-ary relation* on a set A is a subset of A^r. A *signature* $\sigma := \{R_1, R_2, \dots\}$ is a finite set of *relation symbols* R_i. Each relation symbol $R \in \sigma$ is assigned an *arity*, $ar(R) \in \mathbb{N}$. A *σ-structure*

$$A := (U(A), R_1(A), \dots, R_n(A))$$

is a tuple consisting of a finite set $U(A)$, the *universe of A*, where each $R_i(A)$ is an r_i-ary relation on $U(A)$, with $r_i := ar(R_i)$.

Here are two examples of relational structures.

Example 3.2.1 (Graphs)**.** Let G be a simple graph. Let $\mathbf{adj}$ be a relation symbol with arity 2. Consider the $\{\mathbf{adj}\}$-structure $A := A(G)$ with universe $U(A) := V(G)$ and $\mathbf{adj}(A) := \{(u, v) : uv \in E(G)\}$. Hence, G is naturally represented as an $\{\mathbf{adj}\}$-structure.

Observe that there is more than one way to represent a graph as a relational structure. For example, we can instead use the universe $V(G) \cup E(G)$, and encode incidences between vertices and edges. For group-labelled graphs, we certainly care about the edge structure and the group-labels, so we use the following description.

Example 3.2.2 (Γ-labelled Graphs)**.** Let G be a Γ-labelled graph. Let $\mathbf{graph}$ be a relation symbol with arity 4. Consider the $\{\mathbf{graph}\}$-structure $A := A(G)$ with universe $U(A) := V(G) \cup E(G) \cup \Gamma$ and

$$\mathbf{graph}(A) := \{(u, v, e, \gamma) : e \in E(G), u = tail_G(e), v = head_G(e), \gamma = \gamma_G(e)\}$$

We can thus regard G as a $\{\mathbf{graph}\}$-structure.

3.3 Monadic Second-Order Logic

Let σ be a signature. A tuple $\bar{\mathbf{x}} := x_1, \dots, x_n$ will be denoted by a boldface letter. We assume a countably infinite set $\{x, y, \dots\}$ of *first-order variables*, and a countably infinite set $\{X, Y, \dots\}$ of *set variables*. By convention, first-order variables are always in lowercase, and set variables in uppercase. We define the class of formulas of *first-order logic over σ*, $\mathrm{FO}[\sigma]$, inductively as follows.

- If a and b are first-order variables, then $a = b$ is in $\mathrm{FO}[\sigma]$.
- If ϕ and τ are both in $\mathrm{FO}[\sigma]$, then so are $(\phi \vee \tau)$, $(\phi \wedge \tau)$, and $\neg\phi$.
- If $R \in \sigma$ with arity r, and $\bar{\mathbf{x}}$ is an r-tuple, then $\bar{\mathbf{x}} \in R$ is in $\mathrm{FO}[\sigma]$.
- If ϕ is in $\mathrm{FO}[\sigma]$, and x is a first-order variable such that neither $\exists x$ nor $\forall x$ appear in ϕ, then both $\exists x\phi$ and $\forall x\phi$ are in $\mathrm{FO}[\sigma]$.

The class of formulas of *monadic second-order logic over σ*, $\mathrm{MSO}[\sigma]$, is an extension of $\mathrm{FO}[\sigma]$, with the following additional rules.

- If x is a first-order variable and X is a set variable, then $x \in X$ is in $\mathrm{MSO}[\sigma]$.
- If X is a set variable and ϕ is in $\mathrm{MSO}[\sigma]$ such that neither $\exists X$ nor $\forall X$ appear in ϕ, then $\exists X \phi$ and $\forall X \phi$ are both in $\mathrm{MSO}[\sigma]$.

Finally, we define the class of formulas of *monadic second-order logic*, MSO, as $\bigcup \mathrm{MSO}[\sigma]$, where the union ranges over all signatures. First-order variables range over elements of σ-structures, and set variables range over sets of elements.

Loosely speaking, monadic second-order logic is a logic that allows quantification over elements and sets of elements. Note that it is quite relevant how we choose to encode a given object as a relational structure. For example, $\mathrm{MSO}[\mathbf{adj}]$ formulas only have quantifications over vertices and subsets of vertices, while $\mathrm{MSO}[\mathbf{graph}]$ formulas have quantifications over subsets of edges (and group elements) as well.

3.4 The Model-Checking Problem

Let $\phi \in \mathrm{MSO}[\sigma]$ and let A be a σ-structure. By interpreting the symbols $=, \neg, \vee, \wedge, \exists, \forall$, and $\in$ in the usual way, we can inductively ascertain if ϕ *is true in* A. For example, $\phi_1 \wedge \phi_2$ is true in A if and only if both ϕ_1 and ϕ_2 are true in A. Similarly, let x be a first-order variable and X be a set variable. We say $\forall x \phi$ is *true in* A, if for all $a \in U(A)$, ϕ is true when we interpret a for x in ϕ. Analogously, $\exists X \phi$ is true in A if there exists a set $S \subseteq U(A)$ such that ϕ is true when we interpret S for X.

A variable x is *free in* ϕ if x occurs in ϕ but neither $\exists x$ nor $\forall x$ do. We will write $\phi(\bar{\mathbf{x}})$ to indicate that the variables in $\bar{\mathbf{x}}$ occur free in ϕ. A formula without a free variable is a *sentence*. If ϕ is a sentence and ϕ is true in A, we write $A \models \phi$. If ϕ has free variables $\bar{\mathbf{x}}$, and $\bar{\mathbf{a}}$ is a tuple of elements from A of the same length as $\bar{\mathbf{x}}$, we write $A \models \phi(\bar{\mathbf{a}})$, if ϕ is true when the variables in $\bar{\mathbf{x}}$ are interpreted by $\bar{\mathbf{a}}$.

Problem 3.4.1. The *model-checking problem* is: given a sentence $\phi \in \mathrm{MSO}$ and a σ-structure A, determine if $A \models \phi$.

Similarly, we can define model-checking for MSO formulas that are not sentences.

Problem 3.4.2. The *evaluation problem* is: given a formula $\phi(\bar{\mathbf{x}}) \in \mathrm{MSO}$, a σ-structure A, and a tuple $\bar{\mathbf{a}}$ of elements from A of the same length as $\bar{\mathbf{x}}$, determine if $A \models \phi(\bar{\mathbf{a}})$.

3.5 Some MSO Formulas

To keep the length of formulas manageable, we will use obvious abbreviations such as $x \neq y$, $\to$, $\bigwedge_{i=1}^{n} \phi_i$, and $\exists_{i=1}^{n} X_i$.

We will also make some less obvious, but still natural, substitutions such as

- Replace $\forall x(x \notin X)$ by $X = \emptyset$.
- Replace $\exists x\big((x \in X) \wedge (x \in Y)\big)$ by $\exists x \in X \cap Y$.
- Replace $\forall z\big((z \in X) \to (z \in Y)\big)$ by $X \subseteq Y$.
- Replace $\forall x\big((x \in X) \to \phi(x)\big)$ by $(\forall x \in X)\phi(x)$.
- Replace $\exists x\Big(\phi(x) \wedge \forall y\big(\phi(y) \to (x = y)\big)\Big)$ by $\exists^{=1} x \phi(x)$.

Finally we make some abbreviations that are particular to formulas in $\mathrm{MSO}[\{\mathbf{graph}\}]$ such as

- Replace $\exists u \exists v \exists \gamma\big((u, v, e, \gamma) \in \mathbf{graph}\big)$ by $e \in \mathrm{edg}$.
- Replace $\exists v \exists \gamma\big((u, v, e, \gamma) \in \mathbf{graph}\big)$ by $(u, e) \in \mathrm{inc}$.
- Replace $\exists u \exists \gamma\big((u, v, e, \gamma) \in \mathbf{graph}\big)$ by $(e, v) \in \mathrm{inc}$.
- Replace $\exists \gamma\big((u, v, e, \gamma) \in \mathbf{graph}\big)$ by $(u, v) = \mathrm{ends}(e)$.
- Replace $\exists u \exists v\big((u, v, e, \gamma) \in \mathbf{graph}\big)$ by $\gamma = \mathrm{lab}(e)$.
- Replace $(\exists e \in F \subseteq \mathrm{edg})\big((e, x) \in \mathrm{inc} \vee (x, e) \in \mathrm{inc}\big)$ by $x \in V(F)$.

We now proceed to describe some MSO formulas. All of the formulas we describe are actually $\mathrm{MSO}[\mathbf{graph}]$ formulas. Let $\phi \in \mathrm{MSO}[\mathbf{graph}]$ with free variables $\bar{\mathbf{x}}$. Let G be a Γ-labelled graph, and let $\bar{\mathbf{a}}$ be a tuple of elements from $V(G) \cup E(G) \cup \Gamma$ of the same length as $\bar{\mathbf{x}}$. By regarding G as a $\{\mathbf{graph}\}$-structure, it makes sense to ask whether $G \models \phi(\bar{\mathbf{a}})$. If so, we say that G *models* $\phi(\bar{\mathbf{a}})$, or that $\phi(\bar{\mathbf{a}})$ *is true in* G.

Formula 1 (Connectedness)**.** The following formula $C(F)$ is true in G if and only if F is a subset of edges of G which induce a connected subgraph of G.

$$(F \subseteq \text{edg}) \wedge \forall X \forall Y \Big(((X \neq \emptyset \neq Y) \wedge (X \cup Y = F)) \rightarrow \exists x \in V(X) \cap V(Y) \Big)$$

Note that we regard the empty set of edges as connected.

Formula 2 (Degree 1 Vertices)**.** The following formula $d_1(F, x)$ is true in G if and only if $F \subseteq E(G)$, $x \in V(G)$, and x is a vertex of degree one in the subgraph of G induced by F.

$$(F \subseteq \text{edg}) \wedge \exists^{=1} e \Big(e \in F \wedge ((x, e) \in \text{inc} \vee (e, x) \in \text{inc}) \Big)$$

Formula 3 (Leaf Edges)**.** The following formula $l(F, e)$ is true in G if and only if $F \subseteq E(G)$ and e is a leaf edge in the subgraph of G induced by F.

$$(e \in F \subseteq \text{edg}) \wedge \exists x \Big(((x, e) \in \text{inc} \vee (e, x) \in \text{inc}) \wedge d_1(F, x) \Big)$$

Formula 4 (Trees)**.** The following formula $T(F)$ is true in G if and only if $F \subseteq E(G)$ and the subgraph of G induced by F is a tree.

$$C(F) \wedge (\forall e \in F) \Big(C(F \backslash e) \rightarrow l(F, e) \Big)$$

Formula 5 (Paths)**.** The following formula $P(F)$ is true in G if and only if F is the set of edges of a path in G.

$$T(F) \wedge \neg \Big(\exists x_1 \exists x_2 \exists x_3 (\bigwedge_{1 \leq i < j \leq 3} (x_i \neq x_j) \wedge \bigwedge_{i=1}^{3} l(F, x_i) \Big)$$

We remark that we regard the empty set of edges as a path.

Formula 6 (Paths with Distinct Ends)**.** The following formula $P(F,x,y)$ is true in G if and only if F is the set of edges of a path in G with distinct ends x and y.

$$P(F) \wedge x \neq y \wedge d_1(F,x) \wedge d_1(F,y).$$

Formula 7 (Paths with Group-Values)**.** There is a formula $P(F,x,y,\gamma)$ which is true in G if and only if F is the set of edges of a path with ends x and y and with group-value γ.

This is the only formula which we do not fully write out, since it is too lengthy to do so. However, we will describe how $P(F,x,y,\gamma)$ is constructed. Recall that we have already encoded paths with ends via the formula $P(F,x,y)$. Thus, $P(F,x,y,\gamma)$ consists of $P(F,x,y) \wedge \tau$, where τ is some MSO formula forcing the path F from x to y to assume the group-value γ. It remains to describe τ. The key idea is that F induces a partition $\mathcal{P}(F) := \{V_\alpha : \alpha \in \Gamma\}$ of $V(F)$, where $u \in V_\alpha$ if and only if the subpath of F from x to u has group-value α. Note that some members of $\mathcal{P}(F)$ may be empty and by convention $x \in V_0$. With respect to $\mathcal{P}(F)$, observe that F has group-value γ if and only if $y \in V_\gamma$. Therefore, it suffices to describe how to construct $\mathcal{P}(F)$, given F. We can do this by quantifying over F and Γ. For each $\alpha \in \Gamma$, and each $e \in F$ we proceed as follows. If the tail, head and group-value of e are u, v, and β respectively, we require $u \in X_\alpha$ if and only if $v \in X_{\alpha+\beta}$.

Lemma 3.5.1. *Let G be a Γ-labelled graph regarded as a $\{\mathbf{graph}\}$-structure. There is an* MSO *formula $\phi(s_1,t_1,\gamma_1,\ldots,s_k,t_k,\gamma_k)$ that is true in G if and only if G has a Π-linkage where $\Pi := \{(s_i,t_i,\gamma_i) : i \in [k]\}$.*

Proof. This is easy given the formulas we have already constructed. Let $P(F,x,y,\gamma)$ be as in Formula 7. Then the required formula $\phi(s_1,t_1,\gamma_1,\ldots,s_k,t_k,\gamma_k)$ is

$$\exists_{i=1}^{k} F_i \Big(\bigwedge_{i=1}^{k} P(F_i,s_i,t_i,\gamma_i) \wedge \bigwedge_{1 \leq i < j \leq k} V(F_i) \cap V(F_j) = \emptyset \Big)$$

$\square$

3.6 Courcelle's Theorem

We are now near the goal we set for ourselves at the beginning of this chapter. That is, we will promptly show that we can efficiently solve linkage problems over any class of Γ-labelled graphs of bounded branch-width. We do this by exploiting a powerful theorem of Courcelle [5], which asserts that for any fixed formula $\phi \in \mathrm{MSO}$, the model-checking problem for ϕ can be solved in linear-time over any class of graphs of bounded branch-width. Actually, we require a mild extension of Courcelle's Theorem.

Theorem 3.6.1 (Arnborg, Stefan, Seese [1]). *Fix $n \in \mathbb{N}$ and $\phi(\bar{\mathbf{x}}) \in \mathrm{MSO}$. If G is a Γ-labelled graph of branch-width at most n, and $\bar{\mathbf{a}}$ is a tuple from $V(G) \cup E(G) \cup \Gamma$ of the same length as $\bar{\mathbf{x}}$ then there is a linear-time algorithm that determines if $G \models \phi(\bar{\mathbf{a}})$.*

Corollary 3.6.2. *Fix $w, k \in \mathbb{N}$ and Γ a finite abelian group. Then for any Γ-labelled graph G of branch-width at most w, and any k-pattern Π in G, there is a linear-time algorithm that determines if G has a Π-linkage.*

Proof. Immediate from Lemma 3.5.1 and Theorem 3.6.1. □

Chapter 4

Tangles

Tangles were first introduced by Robertson and Seymour in Graph Minors X [37]. Roughly speaking, a tangle corresponds to a highly connected portion of a graph. Tangles can also be viewed as a dual notion to branch-width, introduced in Chapter 3. They turn out to be a remarkably effectual idea, and we use them as a unifying framework throughout.

4.1 Basics

Let G be a graph, and let A and B be subgraphs of G. We define $A \cup B$ to be the subgraph of G with vertex set $V(A) \cup V(B)$ and edge set $E(A) \cup E(B)$. We define $A \cap B$ analogously. A *separation* of G is an ordered pair (A, B) of edge-disjoint subgraphs of G with $A \cup B = G$. The *order* of a separation (A, B), denoted $ord(A, B)$, is $|V(A \cap B)|$. The (vertex) *boundary* of (A, B) is $V(A \cap B)$, which we denote $bd(A, B)$. If $X \subseteq V(G)$, we occasionally abuse notation and let X also denote the subgraph of G with vertex set X and no edges. In contrast, $G[X]$ denotes the subgraph of G with vertex X and all edges of G with both ends in X. We say $G[X]$ is the subgraph of G *induced by* X. Lastly, for each subgraph A of G, $G \ominus A$ denotes the subgraph of G with edge set $E(G)\backslash E(A)$ and with vertex set

$$(V(G)\backslash V(A)) \cup \{v \in V(G) : \ v \text{ is an end of an edge in } E(G)\backslash E(A)\}.$$

Thus, $(A, G \ominus A)$ is a separation of G.

A *tangle of order* $n \geq 1$ is a set $\mathcal{T}$ of separations of G, such that

(T1) $ord(A, B) < n$, for each $(A, B) \in \mathcal{T}$;

(T2) if $ord(A, B) < n$, then either $(A, B) \in \mathcal{T}$ or $(B, A) \in \mathcal{T}$;

(T3) if $(A, B) \in \mathcal{T}$, then $V(A) \neq V(G)$;

(T4) if $(A_i, B_i) \in \mathcal{T}$ for each $i \in [3]$, then $A_1 \cup A_2 \cup A_3 \neq G$.

A *tangle* in a group-labelled graph is simply a tangle in the underlying graph. Let $\mathcal{T}$ be a tangle of order n in G. We write $ord(\mathcal{T}) = n$. A subgraph A of G is $\mathcal{T}$*-small* if $(A, B) \in \mathcal{T}$, for some B. On the other hand, A is $\mathcal{T}$*-big* if $(B, A) \in \mathcal{T}$, for some B.

As alluded to earlier, tangles are a dual notion to branch-width. We have the following exact min-max relation [37, Theorem 4.3].

Theorem 4.1.1. *Let G be a graph. The maximum order of a tangle in G is equal to the branch-width of G.*

We now describe a matroid that is naturally associated to a tangle $\mathcal{T}$ of order n in G. For $X \subseteq V(G)$, we let $r_{\mathcal{T}}(X)$ denote the minimum order amongst all separations $(A, B) \in \mathcal{T}$, with $X \subseteq A$. If no such separation exists, we define $r_{\mathcal{T}}(X) = n$. It was first shown in [37] that $r_{\mathcal{T}}$ is indeed the rank function of a matroid on $V(G)$.

Lemma 4.1.2. *Let G be a graph and let $\mathcal{T}$ be a tangle of order n in G. $M_{\mathcal{T}} := (V(G), r_{\mathcal{T}})$ is a matroid.*

Proof. Let $X, Y \subseteq V(G)$. Obviously, $0 \leq r_{\mathcal{T}}(X) \leq r_{\mathcal{T}}(Y)$, for $X \subseteq Y$. If $|X| \geq n$, then evidently $r_{\mathcal{T}}(X) \leq |X|$. Otherwise, consider the separation (X, G). As $ord(X, G) = |X|$, **(T2)** and **(T3)** imply that $(X, G) \in \mathcal{T}$. Therefore, $r_{\mathcal{T}}(X) \leq |X|$ in this case as well. Finally, let us show that $r_{\mathcal{T}}$ is submodular. If $r_{\mathcal{T}}(X) = n$, then

$$\begin{aligned} r_{\mathcal{T}}(X) + r_{\mathcal{T}}(Y) &= n + r_{\mathcal{T}}(Y) \\ &\geq r_{\mathcal{T}}(X \cup Y) + r_{\mathcal{T}}(X \cap Y). \end{aligned}$$

By symmetry, we may now assume that $r_{\mathcal{T}}(X) < n$ and $r_{\mathcal{T}}(Y) < n$. Choose $(A, B) \in \mathcal{T}$ with $X \subseteq A$ and $r_{\mathcal{T}}(X) = ord(A, B)$. Choose $(C, D) \in \mathcal{T}$ with

$Y \subseteq C$ and $r_{\mathcal{T}}(Y) = ord(C, D)$. We have

$$\begin{aligned} r_{\mathcal{T}}(X) + r_{\mathcal{T}}(Y) &= ord(A, B) + ord(C, D) \\ &\geq ord(A \cup C, B \cap D) + ord(A \cap C, B \cup D) \\ &\geq r_{\mathcal{T}}(X \cup Y) + r_{\mathcal{T}}(X \cap Y). \end{aligned}$$

□

We call $M_{\mathcal{T}}$ the *tangle matroid* of G associated to $\mathcal{T}$. A subset X of $V(G)$ is *$\mathcal{T}$-independent*, if it is independent in $M_{\mathcal{T}}$. We abuse terminology and say that a separation (A, B) is *$\mathcal{T}$-independent*, if $(A, B) \in \mathcal{T}$ and $r_{\mathcal{T}}(V(A)) = ord(A, B)$. Since $M_{\mathcal{T}}$ is a matroid, it follows that for each $X \subseteq V(G)$ there is a unique maximal $Y \subseteq V(G)$ such that $r_{\mathcal{T}}(X) = r_{\mathcal{T}}(Y)$. The set Y is of course, simply the matroid closure of X. Hence, we call Y the *$\mathcal{T}$-closure* of X. A set which is equal to its $\mathcal{T}$-closure, is *$\mathcal{T}$-closed*.

Lemma 4.1.3. *If $\mathcal{T}$ is a tangle in G and $X, Y \subseteq V(G)$ are both $\mathcal{T}$-independent, then there are $n := \min\{|X|, |Y|\}$ vertex disjoint paths between X and Y.*

Proof. Suppose not. Then by Menger's theorem, there is a separation (A, B) of G where $X \subseteq V(A)$, $Y \subseteq V(B)$, and $ord(A, B) < n$. So, either $(A, B) \in \mathcal{T}$ or $(B, A) \in \mathcal{T}$. If $(A, B) \in \mathcal{T}$, then X is not $\mathcal{T}$-independent. If $(B, A) \in \mathcal{T}$, then Y is not $\mathcal{T}$-independent. □

Lemma 4.1.4. *If $\mathcal{T}$ is a tangle in G and $(A, B) \in \mathcal{T}$ is $\mathcal{T}$-independent, then the boundary of (A, B) is $\mathcal{T}$-independent.*

Proof. Let Y be the boundary of (A, B) and assume the lemma is false. Then there is a separation $(C, D) \in \mathcal{T}$ such that $Y \subseteq C$ and $ord(C, D) < |Y| = r_{\mathcal{T}}(V(A))$. Therefore,

$$\begin{aligned} 2|Y| &> ord(A, B) + ord(C, D) \\ &\geq ord(A \cup C, B \cap D) + ord(A \cap C, B \cup D) \\ &\geq ord(A \cup C, B \cap D) + |Y|, \end{aligned}$$

where the last inequality follows since $Y \subseteq V((A \cap C) \cap (B \cup D))$. Subtracting $|Y|$ gives

$$|Y| > ord(A \cup C, B \cap D).$$

Thus, either $(A \cup C, B \cap D) \in \mathcal{T}$ or $(B \cap D, A \cup C) \in \mathcal{T}$. If $(B \cap D, A \cup C) \in \mathcal{T}$, then $G = A \cup C \cup (B \cap D)$, contradicting **(T4)**. So, $(A \cup C, B \cap D) \in \mathcal{T}$. However, this contradicts $r_{\mathcal{T}}(V(A)) = |Y|$. □

4.2 Tangle Constructions

Let $\mathcal{T}$ be a tangle of order θ in a graph G. Let $1 \leq \theta' \leq \theta$. Define

$$\mathcal{T}' := \{(A, B) \in \mathcal{T} : ord(A, B) < \theta'\}$$

It follows readily that $\mathcal{T}'$ is a tangle of order θ' in G. We say that $\mathcal{T}'$ is the *truncation of $\mathcal{T}$ to order θ'*.

Let H be a subgraph of G, and let $\mathcal{T}_H$ be a tangle of order θ in H. Let $\mathcal{T}_G$ be the set of all separations (A, B) of G, with $ord(A, B) < \theta$ and such that $(A \cap H, B \cap H) \in \mathcal{T}_H$. It is easy to check that T_G is a tangle of order θ in G.

Let H be a minor of G and let $\mathcal{T}_H$ be a tangle of order θ in H. Define $\mathcal{T}_G$ to be the set of all separations (A, B) of G, with $ord(A, B) < \theta$ and such that $E(A) \cap E(H) = E(A')$ from some $(A', B') \in \mathcal{T}_H$. Again, it is readily checked that $\mathcal{T}_G$ is a tangle of order θ in G.

If H is a subgraph or minor of G, and $\mathcal{T}_H$ is a tangle in H of order θ, we say that $\mathcal{T}_G$ is the tangle *induced by $\mathcal{T}_H$ in G*. If $\mathcal{T}$ is a tangle of order $\theta \geq \theta'$ in G, we say that $\mathcal{T}$ *controls* $\mathcal{T}_H$ if the tangle induced by $\mathcal{T}_H$ in G is a truncation of $\mathcal{T}$.

4.3 Some Tangle Lemmas

Lemma 4.3.1. *Let G be a graph and let $\mathcal{T}$ be a tangle of order $\theta > \theta'$ in G. If X is a subset of $V(G)$ of size θ', then $\mathcal{T}$ controls a tangle $\mathcal{T}'$ of order $\theta - \theta'$ in $G \backslash X$.*

Proof. For disjoint subsets U and W of $V(G)$, we define $G[U, W]$ to be the maximal bipartite subgraph of G with bipartition (U, W). Now for each subgraph A of $G \backslash X$, we define A^+ to be $A \cup G[X] \cup G[V(A), X]$. We let $\mathcal{T}'$ denote the collection of all separations (A, B) of $G \backslash X$ with $ord(A, B) < \theta - \theta'$ and such that $(A^+, B^+) \in \mathcal{T}$. It is easy to check that $\mathcal{T}'$ is a tangle of order $\theta - \theta'$ in $G \backslash X$ and that $\mathcal{T}$ controls $\mathcal{T}'$. □

The next lemma is a modest generalization of Lemma 4.3.1.

Lemma 4.3.2. *Let G be a graph and let $\mathcal{T}$ be a tangle of order $\theta > \theta'$ in G. If X is a subset of $V(G)$ with $r_{\mathcal{T}}(X) = \theta'$, then $\mathcal{T}$ controls a tangle $\mathcal{T}'$ of order $\theta - \theta'$ in $G \backslash X$.*

Proof. Let $(A, B) \in \mathcal{T}$ be a separation in G of order θ' with $X \subseteq V(A)$. Let Y be the boundary of (A, B). By Lemma 4.3.1, $\mathcal{T}$ controls a tangle $\mathcal{T}_1$ of order $\theta - \theta'$ in $G \backslash Y$. Since $(A \backslash Y, B \backslash Y)$ is a separation of order 0 in $\mathcal{T}_1$, it follows that $\mathcal{T}_1$ controls a tangle $\mathcal{T}_2$ of order $\theta - \theta'$ in $B \backslash Y$. Let $\mathcal{T}'$ be the tangle of order $\theta - \theta'$ in $G \backslash X$ induced by $\mathcal{T}_2$. It is easy to check that $\mathcal{T}$ controls $\mathcal{T}'$. $\square$

Lemma 4.3.3. *Let G be a graph and let $\mathcal{T}$ be a tangle of order θ in G. Let (A, B) be a $\mathcal{T}$-independent separation of order θ'. If $\mathcal{T}'$ is a tangle in B of order $\geq \theta'$ that is controlled by $\mathcal{T}$, then the boundary of (A, B) is $\mathcal{T}'$-independent.*

Proof. Let Y be the boundary of (A, B). If Y is not $\mathcal{T}'$-independent, then there is a separation $(C, D) \in \mathcal{T}'$ of B such that $Y \subseteq C$ and $ord(C, D) < |Y| = \theta'$. But then, since $\mathcal{T}$ controls $\mathcal{T}'$ we have $(A \cup C, D) \in \mathcal{T}$, which contradicts the fact that (A, B) is $\mathcal{T}$-independent. $\square$

4.4 A Tangle in a Grid

The $n \times n$ *grid*, denoted G_n, is the graph with vertex set

$$V(G_n) := \{(i, j) : i \in [n], j \in [n]\},$$

where two vertices (i, j) and (i', j') are adjacent if and only if

$$|i - i'| + |j - j'| = 1.$$

The aim of this section is to define a natural tangle $\mathcal{T}_n$ of order n in G_n. For $i \in [n]$, let P_i be the path in G_n with vertex set $\{(i, j) : j \in [n]\}$, and let Q_i be the path in G_n with vertex set $\{(j, i) : j \in [n]\}$.

We let $\mathcal{T}_n$ consists of all separations (A, B) of G_n of order less than n such that $E(A)$ does not contain $E(P_i)$ for any $i \in [n]$. Kleitman and Saks showed $\mathcal{T}_n$ is indeed a tangle; see Graph Minors X [37, Theorem 7.2].

Theorem 4.4.1. *$\mathcal{T}_n$ is a tangle of order n in G_n.*

In the next easy lemma we exhibit an archetypal $\mathcal{T}_n$-independent subset of $V(G_n)$

Lemma 4.4.2. *Let G_n be the $n \times n$ grid and let $\mathcal{T}_n$ be the tangle of order n defined above. Then the set $\{(i,i) : i \in [n]\}$ is independent in $M_{\mathcal{T}_n}$.*

Proof. We call $\{(i,i) : i \in [n]\}$ the *diagonal* of G_n. As above we let P_i be the path in G_n with vertex set $\{(i,j) : j \in [n]\}$. Towards a contradiction assume that the diagonal of G_n is not $\mathcal{T}_n$-independent. That is, there is a separation $(A,B) \in \mathcal{T}_n$ such that $V(A)$ contains the diagonal of G_n. Observe that, for any $i \in [n]$, the diagonal of G_n contains a vertex of P_i, and that $E(A)$ does not contain $E(P_i)$. It follows that the boundary of (A,B) must contain a vertex of P_i for each $i \in [n]$. Thus, $ord(A,B) \geq n$, a contradiction since $ord(\mathcal{T}_n) = n$. $\square$

4.5 A Tangle in a Clique

Let K_n be the complete graph on n vertices. In this section, we show that there is a natural tangle $\mathcal{T}$ of order $\lceil \frac{2n}{3} \rceil$ in K_n, first shown in [37, Theorem 4.4]. We then prove some basic lemmas concerning $\mathcal{T}$.

We define $\mathcal{T}$ such that $(A,B) \in \mathcal{T}$ if and only if $ord(A,B) < \lceil \frac{2n}{3} \rceil$ and $|V(A)| < n$.

Lemma 4.5.1. *$\mathcal{T}$ is a tangle of order $\lceil \frac{2n}{3} \rceil$ in K_n.*

Proof. Evidently, $\mathcal{T}$ satifies **(T1)**, **(T2)**, and **(T3)**. For **(T4)**, let $(A_i, B_i) \in \mathcal{T}$ for $i \in [3]$ with $K_n = A_1 \cup A_2 \cup A_3$. Observe that

$$|V(A_1)| + |V(A_2)| + |V(A_3)| \leq 3 \left\lceil \frac{2n}{3} \right\rceil - 3 \leq 2n - 1.$$

Hence some vertex $v \in K_n$ is in exactly one A_i, say A_1. Since v is only in A_1 and $K_n = A_1 \cup A_2 \cup A_3$, it follows that all the neighbours of v must also be in A_1. Thus, $|V(A_1)| = n$, which contradicts **(T3)**. $\square$

Remark 4.5.2. Let $M_{\mathcal{T}}$ be the tangle matroid on $V(K_n)$ associated with $\mathcal{T}$. Clearly, *any* subset of $V(K_n)$ of size at most $\lceil \frac{2n}{3} \rceil$ is independent in $M_{\mathcal{T}}$. That is, $M_{\mathcal{T}}$ is the uniform matroid on $V(G)$ of rank $\lceil \frac{2n}{3} \rceil$.

The proof of Lemma 4.5.1 also shows that $K(\Gamma, n)$ has a tangle of order $\lceil \frac{2n}{3} \rceil$. So, let G be a Γ-labelled graph which has a $K(\Gamma, n)$-minor, H. Let $\mathcal{T}_H$ be the tangle in H of order $\theta := \lceil \frac{2n}{3} \rceil$ described above. For the rest of this section let $\mathcal{T}$ be the tangle in G induced by H. We now prove some lemmas regarding $\mathcal{T}$. For $X \subseteq V(G)$, we say that we can *contract a Γ'-labelled clique onto X* if G has a $K(\Gamma', |X|)$-minor H, such that we can contract H onto X (recall Definition 1.3.4).

Lemma 4.5.3. *Let G be a Γ-labelled graph and with a $K(\Gamma', n)$-minor, H. Let $\mathcal{T}$ be the tangle of order θ in G induced by H. If $X \subseteq V(G)$ is $\mathcal{T}$-independent, with $|X| < \theta$, then we can contract a Γ'-labelled clique onto X.*

Proof. We proceed via induction on $|V(G)|$. The lemma clearly holds if $|V(G)| = 1$. Consider a counterexample (G, H, X) with $|V(G)|$ minimal. Since G has a $K(\Gamma', n)$-minor, H, we may shift and assume that there exist vertex disjoint trees $\{T_v \mid v \in V(H)\}$ in G such that

- $\gamma_{G'}(e) = 0$ for all $v \in V(H)$ and $e \in E(T_v)$,
- $head_{G'}(e) \in V(T_{head_H(e)})$ and $tail_{G'}(e) \in V(T_{tail_H(e)})$ for each $e \in E(H)$, and
- $\gamma_{G'}(e) = \gamma_H(e)$, for each $e \in E(H)$.

Let $\mathcal{T}$ be the tangle induced by H in G. Let Y be the $\mathcal{T}$-closure of X and consider the separation $(A, B) := (G[Y], G \ominus G[Y])$.

Case 1. *For some $v \in V(H)$ there exists $e = xy \in E(T_v)$ such that $e \in E(B)$ and $\{x, y\} \nsubseteq V(A \cap B)$.*

Let $G' := G/e$. Evidently, G' still has a $K(\Gamma', n)$-minor, H'. Let $\mathcal{T}'$ be the tangle in G' induced by H'.

Claim. *X is $\mathcal{T}'$-independent.*

SUBPROOF. Towards a contradiction suppose X is not $\mathcal{T}'$-independent. Then there is a separation (C', D') in G' such that

- $(C', D') \in \mathcal{T}'$.
- $X \subseteq V(C')$.
- $ord_{G'}(C', D') < |X|$

Therefore, by uncontracting e, there is a separation (C, D) in G such that

- $(C, D) \in \mathcal{T}$.
- $X \subseteq V(C)$.
- $ord_G(C, D) \leq |X|$.

Note that in fact $ord_G(C, D) = |X|$, and that we may assume $e \in E(C)$. By submodularity, we have that

$$\begin{aligned} 2|X| &= ord_G(A, B) + ord_G(C, D) \\ &\geq ord_G(A \cap C, B \cup D) + ord_G(A \cup C, B \cap D) \end{aligned}$$

First suppose $ord_G(A \cup C, B \cap D) > |X|$. Then $ord_G(A \cap C, B \cup D) < |X|$. Thus, $(A \cap C, B \cup D) \in \mathcal{T}$ or $(B \cup D, A \cap C) \in \mathcal{T}$. If $(B \cup D, A \cap D) \in \mathcal{T}$, then $G = A \cup (B \cup D)$, a contradiction. So, $(A \cap C, B \cup D) \in \mathcal{T}$. But $X \subseteq A \cap C$, which contradicts the fact that X is $\mathcal{T}$-independent.

Therefore, $ord_G(A \cup C, B \cap D) \leq |X|$. It now follows that $(A \cup C, B \cap D) \in \mathcal{T}$, for otherwise, G is the union of A, C, and $B \cap D$, each of which is $\mathcal{T}$-small. However, $(A \cup C, B \cap D)$ contradicts the choice of (A, B). So, X is indeed $\mathcal{T}'$-independent, proving the claim. $\square$

By induction, we can contract a Γ'-labelled clique onto X in G/e and hence also in G. This completes Case 1.

Case 2. *For all $v \in V(H)$ and all $e = xy \in E(T_v \cap B)$, $\{x, y\} \subseteq V(A \cap B)$.*

In this case we consider a tree T_v to be *small* if $|V(T_v \cap B)| = 1$, and *big* otherwise.

Claim. *There are at least $|X|$ small trees T_v.*

SUBPROOF. We are assuming that each big tree T_v satisfies $V(T_v \cap B) \subseteq V(A \cap B)$. Therefore, there are at most $|V(A \cap B)|/2 = |X|/2$ big trees. Thus, there are at least

$$n - \frac{|X|}{2} = \left(\frac{3}{2}\right)\left(\frac{2n}{3}\right) - \frac{|X|}{2} \geq \frac{3|X|}{2} - \frac{|X|}{2} = |X|$$

small trees. □

Let

$$Z = \bigcup_{T_v \text{ is small}} V(T_v \cap B)$$

By construction $|Z| \geq |X|$ and $G[Z] = K(\Gamma', |Z|)$. Also, as X and Z are both $\mathcal{T}$-independent, by Lemma 4.1.3 there are $|X|$ vertex-disjoint paths in G between X and Z. By shifting (at vertices not in Z), we may assume that the edges of these paths are all zero-labelled. Thus, we can clearly contract a Γ'-labelled clique onto X, as claimed. This completes the second case and hence the proof. □

4.6 Tangles in Connectivity Systems

In this section we define tangles for arbitrary connectivity systems. Let $K = (E, \lambda)$ be a connectivity system. A *separation* is an ordered partition (A, B) of E. The *order* of (A, B) is defined to be $\lambda(A)$ (which is equal to $\lambda(B)$). A *tangle* $\mathcal{T}$ *of order* n is a collection of subsets of E satisfying

(T1) $\lambda(A) < n$ for all $A \in \mathcal{T}$;

(T2) if $\lambda(A) < n$, then either $A \in \mathcal{T}$ or $E \backslash A \in \mathcal{T}$;

(T3) $E \backslash \{e\} \notin \mathcal{T}$, for all $e \in E$;

(T4) if $A_i \in \mathcal{T}$ for $i \in [3]$, then $A_1 \cup A_2 \cup A_3 \neq E$.

In [16], it is proved that Theorem 4.1.1 extends to arbitrary connectivity systems, although this is implicit in [37].

Theorem 4.6.1. *Let K be a connectivity system. The maximum order of a tangle in K is equal to the branch-width of K.*

Let $K = (E, \lambda)$ be a connectivity system and let $\mathcal{T}$ be a tangle of order n in K. For $X \subseteq E$, we let $r_{\mathcal{T}}(X)$ denote the minimum order amongst all separations $(A, B) \in \mathcal{T}$, with $X \subseteq A$. If no such separation exists, we define $r_{\mathcal{T}}(X) = n$. The following lemma has the same proof as Lemma 4.1.2.

Lemma 4.6.2. *$M_{\mathcal{T}} := (E, r_{\mathcal{T}})$ is a matroid.*

We also call M the *tangle matroid* of K associated with $\mathcal{T}$.

Remark 4.6.3. Lemma 4.6.2 can be viewed as a generalization of Lemma 4.1.2 as follows. Let G be a graph and let $K = (V(G) \cup E(G), \lambda_G)$ be the connectivity function described in Example 3.1.2. Let $\mathcal{T}$ be a tangle of order n in K. Then the restriction of $M_{\mathcal{T}}$ to $V(G)$ is the usual tangle matroid given in Lemma 4.1.2.

Let $K = (E, \lambda_K)$ and $K' = (E, \lambda_{K'})$ be connectivity systems. We say that K' is a *tie-breaker for K* if for all $X, Y \subseteq E$,

1. $\lambda_{K'}(X) \neq \lambda_{K'}(Y)$, unless $X = Y$ or $X = E \backslash Y$.
2. If $\lambda_K(X) < \lambda_K(Y)$, then $\lambda_{K'}(X) < \lambda_{K'}(Y)$.

The following was proved in [16, Lemma 9.2].

Lemma 4.6.4. *Every connectivity system $K = (E, \lambda_K)$ has a tie-breaker.*

Proof. We may assume $E = [n]$. We first define $\lambda_L : 2^E \to \mathbb{N}$ as

$$\lambda_L(X) := \begin{cases} \sum_{x \in X} 2^x & \text{if } X \subseteq [n-1], \\ \lambda_L(E \backslash X) & \text{if } n \in X. \end{cases}$$

We claim that λ_L is a connectivity function on E. It is trivially symmetric from its definition. Let $X, Y \subseteq E$. We must show

$$\lambda_L(X) + \lambda_L(Y) \geq \lambda_L(X \cup Y) + \lambda_L(X \cap Y).$$

This clearly holds with equality if both $X, Y \subseteq [n-1]$. Let us now consider the case $n \in X \backslash Y$. By definition,

$$\begin{aligned} \lambda_L(X) + \lambda_L(Y) &= \sum_{i \in Y \backslash X} 2^{i+1} + \sum_{i \in X \cap Y} 2^i + \sum_{i \in E \backslash (X \cup Y)} 2^i \\ &\geq \sum_{i \in X \cap Y} 2^i + \sum_{i \in E \backslash (X \cup Y)} 2^i \\ &= \lambda_L(X \cup Y) + \lambda_L(X \cap Y). \end{aligned}$$

The remaining case $n \in X \cap Y$ is even easier, so we omit it.

We now define $\lambda_{K'}(X) := \lambda_L(X) + 2^n \lambda_K(X)$, for all $X \subseteq E$. Since $\lambda_{K'}$ is the sum of two connectivity functions, it is a connectivity function. Moreover, it is easy to verify that it is indeed a tie-breaker for K. $\square$

We remark that if $\mathcal{T}$ is a tangle in a connectivity function K, then $\mathcal{T}$ is also a tangle in any tie-breaker K' for K.

4.7 Tree-decompositions and Laminar Families

Let E be a finite set. A *separation of* E is an ordered partition (A, B) of E. Two separations (A_1, B_1) and (A_2, B_2) *cross*, if $A_i \cap B_j \neq \emptyset$ for all $i, j \in [2]$. A *laminar family* is a collection $\mathcal{L}$ of separations of E such that no two separations in $\mathcal{L}$ cross.

A *tree-decomposition* of E is a pair $(T, \mathcal{S})$, where T is a tree and $\mathcal{S} := \{S_v : v \in V(T)\}$ is a partition of E. For $X \subseteq E$ we define $\mathcal{S}(X) := \bigcup_{x \in X} S_x$. For all $e \in E(T)$, notice that $T \backslash e$ has two components T_1 and T_2. The *separation of* E *displayed by* e is defined to be $(\mathcal{S}(V(T_1)), \mathcal{S}(V(T_2)))$.

The next two lemmas are due to Edmonds and Giles [11].

Lemma 4.7.1. *If $(T, \mathcal{S})$ is a tree-decomposition of E, then the collection of all separations of E displayed by $(T, \mathcal{S})$ is a laminar family.*

Conversely, every laminar family arises from a tree-decomposition.

Lemma 4.7.2. *Let $\mathcal{L}$ be a laminar family with ground set E. Then there is a tree-decomposition $(T, \mathcal{S})$ of E such that the collection of all separations of E displayed by T is precisely $\mathcal{L}$.*

Let $K = (E, \lambda)$ be a connectivity system. A subset $A \subseteq E$ is *robust* if for every separation (A_1, A_2) of A, $\lambda(A_1) > \lambda(A)$ or $\lambda(A_2) > \lambda(A)$. A separation (A, B) of E is *robust* if both A and B are robust. It turns out that the collection of all robust separations of E is a laminar family.

Lemma 4.7.3. *Let $K = (E, \lambda)$ be a connectivity system. The collection of all robust separations of E is a laminar family.*

Proof. See [16, Lemma 8.3]. $\square$

4.8 The Tree of Tangles

Let $K = (E, \lambda)$ be a connectivity system and let $\mathcal{T}_1$ and T_2 be tangles in K. We say that a separation (A, B) *distinguishes* $\mathcal{T}_1$ and $\mathcal{T}_2$, if $A \in \mathcal{T}_1$ and $B \in \mathcal{T}_2$. Tie-breakers are a convenient way to choose canonical separations that distinguish tangles. Recall that $\mathcal{T}_1$ is a *truncation* of $\mathcal{T}_2$, if $\mathcal{T}_1 \subseteq \mathcal{T}_2$.

Lemma 4.8.1. *Let $K = (E, \lambda)$ be a connectivity system, and let $\mathcal{T}_1$ and $\mathcal{T}_2$ be tangles of K that are incomparable by truncation. Let $K' = (E, \lambda')$ be a tie-breaker for K. Among all separations which distinguish $\mathcal{T}_1$ and $\mathcal{T}_2$, let (A, B) be the separation of minimum K'-order. Then (A, B) is robust.*

Proof. See [16, Lemma 9.3]. □

From here, it is not too difficult to show that every connectivity system has a canonical tree-decomposition which we call its "tree of tangles".

Theorem 4.8.2. *Let $K = (E, \lambda)$ be a connectivity system and let $\{\mathcal{T}_1, \dots, \mathcal{T}_n\}$ be a collection of tangles of K, which is pairwise incomparable by truncation. Then there is a tree decomposition $(T, \mathcal{S})$ of E such that*

- *$V(T) = [n]$,*
- *for each $i \in V(T)$, and $e \in E(T)$, if T' is the component of $T \backslash e$ containing i, then $\mathcal{S}(V(T')) \notin \mathcal{T}_i$, and*
- *For all distinct $i, j \in [n]$, there is a minimum order separation distinguishing $\mathcal{T}_i$ and $\mathcal{T}_j$ that is displayed by T.*

A tangle that corresponds to a leaf in the tree of tangles will be called a *peripheral tangle*. We will see that when applying structure theory results, it is more advantageous to work with peripheral tangles, rather than arbitrary ones. Therefore, a key subroutine of the algorithm will be dedicated to finding peripheral tangles.

A very attractive feature of tangles, as opposed to other certificates of high branch-width such as *brambles* (see Reed [32]), is that every connectivity system only has a few maximal tangles.

Corollary 4.8.3. *Let $K = (E, \lambda)$ be a connectivity system. Then K has at most $(|E| - 2)/2$ maximal tangles.*

Proof. See [16, Corollary 9.4]. □

4.9 Algorithms and Tangles

In this section we discuss algorithmic questions related to tangles. Let $\mathcal{T}$ be a tangle in a graph G. The first observation is that $|\mathcal{T}|$ may be too large, thus rendering any complexity questions meaningless. If however, we only wish to decide if $(A, B) \in \mathcal{T}$, when $ord(A, B)$ is low (in comparison to $ord(\mathcal{T})$), then there is an elementary procedure we can follow.

Lemma 4.9.1. *Let G be a graph, $\mathcal{T}$ be a tangle in G, and Y be a $\mathcal{T}$-independent subset of $V(G)$. If (A, B) is a separation of G of order at most $|Y|/2$, then $(A, B) \in \mathcal{T}$ if and only if $|Y \cap A| < |Y|/2$.*

Proof. Let (A, B) be a separation of G of order at most $|Y|/2$. First suppose that $|Y \cap A| < |Y|/2$. If $(A, B) \notin \mathcal{T}$, then $(B, A) \in \mathcal{T}$. Let H be the subgraph of G with no edges and only those vertices in $Y \cap A$. Then $(B \cup H, A)$ is a separation of G of order $< |Y|$. Since Y is $\mathcal{T}$-independent, this implies that $(A, B \cup H) \in \mathcal{T}$. But then $G = A \cup B$, and both A and B are $\mathcal{T}$-small, a contradiction. The converse is similar. □

We will need to solve the following algorithmic problem.

Problem 4.9.2. Let $K = (E, \lambda)$ be a connectivity system and let s and t be distinct elements of E. The *minimum s-t cut problem* is to find a subset A of E such that $s \in A$, $t \notin A$, and $\lambda(A)$ is minimum.

We remark that the *global* minimum of λ is always assumed by $\emptyset$ since for any $X \subseteq E$ we have

$$2\lambda(X) = \lambda(X) + \lambda(E \backslash X) \geq \lambda(\emptyset) + \lambda(E) = 2\lambda(\emptyset).$$

Queyranne [31] shows that the minimum s-t cut problem for symmetric submodular functions is polynomially equivalent to the problem of minimizing a submodular (but not necessarily symmetric) function. Grötschel, Lovász, and Schrijver [19] exhibited the first polynomial-time algorithm for minimizing a submodular function, via the ellipsoid method. Later, combinatorial strongly polynomial algorithms were developed independently by Schrijver [46] and by Iwata, Fleischer, and Fujishige [21]. Therefore, there is a strongly polynomial algorithm that finds a minimum s-t cut problem for any symmetric submodular function.

Theorem 4.9.3. *Let λ be a symmetric submodular function on a finite set E, and let s and t be distinct members of E. There is a strongly polynomial algorithm that outputs a subset A of E such that $s \in A$, $t \notin A$, and $\lambda(A)$ is minimum.*

Using Theorem 4.9.3, we now show how to compute the rank of a set in the tangle matroid of a graph.

Theorem 4.9.4. *Let G be a graph, $K_G = (V(G) \cup E(G), \lambda_G)$ be its connectivity system, and $\mathcal{T}$ be a tangle in K_G. Let Y be an independent subset of vertices in the tangle matroid $M_{\mathcal{T}}$. Then for any $X \subseteq V(G)$ of rank at most $|Y|/2$ in $M_{\mathcal{T}}$, we can compute $r_{\mathcal{T}}(X)$ in polynomial-time.*

Proof. We will need the following operation on connectivity systems.

Definition 4.9.5. Let $K = (E, \lambda)$ be a connectivity system and let $X \subseteq$ E. Define $K \circ X := ((E \backslash X) \cup e_X, \lambda_{K \circ X})$, where

$$\lambda_{K\circ X}(A) := \begin{cases} \lambda(A), & \text{if } A \subseteq E\backslash X \\ \lambda((A\backslash\{e_X\}) \cup X), & \text{if } e_X \in A. \end{cases}$$

It is easy to verify that $K \circ X$ is a connectivity system. Now let X be a subset of $V(G)$ of rank at most $|Y|/2$ in $M_{\mathcal{T}}$. Let Y' be a subset of $Y\backslash X$ of size $\lceil |Y|/2 \rceil$. Consider the connectivity system $K' := (K_G \circ X) \circ Y'$. Let $s := e_X \in E(K')$ and $t := e_{Y'} \in E(K')$. By Theorem 4.9.3, we can find a minimum s-t cut A' of K' in polynomial-time. Letting $A := A'\backslash\{e_X\} \cup X$, we note that Lemma 4.9.1 implies that A is $\mathcal{T}$-small. By letting Y' range over all subsets of $Y\backslash X$ of size $\lceil |Y|/2 \rceil$, we will find the minimum order separation $(C, D) \in \mathcal{T}$, with $X \subseteq V(C)$, as required. $\square$

Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be distinct tangles of order n in a graph G. Using the same idea, we can compute a minimum order separation (A, B) distinguishing T_1 from T_2 provided that we know that $ord(A, B)$ is low in comparison to n.

Theorem 4.9.6. *Let G be a graph and let $\mathcal{T}_1, \ldots, \mathcal{T}_m$ be distinct tangles of order $\geq 2n + 1$ in K_G. Let $Y_1, \ldots, Y_m$ be independent subsets of vertices in $M_{\mathcal{T}_1}, \ldots, M_{\mathcal{T}_2}$ respectively, each of size $2n + 1$. Let K' be a tie-breaker for K_G. If for all distinct $i, j \in [m]$, $\mathcal{T}_i$ and $\mathcal{T}_j$ are distinguished by a separation of order at most n in K, then we can construct the tree of tangles for $\mathcal{T}_1, \ldots, \mathcal{T}_m$ (with respect to K') in polynomial-time.*

Proof. It suffices to show that we can compute minimum K'-order distinguishing separations between tangles in polynomial-time. Let (A, B) be the unique separation distinguishing $\mathcal{T}_i$ and $\mathcal{T}_j$ with minimum order in K'. By Lemma 4.9.1, we have that $|V(A) \cap Y_i| \leq n$ and $|V(B) \cap Y_j| \leq n$. This implies $|Y_i \cap Y_j| \leq n$, since $ord_K(A, B) \leq n$. Let $Z_i \subseteq (V(B) \cap Y_i) \backslash V(A)$ and $Z_j \subseteq (V(A) \cap Y_j) \backslash V(B)$ each be of size $n + 1$. Note that Z_i and Z_j are disjoint, and that (A, B) is simply the minimum order separation in K' with $Z_i \subseteq V(B)$ and $Z_j \subseteq V(A)$. Therefore, given Z_i and Z_j we can find (A, B) by considering the connectivity system $(K' \circ Z_i) \circ Z_j$, and applying Theorem 4.9.3. We can find Z_i and Z_j by enumerating over all pairs of subsets of size $n + 1$ of $Y_i \backslash Y_j$ and $Y_j \backslash Y_i$, respectively. $\square$

Remark 4.9.7. In our algorithm, it will not actually be necessary to *construct* the tree of tangles. Rather, it will suffice to find a peripheral tangle, which we can do much quicker, since we can avoid using submodular function minimization.

Chapter 5

Structure Theory

5.1 The Grid Theorem

Recall that the $n \times n$ *grid* is the graph G_n with vertex set

$$V(G_n) = \{(i, j) : i, j \in [n]\}$$

where (i, j) and (i', j') are adjacent if and only if

$$|i - i'| + |j - j'| = 1.$$

In Theorem 4.4.1, we exhibited a tangle $\mathcal{T}_n$ of order n in G_n. It thus follows by Theorem 4.6.1, that G_n has branch-width at least n. It is also easy to find a branch-decomposition of G_n of width n. Therefore,

Lemma 5.1.1. *The $n \times n$ grid has branch-width n.*

The grid theorem provides a partial converse to Lemma 5.1.1. It asserts that graphs with huge branch-width have large grid-minors. It was first proved by Robertson and Seymour in Graph Minors V [35, Theorem 1.5]. Diestel, Gorbunov, Jensen and Thomassen later found a shorter proof [7].

Theorem 5.1.2. *For all $n \in \mathbb{N}$, there exists $f(n) \in \mathbb{N}$ such that every graph with branch-width at least $f(n)$ has a minor isomorphic to the $n \times n$ grid.*

Remark 5.1.3. Let f be the function in Theorem 5.1.2. If we are given a graph G with $bw(G) \geq f(n)$, then it is quite easy to quickly *find* an

$n \times n$ grid-minor in G. To do this, we first use Theorem 3.1.4, to test if $bw(G) \leq f(n)$. If $bw(G) \leq f(n)$, then we can use dynamic programming or algorithms from monadic second-order logic to find an $n \times n$ grid-minor. Otherwise, we arbitrarily choose an edge e_1 of G and then test if $bw(G \backslash e_1) \leq f(n)$. If $bw(G \backslash e_1) \leq f(n)$, then note that $bw(G \backslash e_1) = f(n)$, since $bw(G \backslash e_1) \geq bw(G) - 1$. Thus, Theorem 5.1.2 guarantees that $G \backslash e_1$ still has an $n \times n$ grid-minor and as before we can find this grid-minor efficiently. If $bw(G \backslash e_1) > f(n)$, we choose an edge e_2 of $G \backslash e_1$ and recurse. We thus obtain a sequence of edges $e_1, \ldots, e_k$, such that $G \backslash \{e_1, \ldots, e_k\}$ has branch-width exactly $f(n)$. Therefore, we can efficiently find the required grid-minor in $G \backslash \{e_1, \ldots, e_k\}$.

5.2 Surfaces and Vortices

In this section we define surfaces and vortices, since they are required to state the Graph Minors Structure Theorem. Surfaces are treated in more detail in Chapter 7 and vortices are discussed further in Chapter 9. For more background information on surfaces, please refer to [29] or [6].

A *surface* Σ is a connected compact 2-manifold with (possibly empty) boundary. We let $bd(\Sigma)$ denote the boundary of Σ. The components in $bd(\Sigma)$ are the *holes* of Σ. The *genus* of Σ, denoted $\epsilon(\Sigma)$, is $2m + n$, where m and n are the number of handles and crosscaps of Σ, respectively. We let $h(\Sigma)$ denote the number of holes of Σ. If X is a subset of Σ, the (topological) closure, interior, and boundary of X will be denoted by $\overline{X}, int(X)$, and $bd(X)$ respectively. The surface obtained from Σ by capping each hole by a disk will be denoted $\widehat{\Sigma}$.

Let G be a graph embedded in a surface Σ. We will identify G with its embedding. Thus, $V(G) \subseteq \Sigma$ and each edge xy of G is an (open) arc in Σ connecting x and y. We will always assume that every edge of G is either contained in $bd(\Sigma)$ or disjoint from it. To stress the embedding we sometimes will write (G, Σ) whenever G is embedded in Σ. A *face* of G is a (topological) component of $\Sigma \backslash G$. Let f be a face of G. The *vertices* of f are the vertices of G contained on the boundary of f. The *edges* of f are defined similarly. We denote the vertices and edges of f as $V(f)$ and $E(f)$, respectively. If each face of G is an open disk, then we say that G is *2-cell embedded* in Σ. The *dual graph* G^* of G is the graph whose vertices are the faces of G, where two faces f_1 and f_2 are adjacent in G^* if and only if $E(f_1) \cap E(f_2) \neq \emptyset$.

A *society* is a finite set of points S that are cyclically ordered. If S is a finite subset of a circle, then evidently S can be regarded as a society. An *interval of* S is a proper subset of consecutive vertices of S. A *halving* of S is a partition of S into two intervals. For $u, v \in S$, we let $S(u,v)$ denote those vertices that occur after u but before v in S. We define $S[u,v] := S(u,v) \cup \{u,v\}$. So, if v is not the successor of u, then $\{S(u,v), S[v,u]\}$ is a halving of S. If G is a graph and $S \subseteq V(G)$ is a society, we call the pair (G,S) a *vortex*. A vortex (G,S) has *adhesion at most* n if for any halving of S, there do not exist n vertex disjoint paths in G between the two halves.

Example 5.2.1. Let G be the graph with vertex set $[k]$, and edge set

$$\{ij : |i-j| = 2 \text{ or } |i-j| = k-2\}.$$

Let S be the society in G with vertex set $[k]$ cyclically ordered as $1, 2, \ldots, k, 1$. It is easily seen that (G,S) has adhesion at most 5.

Let L be a linearly ordered set. We recycle our previous notion for cyclically ordered sets. For $u, v \in L$, we let $L(u,v)$ denote those vertices that occur after u but before v in L. We define $L[u,v] := L(u,v) \cup \{u,v\}$. If the ordering L under question is clear, we will occasionally write $[u,v]$ in place of $L[u,v]$.

Remark 5.2.2. Let δ be a hole in a surface Σ and let $x \in \delta$. Observe that we can regard $\delta \backslash \{x\}$ as a linearly ordered set since it is order-isomorphic to the open interval $(0,1)$. Thus, whenever we regard a hole as a linearly ordered set, this is the ordering we are referring to. We assume that x has been chosen a priori, and that if G is embedded in Σ, then G does not contain x.

Let G be a graph and let L be a linearly ordered subset of $V(G)$. A *vortex decomposition* of (G,L) is a collection $\{G_v : v \in L\}$ of subgraphs of G such that for all $x, y \in L$, with $x \leq y$:

(V1) $E(G_x \cap G_y) = \emptyset$, and $\bigcup_{v \in L} G_v = G$;

(V2) $G_x \cap G_y \subseteq \bigcap_{z \in L[x,y]} G_z$;

(V3) if $x \in V(G_y)$, then $y = x$ or u is the successor of x in L.

The *depth* of such a decomposition is $\max\{|V(G_x \cap G_y)| : x \neq y\}$, and its *width* is $\max\{|V(G_v)| - 1 : v \in S\}$. The *depth* of a vortex (G, L) is the minimum depth taken over all vortex decompositions of (G, L). The *width* of a vortex is defined similarly.

Remark 5.2.3. It is also possible to define vortex decompositions with respect to cyclically ordered sets. However, we prefer to work with linear vortex decompositions, since the notion coincides more closely with tree-decompositions (actually path-decompositions).

5.3 The Graph Minors Structure Theorem

The Graph Minors Structure Theorem [41, Theorem 1.3] is the workhorse of the entire Graph Minors Project. It gives a rough description of the class of graphs excluding a fixed minor. It has since been successfully applied to obtain a number of deep and interesting results. See [26] for an excellent survey.

For any graph K, we let ex(K) denote the class of graphs that do not contain a K-minor. We are principally interested in ex(K_n), and this is without loss of generality since every graph is a minor of some clique. Let us contemplate what a "rough description" of ex(K_n) might look like.

We start by considering a surface Σ that K_n does not embed in. We let $\mathcal{C}(\Sigma)$ denote the class of graphs that do embed in Σ. Since $\mathcal{C}(\Sigma)$ is minor-closed, it follows that $\mathcal{C}(\Sigma) \subseteq$ ex(K_n).

Suppose that K_{n-l} also does not embed in Σ for some $l > 0$. It follows that if G is a graph such that $G \backslash X \in \mathcal{C}(\Sigma)$, for some $X \subseteq V(G)$ with $|X| \leq l$, then G must also be ex(K_n). It is easy to check that the class of all such graphs G is minor-closed, called the *l-apex* of $\mathcal{C}(\Sigma)$.

Let G_1 and G_2 be graphs on disjoint vertex sets, and for $i \in [2]$ let $H_i \subseteq G_i$ be a clique of order k in G_i. Let G be a graph obtained from G_1 and G_2 by identifying H_2 with H_1 and then removing (a possibly empty) subset of edges of H_1. We say that G is a *clique-sum of order* k of G_1 and G_2. Note that if G_1 and G_2 are both in ex(H), then so is any clique-sum of G_1 and G_2. We remark that tree-width can be defined via clique-sums. Namely, a graph has tree-width at most w if and only if it can be obtained from graphs with at most $w+1$ vertices by (repeated) applications of clique-sums.

Let H be a planar graph, and let C be a cycle of H bounding a face with k vertices. Label $V(C) = \{1, \ldots, k\}$, and let $(G, V(C))$ be the vortex from Example 5.2.1. Glue H and G together along $V(C)$ and let the resulting graph be H^+. Evidently H has no K_5-minor as it is planar. Seese and Wessel [47] proved that H^+ can have a K_6-minor, but not a K_7-minor. In general, if G is embedded in a surface Σ, then attaching a vortex of bounded adhesion to a face of G will not produce arbitrarily large clique-minors. Furthermore, it is easy to show that such a graph cannot be produced via the previous three ingredients discussed. Thus, in any potential structure theorem, vortices will inevitably appear.

It turns out the four ingredients so far discussed, *clique-sums, surfaces, apex vertices, and vortices* are indeed sufficient to describe all graphs that do not have a K_n-minor.

We say that G can be *l-near embedded* in a surface Σ if there exists $A \subseteq V(G)$ of size at most l, and holes $\delta_1, \ldots, \delta_l$ of Σ such that $G \backslash A = \bigcup_{i=0}^{l} G_i$ satisfies

(N1) G_0 is embedded in Σ.

(N2) The graphs $G_1, \ldots, G_l$ are pairwise disjoint (and possibly empty), and $L_i := V(G_0) \cap V(G_i) = V(G_0) \cap \delta_i$, for each $i \in [l]$.

(N3) For each $i \in [l]$, if L_i is linearly ordered via δ_i, then (G_i, L_i) has a vortex decomposition of depth at most l.

The vertices in A are called the *apex vertices*, the graph G_0 is the *embedded subgraph* and the pairs (G_i, L_i) are the *vortices* of the l-near embedding. We say that each vortex (G_i, L_i) is *attached* to the hole δ_i.

We can now state the Graph Minors Structure Theorem.

Theorem 5.3.1. *For any $n \in \mathbb{N}$ there exists $l \in \mathbb{N}$ such that every graph that does not contain a K_n-minor can be obtained via clique sums of order at most l from graphs that can be l-near embedded in a surface in which K_n cannot be embedded.*

5.4 Structure Relative to a Tangle

In this section we present an alternative form [41, Theorem 3.1] of the Graph Minors Structure Theorem which is more suitable for our purposes. We wish to describe the structure of $G \in \text{ex}(K_n)$ relative to a high order tangle $\mathcal{T}$ of G. First we need some definitions.

Let G be a graph and let (A, B) be a separation of G of order $t \in [3]$, such that $V(B) \backslash V(A)$ is non-empty. We define G' to be the graph derived from A by placing a clique on the boundary of (A, B). We say that G' is obtained from G by an *elementary reduction*. If $t = 3$, we call the new set of edges the *reduction triangle*. A graph H is a *reduction of* G, if H can be obtained from G by any sequence of elementary reductions.

Let Σ be a surface. We say that G can be embedded in Σ *(up to 3-separations)*, if there exists a graph H that is a reduction of G such that H is embedded in Σ, and every reduction triangle bounds a face in Σ.

Let G be a graph with no K_n-minor and let $\mathcal{T}$ be a high order tangle that controls a large grid-minor J of G. The structure of G relative to $\mathcal{T}$ is as follows. The entire graph G embeds in a surface Σ (up to 3-separations) in which K_n does not embed. The grid-minor J is embedded in Σ, and a large portion of it is embedded in a disk of Σ. There is a bounded number of vortices of bounded adhesion attached to holes of Σ, and there is a bounded number of apex vertices that are arbitrarily connected to each other and the rest of G.

5.5 An Algorithmic Structure Theorem

In this section we present an algorithmic version of the Graph Minors Structure Theorem. The main idea here is due to Paul Seymour (communicated to us by Guoli Ding). Again, we present a slightly different version of the structure theorem here.

Let a, g, h, and d be non-negative integers. Let $\mathcal{H}_0(g, h)$ be the set of all pairs (G, Σ) where Σ is a surface of genus g and with h holes and G is a simple graph embedded in Σ.

Let $(G, \Sigma) \in \mathcal{H}_0(g, h)$, let δ be a hole of Σ, and let $v_1, \ldots, v_k$ be the vertices of G on δ (in the natural order). Now let $X_1, \ldots, X_k$ be disjoint sets of size d such that $X_i \cap V(G) = \{v_i\}$ for each $i \in \{1, \ldots, n\}$. Let G' be the simple graph obtained from G by adding the vertices $X_1 \cup \cdots \cup X_k$ and adding all edges internal to each of the sets $X_1 \cup X_2,\ X_2 \cup X_3, \ldots, X_{k-1} \cup X_k,\ X_k \cup X_1$. We say that G' is obtained from G by adding a *complete vortex of depth* d to the hole δ of (G, Σ). Let $\mathcal{H}_1(g, h, d)$ be the set of graphs obtained from the embedded graphs in $\mathcal{H}_0(g, h)$ by adding a complete vortex of depth d to each hole. Let $\mathcal{H}_2(g, h, d)$ be the class of graphs obtained by closing the class $\mathcal{H}_1(g, h, d)$ under minors.

Let $\mathcal{H}(g, h, d, a)$ be the class of all graphs G such that there is a set $X \subseteq V(G)$ with $|X| \leq a$ such that $G \backslash X \in \mathcal{H}_2(g, h, d)$. Finally, for a class $\mathcal{H}$ of graphs, we let $\mathcal{H}^{\oplus}$ denote the closure of the class $\mathcal{H}$ under clique-sums.

The following version of the Graph Minors Structure Theorem is equivalent to the main result of Graph Minors XVI.

Theorem 5.5.1. *There exist functions* $g, h, d, a : \mathbb{N} \to \mathbb{N}$ *such that, for each* $n \in \mathbb{N}$, *if G is a graph with no* K_n*-minor, then*

$$G \in \mathcal{H}(g(n), h(n), d(n), a(n))^{\oplus}.$$

Using Thereom 1.1.7, in $O(|V(G)|^3)$-time, we can determine whether a given graph G has an K_n-minor. Suppose that G does not contain a K_n-minor. Then, by Theorem 5.5.1,

$$G \in \mathcal{H}(g(n), h(n), d(n), a(n))^{\oplus}.$$

This theorem shows that G can be constructed from simple pieces via clique-sums, but it does not give an algorithm to find the construction. The proof of Theorem 5.5.1 can be adapted into a polynomial-time algorithm that either finds a K_n-minor or demonstrates that $G \in \mathcal{H}(g(n), h(n), d(n), a(n))^{\oplus}$. (Here "demonstrating that G is in $\mathcal{H}(g, h, d, n)^{\oplus}$" refers to a description of G as being obtained from embedded graphs in $\mathcal{H}_0(g, h)$ by adding vortices, taking minors, adding apex vertices, and then taking clique-sums.) Due to the length and difficulty of the proof of Theorem 5.5.1, this algorithm has neither been explicitly described nor analysed in the literature. In this section we sketch a proof of the following theorem which (partially) resolves this computational issue (only partially, since the algorithm is not explicit).

Theorem 5.5.2. *For any* $g, h, d, a \in \mathbb{N}$ *there exist* $d', a' \in \mathbb{N}$ *and an algorithm such that, given a graph G in* $\mathcal{H}(g, h, d, a)^{\oplus}$, *the algorithm will, in* $O(|V(G)|^5)$*-time, demonstrate that G is in* $\mathcal{H}(g, h, d', a')^{\oplus}$.

We will only sketch the proof.

For any class $\mathcal{H}$ of graphs and $G \in \mathcal{H}$, we call G *edge-maximal* if, for each pair (u, v) of non-adjacent vertices in G, the graph $G + uv$ is not in $\mathcal{H}$. Seymour observed that it is relatively easy to recover the structure of edge-maximal graphs in $\mathcal{H}(g, h, d, a)^{\oplus}$.

Recall that Theorem 1.1.1 and Theorem 1.1.7 together imply the following theorem.

Theorem 5.5.3. *For any minor-closed class $\mathcal{H}$ of graphs, there* exists *an algorithm that, given a graph G, will, in $O(|V(G)|^3)$-time, determine whether or not G is contained in $\mathcal{H}$.*

We have emphasized the word "exists" in the statement since we do not know the algorithm explicitly. We would need to be able to compute a bound on the size of the largest excluded minor of $\mathcal{H}$ in order to get an explicit algorithm.

The following result is an immediate consequence of Theorem 5.5.3.

Lemma 5.5.4. *For any minor-closed class $\mathcal{H}$ of graphs, there* exists *an algorithm that, given a simple graph $G \in \mathcal{H}$, constructs, in $O(|V(G)|^5)$-time, an edge-maximal graph $G' \in \mathcal{H}$ that contains G as a spanning subgraph.*

The following result is an easy consequence of a theorem of Tarjan [52] on clique cut-sets and of a theorem of Mader [28] on minor-closed classes of graphs.

Lemma 5.5.5. *For any proper minor-closed class of graphs, there is an algorithm that, given a simple edge-maximal graph $G \in \mathcal{H}^{\oplus}$, finds, in $O(|V(G)|^2)$-time, induced subgraphs $H_1, \ldots, H_k$ of G such that $H_1, \ldots, H_k$ are simple edge-maximal graphs in $\mathcal{H}$, none of $H_1, \ldots, H_k$ has a clique cut-set, and G is obtained from $H_1, \ldots, H_k$ by clique sums. Moreover, $|E(H_1)| + \ldots + |E(H_k)|$ is $O(|V(G)|)$.*

By Lemmas 5.5.4 and 5.5.5, to prove Theorem 5.5.2, it suffices to consider a simple edge-maximal graph G in $\mathcal{H}(g, h, d, a)$. We may assume that $d \geq 4$. We will define d' and a' implicitly, but we will take $a' \geq w$ where w is the size of the largest complete graph in $\mathcal{H}(g, h, d, a)$.

We say that a vertex is *universal* in a graph if it is adjacent to all other vertices in the graph. Let X denote the set of universal vertices in H; by computing vertex degrees we can find X in linear-time. Note that

$|X| \le w \le a'$ and that $G \backslash X$ is a graph in $\mathcal{H}_2(g,h,d)$. Thus we have reduced the problem to that of recovering the structure of graphs in $\mathcal{H}_2(g,h,d)$.

Let H be a graph in $\mathcal{H}_2(g,h,d)$. We assume that:

(A1) H is an edge-maximal graph in $\mathcal{H}_2(g,h,d)^{\oplus}$.

(A2) H has no clique cut-set.

(A3) H is not contained in $\mathcal{H}(g-1,h,d,10)$.

(A4) H is not contained in $\mathcal{H}(g,h-1,d,2d+10)$.

Note that we can test each of these assumptions in $O(|V(G)|^5)$-time and, if any is violated, we can inductively simplify the problem of recovering the strucure of H.

Since $H \in \mathcal{H}_2(g,h,d)$, H is a minor of a graph H_1 that is obtained from an embedded graph $(H_0,\Sigma) \in \mathcal{H}_0(g,h)$ by adding complete vortices of depth d to each hole. Consider such a construction of H such that $|V(H_0)|$ is minimum. By the edge-maximality of H, H is obtained from H_1 by contracting a set C of edges and then suppressing parallel pairs (we contract any loops created in the process).

By **(A1)** and **(A3)**, the embedding (H_0,Σ) is a triangulation with "representativity" at least 10. (That is, every non-contractible curve in Σ and every curve joining two distinct holes must intersect H_0 in at least 10 distinct points.) It also follows from **(A3)** that no edge of H_0 connects two non-consecutive vertices on a hole. By the minimality of $|V(H_0)|$, no edge of H_0 is contained in C. Note that, by **(A2)**, H_0 does not contain a subgraph isomorphic to K_4.

Let δ be a hole of Σ, let $v_1,\ldots,v_k$ be the vertices of H_0 on δ (in the natural order), and let $X_1,\ldots,X_k$ be vertex sets of size d that define the vortex on δ in H_1. We claim that there is no edge in C that has both of its ends in one of the sets $X_1,\ldots,X_k$. Suppose otherwise. Let $e \in C$ and suppose that e has both of its ends in X_i. Let v_i be the unique vertex in $X_i \cap V(H_0)$. Let H_0' be the proper subgraph of H_0, obtained by moving the neighbours of v_i in H_0 onto the hole δ. Clearly, $(H_0',\Sigma) \in \mathcal{H}_0(g,h)$. Now observe that H_1/e is a minor of a graph that can be obtained from H_0' by adding complete vortices of depth d to each hole. Since H is a minor of H_1/e, this contradicts the minimality of $|V(H_0)|$. This contradiction verifies our claim that there is no edge in C having both ends in one of the sets $X_1,\ldots,X_k$.

We have shown that each K_4-subgraph of H will contain an edge of one of the vortices. Conversely, since $d \geq 4$, each edge in a vortex is contained in a K_4-subgraph. Let Z denote the set of all vertices of G that are in a K_4-subgraph. One can find Z in $O(|V(H)|^2)$-time. Indeed, it suffices to consider all pairs of edges and the number of edges is $O(|V(H)|)$.

The graph $H\backslash Z$ is a subgraph of H_0 obtained by deleting all boundary vertices as well as some of their neighbours. Since H_0 triangulated Σ, it is straightforward to recover the embedding of $H\backslash Z$. Each edge in $H\backslash Z$ is in at most two triangles, each triangle in $(H\backslash Z, \Sigma)$ bounds a face, and the only faces of $(H\backslash Z, \Sigma)$ that are not triangles are the faces containing holes. Let W be the edges of $H\backslash Z$ that are not in two triangles. These are the edges that are in the boundaries of the faces containing the holes. Now that we know the boundaries of each of the faces in $(H\backslash Z, \Sigma)$ we have the embedding.

The rest of H attaches to $(H\backslash Z, \Sigma)$ as vortices of adhesion at most $2(d+1)$. Realizing the structure of these vortices is a routine matter of uncrossing separations and is omitted.

Using the faster minor-testing algorithm in [25] improves the complexity in Theorem 5.5.2 to $O(|V(G)|^3 \log(|V(G)|))$.

5.6 Excluding a Group-labelled Graph

In this section we discuss structure theorems for Γ-labelled graphs. These results were all proved by Geelen and Gerards in [14].

The first theorem asserts that sufficiently large clique-minors in $\widetilde{G}$ force big $\{0\}$-labelled clique minors in G.

Theorem 5.6.1. *For each finite abelian group Γ and each $n \in \mathbb{N}$ there exists $f := f(n, |\Gamma|) \in \mathbb{N}$ such that for all Γ-labelled graphs G, if $\widetilde{G}$ has a K_f-minor, then G has a $K(\{0\}, n)$-minor.*

Proof. This follows straightforwardly from Ramsey's theorem. See [14, Theorem 2.8]. □

A *block* of a Γ-labelled graph G is a maximal 2-connected subgraph of G. We regard a single edge (but not a single vertex) as a 2-connected graph. The next theorem is the main theorem of [14].

Theorem 5.6.2. *Let Γ be a finite abelian group, let Γ' be a subgroup of Γ, let $m \in \mathbb{N}$, and let $t = n|\Gamma|^2$. If G is a Γ-labelled graph and G has a minor H which is isomorphic to $K(\Gamma', 4t)$, then either*

- *there is a set X of at most t vertices of G such that the unique block of $G\backslash X$ that contains most of $E(H)$ is Γ'-balanced, or*
- *there is a subgroup Γ'' of Γ properly containing Γ' and a minor H' of G such that H' is isomorphic to $K(\Gamma'', m)$ and $E(H') \subseteq E(H)$.*

Remark 5.6.3. Using the algorithm of Chudnovsky, Cunningham, and Geelen [4], the proof of Theorem 5.6.2 is constructive. Regarding m as a constant, we can either find the set X or the minor H' in $O(|V(G)|^5)$-time.

Lemma 5.6.4. *Let G be a loopless Γ-labelled graph and $\mathcal{T}$ be a tangle of order n in G. For any $X \subseteq V(G)$ with $|X| \leq n-2$, there is a unique block H of $G\backslash X$ such that $G[V(H) \cup X]$ is not contained in any $\mathcal{T}$-small subgraph of G.*

Proof. For each block H of $G\backslash X$, define $\mathcal{T}_H$ to be the collection of separations (A, B) of $G\backslash X$, with $ord(A, B) < 2$ and $H \subseteq B$. It is readily checked that $\mathcal{T}_H$ is a tangle of order 2 in $G\backslash X$ and that every tangle in $G\backslash X$ of order 2 arises in this way.

If A is a subgraph of $G\backslash X$, we define A^+ to be $G[V(A) \cup X]$. Let $\mathcal{T}'$ be the collection of all separations (A, B) of $G\backslash X$ of order $< n - |X|$, such that A^+ is $\mathcal{T}$-small. Clearly, $\mathcal{T}'$ is a tangle in $G\backslash X$. Let $\mathcal{T}_2$ be the truncation of $\mathcal{T}'$ to order 2. Thus, there is a unique block H of $G\backslash X$, with $\mathcal{T}_2 = \mathcal{T}_H$. Towards a contradiction, suppose H^+ is contained in a $\mathcal{T}$-small subgraph of G. By definition, $H \subseteq A$, for some $(A, B) \in \mathcal{T}'$. If $ord(A, B) \leq 1$, then $(A, B) \in \mathcal{T}_2 = \mathcal{T}_H$, contradicting the definition of $\mathcal{T}_H$. Otherwise, we slide the separation (A, B) towards H to obtain a separation $(A', B') \in \mathcal{T}'$, with $ord(A', B') \leq 1$ and $H \subseteq A'$. $\square$

We call H the *$\mathcal{T}$-large block* of $G\backslash X$. The following is the main theorem of [14].

Theorem 5.6.5. *For all $n \in \mathbb{N}$ and all finite abelian groups Γ, there exists $l, t \in \mathbb{N}$ such that if G is a Γ-labelled graph and $\mathcal{T}$ is a tangle of order at least $t+2$ in G then either*

- *$\mathcal{T}$ controls a $K(\Gamma, n)$-minor in G,*
- *$\mathcal{T}$ does not control a K_l minor in $\widetilde{G}$, or*
- *There exists $X \subseteq V(G)$, with $|X| \leq t$, such that the $\mathcal{T}$-large block of $G \backslash X$ is Γ'-balanced for some proper subgroup Γ' of Γ.*

In the case that $\mathcal{T}$ does not control a K_l minor in $\widetilde{G}$, we can use the Graph Minors Structure Theorem to describe the structure of $\widetilde{G}$.

Chapter 6

Redundant Vertices in Clique-minors

Let G be a Γ-labelled graph and let Π be a pattern in G. A vertex $v \in V(G)$ is *essential for* Π, if Π is realizable and $v \in V(\mathcal{P})$ for any realization $\mathcal{P}$ of Π in G. In particular, if Π is realizable, then any vertex $v \in V(\Pi)$ is essential. A vertex is *redundant for* Π if it is not essential for Π. This chapter discusses redundant vertices in clique-minors.

6.1 Big Γ-labelled Cliques

Let G be a Γ-labelled graph and let Π be a pattern in G. We will prove that a big Γ-labelled clique-minor of G contains redundant vertices for Π. With respect to finding redundant vertices, big Γ-labelled clique-minors in Γ-labelled graphs are analogous to big clique minors in (unlabelled) graphs.

Here is our main result.

Theorem 6.1.1. *Let G be a Γ-labelled graph and let Π be a pattern in G. Let $X := V(\Pi)$ and let H be a minor of G which is isomorphic to $K(\Gamma, n)$, where $n > 3|X| + 1$. Let $\mathcal{T}$ be the tangle of order $\theta := \lceil 2n/3 \rceil$ in G induced by H and let $Y := cl_{\mathcal{T}}(X)$. Lastly, let (A, B) be the separation $(G[Y], G \ominus G[Y])$ in G. Under these hypotheses, any vertex $v \in V(B) \backslash V(A)$ is redundant for Π.*

Proof. Note that such a vertex v exists from the definition of $\mathcal{T}$. Let $r := r_{\mathcal{T}}(X)$ and let Z be the boundary of (A, B). By Lemma 4.3.1, $\mathcal{T}$

controls a tangle $\mathcal{T}_1$ of order $\theta - 1$ in $G\backslash v$. By Lemma 4.3.2, $\mathcal{T}_1$ controls a tangle $\mathcal{T}_2$ of order $(\theta - 1) - r$ in $B\backslash v$. Since $n > 3|X| + 1$, we have that $\theta - 1 > 2r$. Thus, by Lemma 4.3.3, it follows that Z is $\mathcal{T}_2$-independent. Finally, Lemma 4.5.3 implies that in $B\backslash v$ we can contract a $K(\Gamma, r)$-minor onto Z. It immediately follows that v is indeed redundant for Π. $\square$

In the special case that the vertices of Π are $\mathcal{T}$-independent, the previous proof shows that G actually does have a Π-linkage. Thus, we have the following sufficient conditions.

Theorem 6.1.2. *Let G be a Γ-labelled graph and let Π be a pattern of size k in G. Let H be a $K(\Gamma, n)$-minor in G, where $n > 6k + 1$. Let $\mathcal{T}$ be the tangle in G induced by H. If $V(\Pi)$ is $\mathcal{T}$-independent, then G has a Π-linkage.*

6.2 Big Γ'-labelled Cliques

Let G be a Γ-labelled graph and let Γ' be a subgroup of Γ. In this section we consider instances of Π-linkage problems where G has a large Γ'-labelled clique-minor. To be precise, we aim to prove the following theorem.

Theorem 6.2.1. *Let $f(k, |\Gamma|) = 12k^2|\Gamma|2^{2k|\Gamma|}$. Let G be a Γ-labelled graph and let Π be a pattern in G with $|\Pi| = k$. If $G\backslash V(\Pi)$ is Γ'-balanced for some subgroup Γ' of Γ and $G\backslash V(\Pi)$ has a $K(\Gamma', f(k, |\Gamma|))$-minor, H, then there exists a vertex v of H that is redundant for (G, Π).*

Proof. Let $X = V(\Pi)$. By hypothesis, $G\backslash X$ is Γ'-balanced, and $G\backslash X$ has a minor H which is isomorphic to $K(\Gamma', f(k, |\Gamma|))$. It is easy to see that we may perform shifts so that $G\backslash X$ is Γ'-labelled and such that we can obtain H from $G\backslash X$ without the need to perform further shifts.

By flipping edges, we may assume that all the edges between X and $G\backslash X$ are directed away from X. Now, the $K(\Gamma', f(k, |\Gamma|))$-minor, H, induces a tangle $\mathcal{T}$ in $G\backslash X$ of order

$$\theta := \lceil 2f(k, |\Gamma|)/3 \rceil \geq 2|X|^2|\Gamma|2^{|X||\Gamma|}.$$

For each $x \in X$ and $\gamma \in \Gamma$, we let

$$N_{x,\gamma} := \{u \in V(G\backslash X) : e = xu \in E(G), \gamma_e = \gamma\}.$$

There are $|X||\Gamma|$ such sets. We re-index them $N_1, \dots, N_{|X||\Gamma|}$ where

$$r_{\mathcal{T}}(N_1) \leq \cdots \leq r_{\mathcal{T}}(N_{|X||\Gamma|}).$$

Let j be the minimum index such that $r_{\mathcal{T}}(N_j) > |X|^2|\Gamma|2^{j-1}$. If no such j exists, we set $j := |X||\Gamma| + 1$.

We choose a separation (A, B) of $G \backslash X$, such that

- $\bigcup_{i=1}^{j-1} N_i \subseteq A$,
- $ord(A, B) = r_{\mathcal{T}}(\bigcup_{i=1}^{j-1} N_i)$, and
- (A, B) is $\mathcal{T}$-independent.

Since (A, B) is $\mathcal{T}$-independent, we have that the boundary Y of (A, B) is $\mathcal{T}$-independent. Note that $\mathcal{T}$ controls a tangle $\mathcal{T}'$ in B of order at least $\theta - |Y|$. Since $\theta \geq 2|Y|$, we conclude that Y is in fact $\mathcal{T}'$-independent by Lemma 4.3.3.

Observe that if $j \neq |X||\Gamma| + 1$, then

$$\begin{aligned}
r_{\mathcal{T}}(\bigcup_{i=1}^{j-1} N_i) &\leq \sum_{i=1}^{j-1} r_{\mathcal{T}}(N_i) \\
&\leq \sum_{i=1}^{j-1} |X|^2|\Gamma|2^{i-1} \\
&= |X|^2|\Gamma|(2^{j-1} - 1) \\
&< r_{\mathcal{T}}(N_j) - |X|^2|\Gamma|
\end{aligned}$$

Let $\mathcal{B}$ be the set of indices i such that $j \leq i \leq |X||\Gamma|$. Note that $\mathcal{B} = \emptyset$ if $j = |X||\Gamma| + 1$. By the above inequality, for each $i \in \mathcal{B}$, we can choose a subset V_i of vertices of G such that

- $\mathcal{V} := \{V_i : i \in \mathcal{B}\}$ is a disjoint family,
- $V_i \subseteq N_i \cap B$,
- $|V_i| = |X|$, and
- $Z := Y \cup \bigcup_{i \in \mathcal{B}} V_i$ is $\mathcal{T}'$-independent.

Let v be a vertex of H so that $Z \cup \{v\}$ is $\mathcal{T}'$-independent. By Lemma 4.5.3, we can contract a Γ'-labelled clique onto $Z \cup \{v\}$ in B. It is hence clear that v is redundant for Π since every edge in B has group-value in Γ'. $\square$

Chapter 7

Surfaces

7.1 Curves in Surfaces

Recall that a surface is a connected compact 2-manifold with (possibly empty) boundary. Let Σ be a surface. A *curve* S in Σ is a continuous function $S : [0, 1] \to \Sigma$. It is *simple* if S is injective, and it is *closed* if $S(0) = S(1)$. Abusing terminology, we call $S(0)$ the *tail* of S and $S(1)$ the *head* of S, respectively denoted $tail(S)$ and $head(S)$. The head and tail of S are its *ends*. An *arc* is a simple curve. Two curves in Σ are *internally disjoint* if they are disjoint except possibly at their ends. Let X and Y be subsets of Σ. A curve S is an X-curve if $ends(S) \subseteq X$, and S is otherwise disjoint from X. It is an X-Y curve if one end of S in on X, the other is on Y, and S is otherwise disjoint from $X \cup Y$. A curve in Σ is *normal* if it is a $bd(\Sigma)$-curve. Let S^1 denote the unit circle in $\mathbb{C}$. A *circle in* Σ is a subset of Σ homeomorphic to S^1. A circle C is *separating* if $\Sigma \backslash C$ is disconnected, and is *non-separating* otherwise. A circle is *contractible* if it bounds a closed disk in Σ, otherwise it is *non-contractible*. Certain arcs may also be contractible. We say that an arc A is *contractible* if A is a δ-arc for some hole δ of Σ and A together with some subset of δ bounds a disk in Σ. Otherwise, A is *non-contractible*.

We next define various notions of homotopies between curves in Σ. Two circles C_1 and C_2 in Σ are *freely homotopic* if there is a continuous function $H : S^1 \times [0, 1] \to \Sigma$ with $H(x, 0) = C_1(x)$ and $H(x, 1) = C_2(x)$ for all $x \in S^1$. We call H a *homotopy* which brings C_1 to C_2.

Similarly, we can define homotopy of curves with fixed base points.

Let A_1 and A_2 both be arcs from p to q in Σ (we allow $p = q$). We say that A_1 is *homotopic* to A_2 (relative to p and q) if there is a continuous function $H : [0,1] \times [0,1] \to \Sigma$ with $H(x,0) = A_1(x)$, $H(x,1) = A_2(x)$, $H(0,y) = p$, and $H(1,y) = q$ for all $x, y \in [0,1]$.

Finally, we now define homotopy for normal arcs of Σ. We first consider arcs with both their ends on the same hole. Let δ be a hole of Σ and let A_1 and A_2 be δ-arcs in Σ. We say that A_1 and A_2 are *homotopic* (relative to δ) if $A_1 \cup A_2$ and some subset of δ bounds a disk in Σ. Homotopy for arcs which connect up two holes is defined similarly. Let δ_1 and δ_2 be distinct holes of Σ and let A_1 and A_2 be δ_1-δ_2 arcs. We say that A_1 and A_2 are *homotopic* (relative to δ_1 and δ_2) if $A_1 \cup A_2$ and some subset of $\delta_1 \cup \delta_2$ bounds a disk in Σ.

A family $\mathcal{C}$ of curves in Σ is called a t-family, if any two distinct curves in $\mathcal{C}$ intersect in at most t points. We require the following two theorems of Juvan, Malnič, and Mohar [23].

Theorem 7.1.1. *For any surface Σ and $t \in \mathbb{N}$, there is a constant $N(\Sigma, t)$ so that any t-family of pairwise non-homotopic circles in Σ has at most $N(\Sigma, t)$ members.*

Theorem 7.1.2. *For any surface Σ and $t \in \mathbb{N}$, there is a constant $N^*(\Sigma, t)$ with the following property. If $\mathcal{A}$ is a t-family of arcs connecting p to q in Σ which are pairwise non-homotopic (with respect to p and q), then $|\mathcal{A}| \leq N^*(\Sigma, t)$.*

We actually only require the cases $t \in [2]$.

Lemma 7.1.3. *For any surface Σ, there is a constant $\rho_1(\Sigma)$ such that if δ is a hole of Σ and $\mathcal{A}$ is a family of non-contractible δ-curves in Σ which are pairwise non-homotopic (with respect to δ) and pairwise internally disjoint, then $|\mathcal{A}| \leq \rho_1(\Sigma)$.*

Proof. Let $\mathcal{A}$ be such a family. We begin by performing the following operation on Σ. Cap the hole δ with a disk Δ, and then shrink Δ to a point x. Let Σ' be the resulting surface and $\mathcal{C}'$ be the resulting family of curves in Σ'. Observe that $\mathcal{C}'$ is a 1-family of pairwise non-homotopic circles in Σ'. Thus, letting $N(\Sigma, t)$ be the function from Theorem 7.1.1, we have that

$$|\mathcal{A}| = |\mathcal{C}'| \leq N(\Sigma', 1).$$

□

Lemma 7.1.4. *For any surface Σ, there is a constant $\rho_2(\Sigma)$ with the following property. If δ_1 and δ_2 are distinct holes of Σ and $\mathcal{A}$ is a family of normal δ_1-δ_2 arcs in Σ, which are pairwise non-homotopic (with respect to δ_1 and δ_2) and pairwise internally disjoint, then $|\mathcal{A}| \leq \rho_2(\Sigma)$.*

Proof. Let $\mathcal{A}$ be such a family. Begin by capping δ_1 with a disk Δ_1 and δ_2 with a disk Δ_2, and then shrinking δ_1 to a point x_1 and δ_2 to a point x_2. Let Σ' be the resulting surface, and $\mathcal{A}'$ be the resulting family of curves in Σ'. Note that $\mathcal{A}'$ is a family of internally disjoint arcs from x_1 to x_2 in Σ' that are pairwise non-homotopic with respect to x_1 and x_2. Thus, letting $N^*(\Sigma, t)$ be the function from Theorem 7.1.2, we have

$$|\mathcal{A}| = |\mathcal{A}'| \leq N^*(\Sigma', 2).$$

□

7.2 Linkages in Surfaces

Let Σ be a surface. A *pattern in* Σ is a finite collection Π of disjoint 2-subsets of Σ. The *size* of Π is $|\Pi|$, and we call Π a *k-pattern (in Σ)*, if it has size at most k. Let $\Pi = \{\{s_i, t_i\} : i \in [k]\}$ be a k-pattern in Σ. A *topological linkage (in Σ)* is a collection of pairwise disjoint arcs in Σ. A *topological realization* of Π is a topological linkage $\mathcal{L} := \{L_i : i \in [k]\}$ such that $ends(L_i) = \{s_i, t_i\}$ for each $i \in [k]$. We call such an $\mathcal{L}$ a *topological Π-linkage*. If Π has a topological realization we say that Π is *realizable in* Σ.

Robertson and Seymour made the following observation.

Lemma 7.2.1. *For any surface Σ and any $k \in \mathbb{N}$, there exists $t := t(k, \Sigma) \in \mathbb{N}$ such that if C is any circle in Σ, and Π is any realizable k-pattern in Σ, then there exists a Π-linkage in Σ that intersects C in at most t points.*

Proof. This follows since up to homeomorphism, there are only a finite number of k-patterns in Σ. □

It turns out that the *existence* of the function $t(k, \Sigma)$ is adequate for the proof of Theorem 1.1.7. However, since the proof of Lemma 7.2.1 is highly non-constructive, it is unclear whether $t(k, \Sigma)$ is in fact a computable function of k and Σ. Geelen and Richter [18] remedied this problem by showing:

Theorem 7.2.2. *For any $k \in \mathbb{N}$, if Σ is a surface, C is a circle in Σ and Π is a realizable k-pattern in Σ, then there exists a topological Π-linkage in Σ that intersects C in at most $2k$ points.*

Thus, we obtain explicit algorithms for both the H-minor testing problem and k-linkage problem for graphs. Interestingly, note that $t(k, \Sigma)$ does not actually depend on the surface Σ.

7.3 Representativity

Let G be a graph embedded in a surface Σ, with $\epsilon(\Sigma) \geq 1$ and with no holes. A curve C in Σ is *dual for G* if it intersects G only at vertices. The *vertices* of C are the vertices of G it contains, and the *length* of C is its number of vertices. The *representativity of G in Σ* is the minimum length of a non-contractible dual circle in Σ. We remark that if G is embedded in a surface Σ (possibly with holes), then the *representativity of G in Σ* is simply the representativity of G in $\widehat{\Sigma}$.

7.4 Distance on a Surface

In this section we present two metrics for graphs embedded in a surface.

Definition 7.4.1. Let G be a graph embedded in a surface Σ. For vertices $u, v \in V(G)$ we let $d_\Sigma(u, v)$ be the minimum of $|V(P)| - 1$ where P ranges over all dual curves from u to v. It is clear that d_Σ is a metric on $V(G)$, which we call the *surface metric*.

Example 7.4.2. Let G_{4n} be the $4n \times 4n$ grid embedded in a surface Σ so that G_{4n} is contained in a disk. Consider the vertices $u := (n, n)$ and $v := (3n, 3n)$ of G_{4n}. It is easy to see that $d_\Sigma(u, v) = 2n$.

For many applications, the surface metric is a suitable metric to work with. On the other hand, it is ultimately inadequate due to the following shortcoming. Let G be any graph embedded in a surface Σ and let F be a face of G which bounds a disk Δ. Let G_{4n} be the $4n \times 4n$ grid, when n is large. Let e be any edge of F and let f be any edge on the outerface of G_{4n}. Finally, let G' be the graph obtained from G_{4n} and G by identifying e and f.

We may regard G' as embedded in Σ since we can place G_{4n} inside Δ. By Example 7.4.2, there are vertices u and v of G_{4n} (regarded as a subgraph of G') such that $d_\Sigma(u,v) = 2n$. However, in some sense, u and v are not far apart in Σ since there is a dual circle of length 2 which bounds a disk containing u and v. To overcome this shortcoming, we introduce a second metric using tangles.

Definition 7.4.3. Let G be a 2-connected graph embedded in a surface Σ and let $\mathcal{T}$ be a tangle of order θ in G. For $u, v \in V(G)$, we define $d_\mathcal{T}(u,v)$ to be the minimum of $r_\mathcal{T}(V(P)) - 1$, where P ranges over all dual curves in Σ from u to v. We call the $d_\mathcal{T}$ the *tangle metric* on $V(G)$.

It was first shown in Graph Minors XI [38, Theorem 9.1] that $d_\mathcal{T}$ is indeed a metric on $V(G)$.

Theorem 7.4.4. *Let G be a 2-connected graph embedded in a surface Σ, and let $\mathcal{T}$ be a tangle in G. Then $d_\mathcal{T}$ is a metric on $V(G)$.*

Proof. Evidently, $d_\mathcal{T}$ is symmetric. Since G is 2-connected, it is also clearly non-negative, and $d_\mathcal{T}(u,v) = 0$ if and only if $u = v$. Let $u, v, w \in V(G)$. Let P be a dual curve from u to v with $r_\mathcal{T}(V(P)) - 1 = d_\mathcal{T}(u,v)$, and Q be a dual curve from v to w with $r_\mathcal{T}(V(Q)) - 1 = d_\mathcal{T}(u,v)$. Note that $R := P \cup Q$ is a dual curve from u to w. Note that since G is connected we have $r_\mathcal{T}(\{v\}) = 1$. Thus,

$$
\begin{aligned}
d_\mathcal{T}(u,w) &\leq r_\mathcal{T}(V(R)) - 1 \\
&= r_\mathcal{T}(V(P) \cup V(Q)) - 1 \\
&\leq r_\mathcal{T}(V(P)) + r_\mathcal{T}(V(Q)) - r_\mathcal{T}(V(P) \cap V(Q)) - 1 \\
&\leq r_\mathcal{T}(V(P)) - 1 + r_\mathcal{T}(V(Q)) - 1 \\
&= d_\mathcal{T}(u,v) + d_\mathcal{T}(v,w)
\end{aligned}
$$

□

For the remainder of this section $d \in \{d_\Sigma, d_\mathcal{T}\}$. Let X and Y be subsets of $V(G)$. We define the *distance between X and Y* (with respect to d) to be

$$d(X,Y) := \min\{d(x,y) : x \in X, y \in Y\}.$$

As special cases, the distance between two faces F_1 and F_2 of G is $d(F_1,F_2) := d(V(F_1), V(F_2))$ and the distance between two holes δ_1 and δ_2

is $d(\delta_1, \delta_2) := d(V(\delta_1), V(\delta_2))$. The *ball of radius r, centred at x* (with respect to d) is the set

$$B_d[x, r] := \{y \in V(G) : d(x, y) \leq r\},$$

and the *sphere of radius r, centred at x* is the set

$$S_d[x, r] := \{y \in V(G) : d(x, y) = r\}.$$

Similarly, we can define balls and spheres which are centred at faces of G, or holes of Σ.

7.5 Respectful Tangles

Observe that in the sphere, any circle is the boundary of *two* disks. Thus, a circle on the sphere does not have a well-defined "inside." In this sense, the sphere is peculiar among all surfaces. We overcome this difficulty by using *respectful tangles*, which allow us to completely unify graphs embedded on the sphere with those embedded in other surfaces. Respectful tangles were first introduced in Graph Minors XI [38]. Now to the definition.

Let G be a graph embedded in a surface Σ without holes, and let $\mathcal{T}$ be a tangle of order n in G. We say that $\mathcal{T}$ is *respectful* if for every dual circle C of length less than n, there is a disk Δ in Σ such that $bd(\Delta) = C$ and

$$(G \cap \Delta, G \cap \overline{\Sigma \backslash \Delta}) \in \mathcal{T}.$$

Observe that if G is embedded on the sphere, then every tangle of G is respectful. If Σ is not the sphere, and $\mathcal{T}$ is a respectful tangle of order n, then clearly the representativity of G is at least n. Conversely, Robertson and Seymour [38, Theorem 4.1] proved that if G is 2-cell embedded in a surface Σ and G has representativity at least n, then G has a respectful tangle of order n.

Theorem 7.5.1. *Let G be a graph 2-cell embedded in a surface Σ with no holes which is not the sphere. If the representativity of G is at least $n \geq 1$, then G has a unique respectful tangle $\mathcal{T}$ of order n.*

Remark 7.5.2. Let G be a graph embedded in a surface Σ with holes. By regarding G as embedded in the capped surface $\widehat{\Sigma}$, we can extend the definition of respectful tangles to include surfaces with boundary.

Let $\mathcal{T}$ be a respectful tangle of G. Recall that $\mathcal{T}$ induces a metric $d_{\mathcal{T}}$ on $V(G)$. We will require [39, Theorem 7.5], which describes the effect on $\mathcal{T}$ and $d_{\mathcal{T}}$ if we delete all the vertices of a face of G.

Theorem 7.5.3. *Let G be a graph 2-cell embedded in a surface Σ without holes, and let $\mathcal{T}$ be a respectful tangle of G of order $n \geq 3$. Let F be a face of G in Σ and let $X := V(F)$. Then there is a 2-cell subdrawing G' of $G \backslash X$ and a face F' of G' containing F such that*

- *G' has a respectful tangle $\mathcal{T}'$ of order $n-2$.*
- *G is a subset of $G' \cup F'$ (in Σ).*
- *$d_{\mathcal{T}'}(u,v) \geq d_{\mathcal{T}}(u,v) - 4$, for all $u, v \in V(G')$.*
- *$d_{\mathcal{T}'}(u,F') \geq d_{\mathcal{T}}(u,F) - 2$, for all $u \in V(G')$.*
- *$d_{\mathcal{T}}(u,v) \leq 3$, for all $u, v \in V(G) \cap F'$.*

By repeatedly applying Theorem 7.5.3, we get [39, Theorem 7.6].

Theorem 7.5.4. *Let $n \in \mathbb{N}$, let G be a graph 2-cell embedded in a surface Σ without holes, and let $\mathcal{T}$ be a respectful tangle of G of order $m \geq 2n+1$. Let F be a face of G and let $X \subseteq V(G)$ be the ball of radius $n-1$ (with respect to d_{Σ}) centred at F. Then there is a 2-cell subdrawing G' of $G \backslash X$, and a face F' of G' containing F such that*

- *G' has a respectful tangle $\mathcal{T}'$ of order at least $m-2n$.*
- *G is a subset of $G' \cup F'$.*
- *$d_{\mathcal{T}'}(u,v) \geq d_{\mathcal{T}}(u,v) - 4n$, for all $u, v \in V(G')$.*
- *$d_{\mathcal{T}'}(u,F') \geq d_{\mathcal{T}}(u,F) - 2n$, for all $u \in V(G')$.*
- *$d_{\mathcal{T}}(u,v), \leq 2n+1$, for all $u, v \in V(G) \cap F'$.*

Theorem 7.5.4 has the following simple, but useful corollary.

Corollary 7.5.5. *Let $n \in \mathbb{N}$, let G be a graph 2-cell embedded in a surface Σ without holes, let $\mathcal{T}$ be a respectful tangle of G of order $m \geq 2n+1$, and let F be a face of G. Then for each $i \in [n]$, there is a cycle C_i of G, such that*

- *C_i bounds a disk in Σ containing F,*
- *C_i only passes through vertices at distance exactly $i-1$ from F (with respect to d_Σ).*

The last result of the section asserts that if $\mathcal{T}$ is a respectful tangle in G of high order, then any $\mathcal{T}$-independent subset of vertices of a face F cannot be separated from G by a low order separation "close to F."

Lemma 7.5.6. *Let $n, G, \Sigma, \mathcal{T}, F$, and $\{C_i : i \in [n]\}$ be as in Corollary 7.5.5. If Z is a $\mathcal{T}$-independent subset of $V(F)$, then there are $|Z|$ vertex-disjoint paths in G from Z to $V(C_n)$.*

Proof. We assume for the moment that Σ is not the sphere. Recall that C_n bounds a disk Δ_n in Σ containing F. If the desired paths do not exist, then by Menger's Theorem there is a dual circle D of length less than $|Z|$, which separates Z from C_n. That is D bounds a disk $\Delta \subseteq \Delta_n$ such that $Z \subseteq \Delta$. Consider the separation $(G \cap \Delta, G \cap \overline{\Sigma \backslash \Delta})$. Since Σ is not the sphere and $\mathcal{T}$ is respectful, we must have that $(G \cap \Delta, G \cap \overline{\Sigma \backslash \Delta}) \in \mathcal{T}$. However, this contradicts the fact that Z is $\mathcal{T}$-independent. If Σ is the sphere, then $\overline{\Sigma \backslash \Delta}$ is also a disk, so it possible that

$$(G \cap \overline{\Sigma \backslash \Delta}, G \cap \Delta) \in \mathcal{T},$$

thus avoiding a contradiction.

We now give a proof for when Σ is the sphere. Incidentally, this proof actually works for all surfaces, rendering the previous paragraph obsolete. We proceed by induction on n. For each $i \in [n]$, recall that each vertex in C_i is at distance exactly $i-1$ from F. The main part of the proof is the following claim

Claim. *There is a collection of $|Z|$ disjoint paths in G from Z to C_2.*

SUBPROOF. Suppose not. Let Δ_2 be the disk bounded by C_2 that contains F. As before, there must be a dual circle D of length less than $|Z|$ which separates Z from C_2. So D bounds a disk Δ such that $\Delta \subseteq \Delta_2$ and $Z \subseteq \Delta$. We are done unless Σ is the sphere and

$$(G \cap \overline{\Sigma \backslash \Delta}, G \cap \Delta) \in \mathcal{T}.$$

In this case we choose a separation (A, B) of G such that

(1) $(A, B) \in \mathcal{T}$, and

(2) subject to (1), B is minimal.

Note that $B \subseteq G \cap \Delta$. Furthermore, there exist disks Δ'_1 and Δ'_2 in Σ such that

- $B \subseteq \Delta'_1 \cup \Delta'_2$,
- $bd(\Delta_i)$ is a dual curve of length at most $|Z|$ for each $i \in [2]$, and
- $G \cap \Delta'_i$ is a proper subgraph of B, for each $i \in [2]$.

By the minimality of B, for each $i \in [2]$ we must have

$$(G \cap \Delta'_i, (G \cap \overline{\Sigma \backslash \Delta'_i})) \in \mathcal{T}.$$

But now G is the union of A, $G \cap \Delta'_1$, and $G \cap \Delta'_2$, each of which is $\mathcal{T}$-small. This contradiction proves the claim. $\square$

Let $\mathcal{P}$ be such a collection of paths with $|E(\mathcal{P})|$ minimal and let Z_2 be the set of ends of $\mathcal{P}$ in C_2. Let $X := V(F)$. By Theorem 7.5.3, there is a 2-cell subdrawing G' of $G \backslash X$, such that C_2 is the boundary of a face F' of G' containing F. Moreover, G' has a respectful tangle $\mathcal{T}'$ of order at least $2(n-1)+1$. It is easy to see that Z_2 is $\mathcal{T}'$-independent, else Z would not be $\mathcal{T}$-independent. By induction, there is a collection $\mathcal{Q}$ of $|Z_2|$ vertex disjoint paths in G' from Z_2 to C_n. Combining $\mathcal{Q}$ and $\mathcal{P}$ appropriately gives the desired paths. $\square$

7.6 A Disk with Strips

This section provides a different way to view surfaces. A *strip* S is a surface homeomorphic to $[0, 1] \times [0, 10]$. The *tail* of S is $tail(S) := [0, 1] \times \{0\}$ and the *head* of S is $head(S) := [0, 1] \times \{10\}$. The *ends* of S are $ends(S) := \{tail(S), head(S)\}$. The *corners* of S are $cor(S) := \{(0, 0), (0, 10), (1, 0), (1, 10)\}$. A *disk with n strips* is a surface $\Sigma := \Delta \cup S_1 \cup \cdots \cup S_n$, where Δ is a disk and for each $i, j \in [n]$,

- S_i is a strip.

- $S_i \cap \Delta$ is precisely the union of the ends of S_i.

- S_i and S_j are disjoint, except possibly at corners.

For example, both the cylinder and the Möbius band are disks with 1 strip. If $\Sigma = \Delta \cup S_1 \cup \cdots \cup S_n$ is a disk with n strips, then we say $S_1, \ldots, S_n$ are the *strips* of Σ and that $\Delta(\Sigma) := \Delta$ is the *disk* of Σ.

7.7 A Disk with Γ-Strips

Let Γ be a finite abelian group. A Γ*-strip* is a strip S endowed with an element $\gamma(S) \in \Gamma$. We call $\gamma(S)$ the *group-value* of S. A *disk with* n Γ*-strips* is a disk with n strips $\Sigma := \Delta \cup S_1 \cup \cdots \cup S_n$, so that each S_i is a Γ-strip.

For the remainder of this section $\Sigma := \Delta \cup S_1 \cup \cdots \cup S_n$ is a disk with n Γ-strips. Recall that a Δ-arc in Σ is an arc A with both its ends on Δ that is otherwise disjoint from Δ. Evidently, this implies that the ends of A are on *bd*(Δ), and that the rest of A is contained in a Γ-strip of Σ.

Let $S_j := [0, 1] \times [0, 10] \to \Sigma$ be a Γ-strip of Σ and let A be a Δ-arc of Σ contained in S_j. We define the *orientation* of S_j to be the orientation it inherits from $[0, 10]$. We say that A *passes through* S_j if the ends of A are on different ends of S_j. It passes through S_j in the *positive direction* if $head(A) \in head(S_j)$ and $tail(A) \in tail(S_j)$. Conversely, A passes through S_j in the *negative direction* if $head(A) \in tail(S_j)$ and $tail(A) \in head(S_j)$. The *group-value*, or just *value*, of A is defined to be $\gamma(S_j)$ if A passes through S_j in the positive direction, -$\gamma(S_j)$ if A passes through S_j in the negative direction, and zero otherwise. We let $\gamma_\Sigma(A)$ denote the value of A.

We can extend this notion to arbitrary arcs as follows. Let L be an arc in Σ. The *value* of L is

$$\gamma_\Sigma(L) := \sum \gamma_\Sigma(A),$$

where the sum runs over all $A \subseteq L$ such that A is a Δ-arc.

It now makes sense to introduce group-valued linkage problems in Σ. A *pattern* Π in Σ is any set of triples of the form (x, y, γ), where x and y are distinct points of Σ, no point in Σ occurs in more than one triple of Π, and $\gamma \in \Gamma$. As before, Π is a k*-pattern* if $|\Pi| \leq k$.

Definition 7.7.1. Let $\Pi := \{(s_i, t_i, \gamma_i) : i \in [k]\}$ be a pattern in a disk with Γ-strips Σ. A *topological realization of* Π *in* Σ is a family $\{L_i : i \in [k]\}$ of disjoint arcs in Σ such that for each $i \in [k]$, the tail of L_i is s_i, the head of L_i is t_i, and the value of L_i is γ_i.

A pattern Π in Σ is *topologically realizable* if it has a topological realization.

Remark 7.7.2. We mention that a pattern Π now potentially refers to *four* disparate objects. That is, we have defined patterns for graphs, Γ-labelled graphs, surfaces, and disks with Γ-strips. We will make sure to carefully specify what type of pattern we mean if there is any chance of confusion.

Let L be an arc in Σ, let $\mathcal{L}$ be a family of arcs in Σ, and let S be a Γ-strip of Σ. We define the number of times L *uses* S to be the number of Δ-subarcs of L contained in S. The number of times $\mathcal{L}$ *uses* S is the sum over the number of times each member of $\mathcal{L}$ uses S.

The following lemma is analogous to Lemma 7.2.1.

Lemma 7.7.3. *For all $k, n \in \mathbb{N}$ and all finite abelian groups Γ, there exists $t := t(k, \Gamma, n)$ such that if Σ is a disk with n Γ-strips and Π is a topologically realizable k-pattern in Σ, then there is a topological realization $\mathcal{L}$ of Π in Σ such that for each strip S of Σ, $\mathcal{L}$ uses S at most t times.*

Proof. Let $\Sigma := \Delta \cup S_1 \cup \cdots \cup S_n$ be a disk with n Γ-strips, and let Π be a topologically realizable k-pattern in Σ. By shrinking Δ, we may assume that $V(\Pi) \cap \Delta$ is contained on the boundary of Δ. Now by enlarging Δ, we may assume that $V(\Pi) \subseteq bd(\Delta)$. Up to homeomorphism, there are only a finite number of such patterns, so we are done. $\square$

7.8 Disjoint Paths Across a Cylinder

The *cylinder* is the surface $\{(x, y) \in \mathbb{R}^2 : 1 \leq x^2 + y^2 \leq 4\}$. In this section we let Σ denote the cylinder. Let G be a graph drawn on Σ. Let δ_1 and δ_2 be the two holes of Σ. Let $Y = V(G) \cap \delta_1$ and $Z = V(G) \cap \delta_2$. For $A \subseteq Y$, we let $\kappa_G(A, Z)$ be the size of a maximum collection of disjoint A-Z paths. Using Menger's Theorem, one can show that κ_G is the rank function of a matroid on Y. For paths P and Q that intersect, the *product* of P with Q is the path $PQ := PxQ$, where x is the first vertex of P also in Q. By convention, if

P and Q are disjoint A-Z paths, the *region between P and Q* is the (closed) clockwise region from P to Q. Recall that a society is a cyclically ordered finite set. Let S be a society. Recall that an interval of S is a proper subset of consecutive vertices of S. A *contiguous partition* of S is a collection of (non-empty) intervals $I_1, \ldots, I_m$ of S such that

- $\bigcup_{j \in [m]} I_j = S$, and
- for each $j \in [m]$ the first element of I_{j+1} is the successor of the last element of I_j (we regard subscripts modulo m).

Note that we can regard Y as a society, ordered clockwise around δ_1. Let I be an interval of Y. We naturally regard I as linearly ordered, and we call this the *clockwise* ordering of I. The *anti-clockwise* ordering on I is the reverse of the clockwise ordering. If $\mathcal{P}$ is a set of I-Z paths in G, then we can order the paths in $\mathcal{P}$ according to their endpoints in I. Thus, we can order $\mathcal{P}$ either clockwise or anti-clockwise.

Theorem 7.8.1. *Let $A_1, B_1, A_2, B_2, \ldots, A_n, B_n$ be a contiguous partition of Y (in that clockwise order). If $\kappa_G(A_i, Z) = |A_i|$ and $\kappa_G(B_i, Z) \geq 2\sum_{j=1}^{n} |A_j|$ for each $i \in [n]$, then $\kappa_G(\bigcup_{i=1}^{n} A_i, Z) = \sum_{i=1}^{n} |A_i|$.*

Proof. By hypothesis, for each $i \in [n]$, there exists a collection $\mathcal{A}_i$ of $|A_i|$ disjoint A_i-Z paths. If the paths in $\mathcal{A} := \bigcup_{i=1}^{n} \mathcal{A}_i$ are disjoint, we are done. Otherwise, for each $i \in [n]$, we will reroute the paths in $\mathcal{A}_i$ to obtain a collection $\mathcal{A}'_i$ of disjoint A_i-Z paths, so that $\bigcup_{i=1}^{n} \mathcal{A}'_i$ is also a collection of disjoint paths.

Let $B := \bigcup_{i=1}^{n} B_i$. Since κ_G is the rank function of a matroid, we can greedily choose a collection $\mathcal{B}$ of disjoint B-Z paths such that

- $|\mathcal{B}| = 2\sum_{i=1}^{n} |A_i|$, and
- For each i, $\mathcal{B}$ contains exactly $|A_i| + |A_{i+1}|$ paths with an endpoint in B_i.

The idea is to use the paths in $\mathcal{B}$ to reroute the paths in $\mathcal{A}$. Let $\mathcal{B}_i$ be the paths in $\mathcal{B}$ with an endpoint in B_i and let $m_i := |A_i|$.

Label the paths of $\mathcal{A}_1 := \{P_1, \ldots, P_{m_1}\}$ clockwise, the paths of $\mathcal{B}_1 := \{R_1, \ldots, R_{m_1+m_2}\}$ clockwise, and the paths of $\mathcal{B}_n := \{L_1, \ldots, L_{m_n+m_1}\}$ anti-clockwise.

We will reroute the paths in $\mathcal{A}_1$ so that they are all between L_{m_1} and R_{m_1}. Towards a contradiction suppose not. The crux of the proof is the following claim.

Claim. *For any $i < j$, if the paths $P_i, \dots, P_j$ are not all in between L_{j-i+1} and R_{j-1+1} then either*

- *$P_i \cap L_{j-i+1} \neq \emptyset$ and $P_i L_{j-i+1} \cap R_{j-i+1} = \emptyset$, or*
- *$P_j \cap R_{j-i+1} \neq \emptyset$, and $P_j R_{j-i+1} \cap L_{j-i+1} = \emptyset$.*

SUBPROOF. Suppose that the paths $P_i, \dots, P_j$ are not all in between L_{j-i+1} and R_{j-1+1}. By planarity, it follows that P_i or P_j must intersect L_{j-i+1} or R_{j-i+1}. By symmetry we may assume that P_i intersects L_{j-i+1} or R_{j-i+1}. First suppose that P_i intersects L_{j-i+1}. Then we are done unless $P_i L_{j-i+1} \cap R_{j-i+1} \neq \emptyset$. However, this implies that P_i also intersects R_{j-i+1}, and that it does so before it intersects L_{j-i+1}. Therefore, $P_j \cap R_{j-i+1} \neq \emptyset$, and $P_j R_{j-i+1} \cap L_{j-i+1} = \emptyset$, as required. The remaining case is if P_i intersects R_{j-i+1}, but not L_{j-i+1}. In this case we again we have $P_j \cap R_{j-i+1} \neq \emptyset$, and $P_j R_{j-i+1} \cap L_{j-i+1} = \emptyset$. $\square$

We now apply the above claim with $i = 1$ and $j = m_1$. We conclude that all paths of $\mathcal{A}_1$ are indeed between L_{m_1} and R_{m_1} unless

- $P_1 \cap L_{m_1} \neq \emptyset$ and $P_1 L_{m_1} \cap R_{m_1} = \emptyset$, or
- $P_{m_1} \cap R_{m_1} \neq \emptyset$ and $P_{m_1} R_{m_1} \cap L_{m_1} = \emptyset$.

By symmetry, we may assume $P_1 \cap L_{m_1} \neq \emptyset$ and $P_1 L_{m_1} \cap R_{m_1} = \emptyset$. We replace P_1 by $P_1 L_{m_1}$. We can then inductively continue rerouting by applying the claim to $P_2, \dots, P_{m_1}$.

By repeating the above argument, for each $i \in [n]$ we obtain a family $\mathcal{A}'_i$ of disjoint A_i-Z paths such that for all i

- $|\mathcal{A}'_i| = |A_i|$, and
- The paths in $\mathcal{A}'_i$ intersect at most $|A_i|$ paths of $\mathcal{B}_i$ and at most $|A_i|$ paths of $\mathcal{B}_{i-1}$.

It immediately follows that the family $\mathcal{A}' := \bigcup_{i=1}^n \mathcal{A}'_i$ is disjoint, since $|\mathcal{B}_i| \geq |A_i| + |A_{i+1}|$ for each i. $\square$

Let C_m be a cycle of length m and P_n be a path of length $n \geq 2$. We define the (m, n)-*cylindrical grid* to be the graph $C_m \times P_{n-1}$. The two cycles of length m in $C_m \times P_{n-1}$ that pass through only degree 3 vertices are called the *boundary cycles*.

Let Π be a pattern, such that $V(\Pi)$ is a society. We say that Π is *cross-free* if there do not exist distinct points a, b, c, d of $V(\Pi)$ (occurring in that cyclic order) such that $\{a, c\} \in \Pi$ and $\{b, d\} \in \Pi$.

Theorem 7.8.2. *Let G be the $(2k, k)$-cylindrical grid with boundary cycles C_1 and C_2. If Π is a k-pattern with $V(\Pi) = V(C_1)$ that is cross-free, then Π is realizable in G.*

Proof. Let Π be a cross-free k-pattern of $X := V(C_1)$. Since Π is cross-free, we can find an element $\{s, t\} \in \Pi$ such that s and t are consecutive vertices of X. We link s and t directly via the edge st. By regarding $H := C_{2k-2} \times P_{k-1}$ as a minor of $G \backslash \{s, t\}$ in the natural way we reduce to a Π'-linkage problem in H, with $|\Pi'| = |\Pi| - 1$. By induction, H does indeed have a Π'-linkage, so G has a Π-linkage. $\square$

Chapter 8

Redundant Vertices in Surfaces

Let G be a Γ-labelled graph and let Π be a pattern in G. Recall that a vertex $v \in V(G)$ is *essential for* Π, if Π is realizable and $v \in V(\mathcal{P})$ for any realization $\mathcal{P}$ of Π in G. A vertex is *redundant for* Π if it is not essential for Π.

In Chapter 6, we discussed redundant vertices in clique-minors. This chapter addresses another instance where we can certify that a vertex is redundant. That is, we consider redundant vertices for Γ-labelled graphs drawn in a surface.

8.1 The Main Theorem

Let G be a Γ-labelled graph embedded in a surface Σ, let Π be a pattern in G, and let f be a face of G. We may regard the boundary of f as a closed walk W in G. We define the *group-value* (or just *value*) of f to be $\gamma_G(W)$. We let $\gamma_G(f)$ denote the value of f. A vertex $v \in V(G)$ is *l-protected in* Σ (with respect to Π) if

- $\gamma_G(f) = 0$, for every face f of G in Σ;
- there are l vertex disjoint cycles $C_1, \ldots, C_l$ of G, bounding discs $\Delta_1, \ldots, \Delta_l$ in Σ with $v \in \Delta_1 \subset \Delta_2 \subset \cdots \subset \Delta_l$;
- $V(\Pi)$ is disjoint from $int(\Delta_l)$.

We refer to $C_1, \ldots, C_l$ as the *circuits protecting* v. Our aim is to prove the following theorem.

Theorem 8.1.1. *For all surfaces Σ, all finite abelian groups Γ and all $k \in \mathbb{N}$, there exists a constant $\mu := \mu(k, \Gamma, \Sigma)$ such that if G is a Γ-labelled graph embedded in Σ, Π is a k-pattern in G, and $v \in V(G)$ is a μ-protected vertex in Σ with respect to Π, then v is redundant.*

8.2 Proof of the Theorem

In this section we prove the main result. This proof is based on an unpublished proof of Thor Johnson and Paul Seymour. To apply induction, it turns out that it is more useful to work with disks with strips rather than surfaces.

Let $\Sigma := \Delta \cup S_1 \cup \ldots S_n$ be a disk with n strips, G be a Γ-labelled graph embedded in Σ, and Π be a pattern in G. We say that a vertex $v \in V(G)$ is *l-insulated in* $\Delta(\Sigma)$ (with respect to Π) if

- $\gamma_G(F) = 0$, for every face F of G in Σ;
- there are l vertex disjoint cycles $C_1, \ldots, C_l$ of G, bounding discs $\Delta_1, \ldots, \Delta_l$ in Δ with $v \in \Delta_1 \subset \Delta_2 \subset \cdots \subset \Delta_l = \Delta$;
- $V(\Pi)$ is disjoint from $int(\Delta_l)$;
- C_i is an induced subgraph of $G \cap \Delta$ for each $i \in [l]$.

In particular, if we regard Σ as a surface, then an l-insulated vertex is an l-protected vertex, but not necessarily vice versa. We refer to $C_1, \ldots, C_l$ as the *circuits insulating* v, and we call v the *insulated vertex*.

We prove Theorem 8.1.1 as a corollary of the following result.

Theorem 8.2.1. *For all finite abelian groups Γ and all $k, n \in \mathbb{N}$ there exists a constant $g := g(k, n, \Gamma)$ such that if G is a Γ-labelled graph embedded in a disk with n strips Σ, Π is a k-pattern in G, $v \in V(G)$ is a g-insulated vertex in $\Delta(\Sigma)$ with respect to Π, and $V(\Pi) \subseteq bd(\Sigma)$, then v is redundant.*

The proof of Theorem 8.2.1 is rather lengthy, so we defer it until the next section. It is however, relatively straightforward to derive Theorem 8.1.1 from Theorem 8.2.1, which we now proceed to do.

PROOF OF 8.1.1 FROM 8.2.1. Let g be the function given in Theorem 8.2.1. We will prove that $\mu(k, \Gamma, \Sigma) := g(k, \Gamma, 4k + N^*(\Sigma, 1))$ suffices, where N^* is the function from Lemma 7.1.3. Let (G, Π, Σ) be a counterexample with $|V(G)| + |E(G)|$ minimal. That is, G is a Γ-labelled graph embedded on a surface Σ, Π is a k-pattern in G and $v \in V(G)$ is a μ-protected ($\mu := g(k, \Gamma, 4k + N^*(\Sigma, 1))$) vertex in Σ with respect to Π, but yet v is essential.

Let $C_1, \ldots, C_\mu$ be the circuits protecting v, bounding disks $\Delta_1, \ldots, \Delta_\mu$ in Σ, such that $\sum_{i=1}^{\mu} |V(C_i)|$ is minimum. Let $\mathcal{P}$ be a realization of Π in G and let H be the subgraph of G composed of $C_1 \cup \cdots \cup C_\mu$.

Claim 1. *$V(G) = V(H)$.*

SUBPROOF. Suppose not. First observe that by deleting any vertices of G which are not in $V(H) \cup V(\mathcal{P})$ we contradict the minimality of G. Similarly, if $e = ab$ is an edge of a path in $\mathcal{P}$ and $a \notin V(H)$, then we can shift to make $\gamma_G(e) = 0$ and then contract e onto b. □

Observe that Claim 1 implies that $V(\Pi) \subseteq V(C_\mu)$.

Claim 2. *If e is an edge of G contained in Δ_μ, then either $e \in E(H)$ or the ends of e are not contained in $V(C_i)$ for any $i \in [\mu]$.*

SUBPROOF. Towards a contradiction suppose $e \subseteq \Delta_\mu$, $e \notin E(H)$, and $ends(e) \subseteq V(C_j)$ for some $j \in [\mu]$. Since $\sum_{i=1}^{\mu} |V(C_i)|$ is minimum, it must be that the ends of e are consecutive vertices of $V(C_j)$. Recall that by hypothesis, every face of G has group-value zero. Therefore, by flipping if necessary, there exists another edge $e' \in E(G)$ with the same head, tail, and group-value as e. Therefore, $(G \backslash e, \Pi, \Sigma)$ is a smaller counterexample. □

We now consider an edge e of G outside Δ_μ. We say that e is *contractible*, if e and a subpath of C_μ bounds a disk in Σ. Otherwise, e is *non-contractible*.

Claim 3. *There are at most $2k$ homotopy classes of contractible edges.*

SUBPROOF. Let e be a contractible edge. Let Q be a subpath of C_μ such that $Q \cup \{e\}$ bounds a disk in Σ. Note that some internal vertex of Q must be in $V(\Pi)$, for otherwise we can delete $E(Q)$ from G and replace C_μ by $(C_\mu \backslash E(Q) \cup \{e\}$. Now let $\{e_1, \ldots, e_n\}$ be a set of contractible edges that are pairwise non-homotopic. Since each e_i is contractible, no two of these edges can cross. For each i, let Q_i be a subpath of C_μ such that $Q_i \cup \{e_i\}$ bounds a disk in Σ. Observe that if $Q_i \subseteq Q_j$, then e_i and e_j are homotopic. Letting I_i be the set of internal vertices of Q_i for each $i \in [n]$, we have that

$$\mathcal{I} := \{I_i : i \in [n]\}$$

is a family of disjoint subsets, each of which contains a vertex in $V(\Pi)$. So,

$$n = |\mathcal{I}| \leq |V(\Pi)| \leq 2k,$$

as required. $\square$

Claim 4. *There are at most $N^*(\Sigma, 1)$ homotopy classes of non-contractible edges.*

SUBPROOF. Let $\mathcal{E}$ be a maximal family of representatives for the homotopy classes of non-contractible edges. Contract Δ to a point x, and let Σ^* be the resulting surface and $\mathcal{E}^*$ be the resulting family of curves. Evidently, Σ^* is homeomorphic to Σ, and $\mathcal{E}^*$ is a 1-family of pairwise non-homotopic arcs in Σ^* (with respect to the basepoint x). By definition of N^*, we have that

$$|\mathcal{E}| = |\mathcal{E}^*| \leq N^*(\Sigma, 1),$$

as required. $\square$

At this point, we may regard G as being embedded in a disk with $t := 2k + N^*(\Sigma, 1)$ strips $\Sigma' := \Delta \cup S_1 \cup \cdots \cup S_t$, where $\Delta_\mu = \Delta$. Unfortunately, some vertices of $V(\Pi)$ may not be on the boundary of Σ'. However, if some vertex $x \in V(\Pi)$ is on a non-corner point of S_i, we can split S_i into two strips and place x on the boundary of one of the new strips. Note that there are at most $2k$ such vertices x. Therefore, we have shown

Claim 5. *G is embedded in a disk with at most $t + 2k = 4k + N^*(\Sigma, 1)$ strips $\Sigma'' := \Delta \cup S_1 \cup \cdots \cup S_{t+2k}$, v is a $g(k, \Gamma, t + 2k)$-insulated vertex in $\Delta(\Sigma'')$ with respect to Π, and $V(\Pi) \subseteq bd(\Sigma'')$.*

By definition of the function g, we have that v is indeed redundant for Π. $\square$

8.3 Proof of the Auxiliary Result

As promised, we give the proof of the following auxiliary result, which we restate here.

Theorem 8.2.1. *For all finite abelian groups* Γ *and all* $k, n \in \mathbb{N}$ *there exists* $g := g(k, \Gamma, n)$ *such that if* G *is a* Γ*-labelled graph embedded in a disk with* n *strips* Σ, Π *is a* k*-pattern in* G, $v \in V(G)$ *is a* g*-insulated vertex in* $\Delta(\Sigma)$ *with respect to* Π, *and* $V(\Pi) \subseteq bd(\Sigma)$, *then* v *is redundant.*

Proof. We establish the existence of $g(k, \Gamma, n)$ via induction on n. Let $t(k, \Gamma, n)$ be the function given in Lemma 7.2.1, and let

$$m := (4n + 1)t(k, \Gamma, n) + 8k.$$

We will show that any function $g(k, \Gamma, n)$ which satisfies

- $g(k, \Gamma, 0) \geq k^2$, and for each $n \in \mathbb{N}$,
- $g(k, \Gamma, n) - g(2k + m(2n + 1)^{2n(m-1)}, \Gamma, n - 1) \geq 2k + nt(k, \Gamma, n)$

will suffice.

Let (G, Π, Σ) be a counterexample with $|E(G)|$ minimal. That is, G is a Γ-labelled graph embedded on a disk with n strips Σ, Π is a k-pattern in G, $v \in V(G)$ is a g-insulated ($g := g(k, \Gamma, n)$) vertex in Σ with respect to Π, $V(\Pi) \subseteq bd(\Sigma)$, but yet v is essential.

Let $\Sigma = \Delta \cup S_1 \cup \cdots \cup S_n$ and let $C_1, \ldots, C_g$ be circuits insulating v, bounding disks $\Delta_1, \ldots, \Delta_g$ in Σ, where $\Delta_g = \Delta$. Let H be the subgraph of G composed of $C_1 \cup \cdots \cup C_g$. Let $\mathcal{P}$ be a realization of Π in G.

Claim 1. $E(H) \cap E(\mathcal{P}) = \emptyset$.

SUBPROOF. Shift and then contract any edges e in $E(H) \cap E(\mathcal{P})$. $\square$

Claim 2. $V(G) = V(H) \cup (V(G) \cap bd(\Sigma))$.

SUBPROOF. If e has ends a and b and $b \notin V(H) \cup (V(G) \cap bd(\Sigma))$, then we shift to make $\gamma_G(e) = 0$ and contract e onto a. $\square$

Claim 3. $V(G) = V(H)$.

SUBPROOF. Let $e \in E(G)$ be contained in a strip S_i, and let a and b be the ends of e. If both a and b are in $V(H)$, there is nothing to prove. Otherwise, by the previous claim, we may assume that $b \in bd(\Sigma)$, but that b is not on either end of S_i. We shift to make $\gamma_G(e) = 0$, pull e slightly away from $bd(\Sigma)$, and then contract e onto a. $\square$

Claim 4. *Every cycle in $G \cap \Delta$ has value 0.*

SUBPROOF. Let C be a cycle in $G \cap \Delta$. Note that C bounds a disk Δ' in Δ. The value of C is equal to the sum of the values of all the faces contained in Δ'. Hence, C has value 0. $\square$

Claim 5. *We can shift so that every edge in $\Delta(\Sigma)$ is zero-labelled.*

SUBPROOF. Follows from Claim 4 and Lemma 1.2.2. $\square$

Henceforth, we assume that all the edges in $\Delta(\Sigma)$ are zero-labelled. We now examine how the paths in $\mathcal{P}$ go through Δ. If $x \in V(C_j)$, we define the *level* of x, denoted $l(x)$, to be j. A path $P = x_0x_1 \ldots x_q$ in G is *decreasing* if $P \subseteq \Delta$ and

$$l(x_0) \geq l(x_1) \cdots \geq l(x_q).$$

A *hill* of a path P is a subpath $J := aPc$ of P where

- $l(a) = l(c)$, and
- $l(b) > l(a)$ for all $b \in V(P)$ strictly between a and c.
- J and a subpath of $C_{l(a)}$ bound a disk in Σ.

The *sea level* $l(J)$ of J, is defined to be $l(a)$.

Claim 6. *No path P in $\mathcal{P}$ contains a hill.*

SUBPROOF. Suppose not. Among all hills of all paths in $\mathcal{P}$, choose J with the lowest sea level $l(J)$. Let J have ends a and b and suppose that J is a hill of $P \in \mathcal{P}$. By choice of J, there is a subpath K of $C_{l(J)}$ also with ends a and b, such that no path in $\mathcal{P}$ uses an internal vertex of K. Otherwise, there would be a hill of a path in $\mathcal{P}$ with lower sea level than J. Therefore, we can replace P in $\mathcal{P}$ by $(P \backslash J) \cup K$. Letting e be the edge of J incident to a, we conclude that $G \backslash e$ is a smaller counterexample, a contradiction. $\square$

Let us now analyze the edges that are outside Δ. For each strip S_i of Σ, we let $E(S_i)$ be the edges of G contained in S_i. Recall that an edge e passes through S_i if $e \subseteq S_i$ and its ends are on different ends of S_i.

Claim 7. *For each $i \in [n]$, if $e \in E(S_i)$, but e does not pass through S_i, then $\gamma_G(e) = 0$.*

SUBPROOF. Enlarge the disk Δ and apply the argument in Claim 5. $\square$

By flipping edges if necessary we may assume that for each strip S_i, the edges that pass through S_i each pass through in the positive direction.

Claim 8. *For each strip S_i, all the edges that pass through S_i have the same group-value.*

SUBPROOF. For each strip S_i we construct a graph $G^*(S_i)$ as follows. The vertex set of $G^*(S_i)$ is the set of edges of G that pass through S_i. We define e to be adjacent to f in $G^*(S_i)$ if and only if e and f are both on the boundary of the same face of G. If e and f are adjacent in $G^*(S_i)$, then we claim they have the same group-label. To see this, let F be the face of G such that $e \cup f \subseteq \overline{F}$. Note that all other edges of G on the boundary of F have group-label zero. Thus, $\gamma_G(e) = \gamma_G(f)$, as required. It now follows that all edges that pass through S_i have the same group-label since $G^*(S_i)$ is connected (actually a path). $\square$

At this point we now regard each strip S_i as a Γ-labelled strip, where the group-value of S_i is the group-value of any edge that passes through S_i. If we regard Π as a pattern in Σ instead of a pattern in G, then evidently there is a topological realization of Π in Σ, since there is realization of Π in G. Lemma 7.7.3 asserts that there is a topological realization of Π that passes through each strip only a few times. Let $t := t(k, \Gamma, n)$ be the function given in Lemma 7.7.3, and let $\mathcal{L}$ be such a topological realization of Π in Σ. The pivotal idea is to try and realize $\mathcal{L}$ in G.

Let $Y = V(C_g)$, $N := g(k + m(2n+1)^{2n(m-1)}, \Gamma, n-1)$ and $Z = V(C_N)$. We define a matroid M on Y by $r_M(A) = \kappa_{G \cap \Delta}(A, Z)$ for all $A \subseteq Y$.

For each strip S of Σ, we choose a maximum matching $m(S)$ contained in the edges that pass through S. We let $V(S)$ be the vertices covered by $m(S)$. We partition $V(S)$ as $V_0(S) \cup V_1(S)$, according to the end of S a vertex belongs to. For $i \in \mathbb{Z}_2$, we let $M_i(S)$ be the restriction of M to $V_i(S)$

respectively. If we identify the endpoints of each edge that passes through S, then we can naturally regard $M_0(S)$ and $M_1(S)$ as matroids on the same ground set. For $X \subseteq V_i(S)$ we let $clone(X)$ be the copy of X in $V_{i+1}(S)$.

We first consider the case when $M_0(S_i)$ and $M_1(S_i)$ have a large common independent set, for each strip S_i of Σ.

Case 1. *$M_0(S_i)$ and $M_1(S_i)$ have a common independent set of size $m := (4n+1)t(k,\Gamma,n) + 8k$ for each strip S_i of Σ.*

We will need the following claim.

Claim 9. *If $A \subseteq Y$ is independent in M, then there is a family of $|A|$ disjoint decreasing A-Z paths in $G \cap \Delta$.*

SUBPROOF. Choose a family $\mathcal{Q}$ of $|A|$ disjoint A-Z paths in $G \cap \Delta$ with the minimum number of total hills. If $\mathcal{Q}$ has no hills, we are done. Otherwise, among all hills of all paths in $\mathcal{Q}$, choose J with the lowest sea level $l(J)$. Let J have ends a and b and suppose that J is a hill of $Q \in \mathcal{Q}$. By choice of J, there is a subpath K of $C_{l(J)}$ also with ends a and b, such that no path in $\mathbb{Q}$ uses an internal vertex of K. Therefore, we can replace Q in $\mathcal{Q}$ by $(Q\backslash J) \cup K$, contradicting the choice of $\mathcal{Q}$. $\square$

Let us recall some notation. By orienting $bd(\Delta)$ clockwise, we may regard it as a cyclically ordered set. For $a, b \in bd(\Delta)$, $[a, b]$ is the set of points between a and b in the ordering. That is, it is the clockwise subarc of $bd(\Delta)$ from a to b. Since $V(\Pi)$ is disjoint from each strip (except possibly at the corners), it follows that the corners of the strips of Σ induce a contiguous partition $X_1, \ldots, X_l$ of $V(\Pi)$. Consider an arbitrary set X_i of the partition. Label its vertices $x_1, \ldots, x_p$ (clockwise). In particular, this implies that $[x_2, x_{p-1}]$ is disjoint from each strip of Σ.

Recall that $\mathcal{P}$ is a realization of Π in G. For each x_i, let $\mathcal{P}(x_i)$ be the (unique) path in $\mathcal{P}$ starting from x_i. Define $\omega(x_i)$ to be the number of insulating cycles that $\mathcal{P}(x_i)$ intersects before it uses an edge outside of Δ.

Claim 10. *For each $q \in [p]$, $\omega(x_q) \geq \min\{q, p-q+1\}$.*

SUBPROOF. We proceed by induction on $\min\{q, p-q+1\}$. Clearly the claim holds for $q \in \{1, p\}$. Consider an arbitrary x_q. By symmetry we may assume that $q \leq \frac{p}{2}$ and we inductively assume that $\omega(x_{q-1}) \geq q-1$ and $\omega(x_{p-q+2}) \geq q-1$. Towards a contradiction assume that $\omega(x_q) \leq q-1$. Let

a be the second vertex of $\mathcal{P}(x_q)$ which is on $bd(\Delta)$ (x_q is the first). Let Q be the subpath of $\mathcal{P}(x_q)$ from x_q to a. Note that $Q \cup [x_q, a]$ and $Q \cup [a, x_q]$ both bound disks in Δ. We denote them as $\Delta(\Rsh)$ and $\Delta(\Lsh)$, respectively. We say that a disk is *small* if it does not contain v (the insulated vertex). Clearly, exactly one of $\Delta(\Rsh)$ or $\Delta(\Lsh)$ is small. There are various cases depending where a lies on $bd(\Delta)$ and which of $\Delta(\Rsh)$ or $\Delta(\Lsh)$ is small.

Subclaim. $\Delta(\Lsh)$ *is not small.*

SUBPROOF. Towards a contradiction assume $\Delta(\Lsh)$ is small. We first prove that $a \notin [x_{q-1}, x_q]$. If so, then we can reroute $\mathcal{P}(x_q)$ through $[a, x_q]$. Thus, letting e be the first edge of $\mathcal{P}(x_q)$ we see that $G \backslash e$ is a smaller counterexample. So, we have shown that $x_{q-1} \in [a, x_q]$. Since $\omega(x_{q-1}) \geq q-1$, the only way to avoid a contradiction is if $\mathcal{P}(x_q)$ actually connnects x_q to x_{q-1} within Δ. But, we can then delete the first edge of $\mathcal{P}(x_q)$ and connect x_q to x_{q-1} directly via $[x_{q-1}, x_q]$. $\square$

Subclaim. $\Delta(\Rsh)$ *is not small.*

SUBPROOF. Towards a contradiction assume $\Delta(\Rsh)$ is small. As in the proof of the previous subclaim we have $a \notin [x_q, x_{q+1}]$. We now show that $a \notin [x_q, x_p]$. If so, then there must exist an index $r \geq q$ such that some path of $\mathcal{P}$ connects up x_r and x_{r+1} within $\Delta(\Rsh)$. Deleting the first edge of $\mathcal{P}(x_r)$ and rerouting $\mathcal{P}(x_r)$ through $[x_r, x_{r+1}]$ gives a contradiction. The only remaining possibility is if $a \in [x_p, x_1]$. This forces either $\omega(x_{p-q+2}) < q-1$ or $\mathcal{P}(x_q)$ must connect x_q to x_{p-q+2} within Δ. The first possibility contradicts our inductive hypothesis. So, that leaves $\mathcal{P}(x_q) = \mathcal{P}(x_{p-q+2})$. But again, this implies that there is some index $s \in \{q, \ldots, p-q+1\}$ such that some path of $\mathcal{P}$ connects up x_s and x_{s+1} within $\Delta(\Rsh)$, which we have already seen is a contradiction. $\square$

This completes the proof of the claim, since $\Delta(\Lsh)$ and $\Delta(\Rsh)$ cannot both be small. So, $\omega(x_q) \geq \min\{q, p-q+1\}$, as required. $\square$

We remark that the proof of Claim 10 does not actually rely on the hypothesis in Case 1.

Claim 11. X_i *is independent in* M.

SUBPROOF. Let S be an arbitrary strip of Σ. Since we are in Case 1, there exists an $M_0(S)$-independent subset I of size $|X_i| = p$. By Claim 9, there is a family $\mathcal{Q}$ of p disjoint decreasing I-Z paths in $G \cap \Delta$. Label these paths as $Q_1, \dots, Q_p$ (counter-clockwise). We will use $\mathcal{Q}$ to construct p disjoint X_i-Z paths in $G \cap \Delta$. By Claim 10, for each $q \in [p]$, $\omega(x_q) \geq \min\{q, p - q + 1\}$. So for each $q \in \{1, \dots, \lceil p/2 \rceil\}$ we can define a path $\mathcal{R}(x_q)$ as follows:

- follow $\mathcal{P}(x_q)$ until it intersects $C_{g-(q-1)}$;
- follow $C_{g-(q-1)}$ (counter-clockwise) until intersecting $Q_{\lceil p/2 \rceil-(q-1)}$;
- follow $Q_{\lceil p/2 \rceil-(q-1)}$ until reaching Z.

For $q \in \{p, p-1, \dots, \lceil p/2 \rceil + 1\}$ we define $\mathcal{R}(x_q)$ as follows:

- follow $\mathcal{P}(x_q)$ until it intersects C_{g-p+q};
- follow C_{g-p+q} (clockwise) until intersecting $Q_{\lceil p/2 \rceil+p-q+1}$;
- follow $Q_{\lceil p/2 \rceil+p-q+1}$ until reaching Z.

It is easy to verify that

$$\mathcal{R} := \{\mathcal{R}(x_q) : q \in [p]\}$$

is a family of disjoint X_i-Z paths in $G \cap \Delta$. $\square$

Recall that X_i was chosen arbitrarily. Therefore, we have shown

Claim 12. *X_i is M-independent for all $i \in [l]$.*

Next we show that $\bigcup X_i := V(\Pi)$ is actually M-independent. In fact, we prove the following much stronger claim.

Claim 13. *For each strip S_i of Σ there exists a subset K_i of $V_0(S_i)$ of size $t := t(k, \Gamma, n)$ such that $V(\Pi) \cup \bigcup_{i \in [n]}(K_i \cup clone(K_i))$ is independent in M.*

SUBPROOF. Of course we are in the case when $M_0(S_i)$ and $M_1(S_i)$ have a large common independent set for each strip S_i of Σ. So, for each $i \in [n]$ let J_i be an independent set of size $(4n+1)t + 8k$ in $M_0(S_i)$, such that $clone(J_i)$ is also independent in $M_1(S_i)$. We partition J_i into three sets J_i^1, J_i^2 and J_i^3

where J_i^1 are the first $2(nt+2k)$ points, J_i^2 are the next t points and J_i^3 are the last $2(nt+2k)$ points (in the clockwise order). Let

$$\mathcal{A} := \{J_i^2 : i \in [n]\} \cup \{clone(J_i^2) : i \in [n]\} \cup \{X_i : i \in [l]\},$$

and

$$\mathcal{B} := \{J_i^k : i \in [n], k \in \{1,3\}\} \cup \{clone(J_i^k) : i \in [n], k \in \{1,3\}\}.$$

Observe that each set in $\mathcal{A}$ is indeed M-independent, and that for any $B \in \mathcal{B}$ we have

$$r_M(B) = 2(nt+2k) = 2\sum_{A \in \mathcal{A}} |A|.$$

Therefore, we are in perfect position to apply Theorem 7.8.1, and conclude that $\bigcup_{A\in\mathcal{A}} A$ is M-independent. Setting $K_i = J_i^2$ for each $i \in [n]$ gives the result. □

We can now attempt to realize the topological linkage $\mathcal{L}$ in G. We may assume that $\mathcal{L}$ intersects $bd(\Delta)$ only at vertices in $\mathcal{A}$. Let $G' := G \cap (\Delta \backslash int(\Delta_N))$. By removing all the strips from Σ and keeping track of how the paths in $\mathcal{P}$ pass through the strips, we are left with a Π'-linkage problem in the disk Δ, where $V(\Pi') \subseteq V(\mathcal{A})$. Note that this is just a linkage problem in $\widetilde{G'}$, since all edges of G' are zero-labelled.

By Claim 13, we have that $V(\mathcal{A})$ is M-independent. Therefore, by Claim 9, there exists a family of $|V(\mathcal{A})|$ disjoint decreasing $V(\mathcal{A})$-Z paths (recall $Z = V(C_N)$) in G'. These decreasing paths, together with the insulating circuits $C_g, C_{g-1}, \ldots, C_N$ form a large cylindrical-grid minor H' in G'. Since

$$g - N \geq 2k + nt \geq |\mathcal{A}|,$$

Theorem 7.8.2 implies that G' actually has a Π'-linkage. Thus, G' also has a Π-linkage, and v is redundant for Π in G since $v \notin V(G')$. This case is hence complete. □

The remaining case is if $M_0(S_i)$ and $M_1(S_i)$ do not have a large common independent set, for some strip S_i of Σ. By re-indexing, we may assume that $S_i = S_1$.

Case 2. *$M_0(S_1)$ and $M_1(S_1)$ do not have a common independent set of size $m = (4n+1)t(k,\Gamma,n) + 8k$.*

The idea in this case is to reduce the number of strips. Recall that $Y = V(C_g)$, $N := g(k + m(2n+1)^{2n(m-1)}, \Gamma, n-1)$ and $Z = V(C_N)$. Since $M_0(S_1)$ and $M_1(S_1)$ do not have a common independent set of size m, by Theorem 2.6.1 (Matroid Intersection Theorem), there is a partition $\{A, B\}$ of $V_0(S_1)$ such that

$$r_{M_0(S_1)}(A) + r_{M_1(S_1)}(clone(B)) < m.$$

That is, there exist subsets T and U of $V(G \cap \Delta)$ such that

- T separates A from Z in $G \cap \Delta$,
- U separates $clone(B)$ from Z in $G \cap \Delta$, and
- $|T| + |U| < m$.

There are three subcases, depending where T and U lie with respect to the insulating circuits $C_N, \ldots, C_g$. Recall that the level of a vertex $x \in G \cap \Delta$ is the unique index j such that $x \in C_j$.

Subcase. *The level of each $x \in T \cup U$ is at most $g - m$.*

This implies that $T \cup U$ actually separates Y from Z in $G \cap \Delta$. We will reduce the Π-linkage problem in G to a Π'-linkage problem in the disk Δ_N. We let $G' := G \cap \Delta_N$. It remains to explain how to construct Π'. Let P be a path in $\mathcal{P}$. We say that P is *outside* Δ_N if it is disjoint from $int(\Delta_N)$. We say that P is *inside* Δ_N if it is contained in Δ_N. We remark that a path outside of Δ_N may not be contained entirely within Δ, and we regard a single vertex as a path. We define $\mathcal{O}(P)$ to be the family of maximal subpaths of P among those outside Δ_N, and $\mathcal{I}(P)$ to be the family of maximal subpaths of P among those inside Δ_N. Note that the paths in $\mathcal{O}(P)$ are vertex disjoint, as are those in $\mathcal{I}(P)$. Furthermore, P can be written as $O_1 I_1 \ldots O_r I_r O_{r+1}$, where

- $O_i \in \mathcal{O}(P)$, for each $i \in [r+1]$,
- $I_i \in \mathcal{I}(P)$, for each $i \in [r]$,
- the last vertex of O_i is the first vertex of I_i, for each $i \in [r]$, and
- the last vertex of I_i is the first vertex of O_{i+1} for each $i \in [r]$.

For each $i \in [r]$ let x_i be the last vertex of O_i, and y_i be the last vertex of I_i. Place $(x_i, y_i, 0)$ into Π' for each $i \in [r]$, and then repeat for each path of $\mathcal{P}$. Clearly, G' has a Π'-linkage if and only if G has a Π-linkage. We prove that $|\Pi'|$ is not too large by appealing to Claim 6. Thus, for each $P \in \mathcal{P}$, every path in $\mathcal{O}(P)$ must use an edge outside Δ since P contains no hills. Since $T \cup U$ separates Z from Y in $G \cap \Delta$, every member of $\mathcal{O}(P)$ must use at least two vertices of $T \cup U$. It follows that $\sum_{P \in \mathcal{P}} 2|\mathcal{O}(P)| \leq |T \cup U|$. Therefore, $|V(\Pi')| \leq |T \cup U| < m$. Since $N \geq g(k+m, \Gamma, 0)$, and v is clearly an s-insulated vertex in Δ_N with respect to Π', we conclude that v is redundant for Π'. Hence, v is also redundant for Π. This subcase is complete.

The next subcase is if no element of $T \cup U$ occurs very deep in the disk Δ.

Subcase. *For each vertex x of $T \cup U$, the level of x is at least $g - m + 1$.*

Let $h := g - m + 1$. Again the idea is to reduce to number of strips. We will reduce to a problem in a disk with $n - 1$ strips, where the disk is Δ_h. First, we recall some notation from Chapter 7. A path P is a Δ_h-path if both its ends belong on Δ_h, and it is otherwise disjoint from Δ_h. Note that this clearly implies that the ends of P are on the boundary of Δ_h. For each path P of $\mathcal{P}$, we define $\mathcal{U}(P)$ to be the family Δ_h-subpaths of P. We then define $\mathcal{U}(\mathcal{P}) := \bigcup_{P \in \mathcal{P}} \mathcal{U}(P)$.

Claim 14. *There are at most $(2n+1)^{2n(m-1)}$ homotopy classes of paths in $\mathcal{U}(\mathcal{P})$.*

SUBPROOF. Let $Q \in \mathcal{U}(\mathcal{P})$. Since Q does not contain any hills, there is no subpath K of C_h such that $Q \cup K$ bounds a disk in Σ. In particular, this implies that Q must use an edge outside of Δ and that the homotopy class of Q is determined by how Q passes through the strips of Σ. We remark that the *homology* class of Q only depends on the number of times Q passes through each strip, but for homotopy, order is relevant. Let $\mathcal{A}$ be the alphabet $\{S_1, \ldots, S_n, S_1^{-1}, \ldots, S_n^{-1}\}$. The homotopy class Q, denoted $\mathcal{H}(Q)$, is then naturally encoded by a string of letters from $\mathcal{A}$. We make the convention that if $S_i S_i^{-1}$ or $S_i^{-1} S_i$ appears in $\mathcal{H}(Q)$ for some $i \in [n]$, then we cancel it. With this convention, we prove that the length of $\mathcal{H}(Q)$ (as a string) is not very long.

Subclaim. *Each letter of $\mathcal{A}$ appears at most $2m - 1$ times in $\mathcal{H}(Q)$.*

SUBPROOF. Towards a contradiction assume that some letter α appears at least $2m$ times in $\mathcal{H}(Q)$. By reversing the direction of Q we may assume $\alpha = S_j$, for some $j \in [n]$. Let $e_1, \ldots, e_{2m}$ be edges of Q corresponding to the occurrences of S_j in $\mathcal{H}(Q)$. That is, for each $i \in [2m]$, e_i passes through the strip S_j and Q passes through each e_i in the forward direction. Furthermore, by cancellation, the next edge of Q after e_i passing through a strip cannot pass through S_j in the backward direction. Now, for each $i \in [2m]$, define $x_i := head_G(e_i)$. We re-index so that $x_1, \ldots, x_{2m}$ occur clockwise along one end of the strip S_j. Either x_m occurs before x_{m+1} along Q or vice versa. By symmetry, we assume the former. Let $Q' := x_m Q$, the subpath of Q starting from x_m. Let y be the first vertex of Q' such that the next edge of Q' passes through a strip S_k. Note that y exists since x_m occurs before x_{m+1} along Q. By cancellation $S_k \neq S_j$, so it follows that $y \in [x_{2m}, x_1]$ (recall that $[x_{2m}, x_1]$ is the clockwise subarc of $bd(\Delta)$ from x_{2m} to x_1). Also recall that a region $\mathcal{R}$ in Δ is *big* if it contains the insulated vertex, and is *small* otherwise. Clearly, exactly one of $x_m Q y \cup [y, x_m]$ or $x_m Q y \cup [x_m, y]$ bounds a small region $\mathcal{R}$. It might be that $\mathcal{R}$ is not a disk, since $x_m Q y$ may contain other vertices on $bd(\Delta)$ besides x_m and y. However, the relevant observation is that either $\{x_1, \ldots, x_m\} \subseteq \mathcal{R}$ or $\{x_m, \ldots, x_{2m}\} \subseteq \mathcal{R}$. In either case we get a contradiction, since $x_m Q y$ intersects at most $m-1$ insulating circuits. This proves the subclaim. $\square$

We conclude that the length of $\mathcal{H}(Q)$ (as a string) is at most $2n(m-1)$. We conclude that there are at most $(2n+1)^{2n(m-1)}$ possibilities for $\mathcal{H}(Q)$, which proves the claim. $\square$

We call a homotopy class *thin* if it contains at most m paths in $\mathcal{U}(\mathcal{P})$, otherwise it is *thick*.

Claim 15. *There are at most $n-1$ thick homotopy classes.*

SUBPROOF. Let $\mathcal{H}$ be a homotopy class, represented as a string of letters from $\{S_1, \ldots, S_n, S_1^{-1}, \ldots, S_n^{-1}\}$. Note that $\mathcal{H}$ is not the empty string by Claim 6. If $\mathcal{H}$ is of length at least two, then $\mathcal{H}$ cannot contain more than m paths from $\mathcal{U}(\mathcal{P})$, since each path in $\mathcal{U}(\mathcal{P})$ intersects at most m insulating circuits. Thus, if $\mathcal{H}$ is thick, then $\mathcal{H}$ must be a string of length 1. Up to inversion, this implies that $\mathcal{H} = S_i$, for some $i \in [n]$. However, consider the homotopy class $\mathcal{H}_1$ represented by the string S_1. Recall that we are in the case where $M_0(S_1)$ and $M_1(S_1)$ do not have a large common independent set. Therefore, if $Q \in \mathcal{U}(\mathcal{P})$ has homotopy type S_1, then Q must use a vertex of $U \cup T$. Since, $|U \cup T| < m$, it follows that there are fewer than m paths in $\mathcal{H}_1$. That is, $\mathcal{H}_1$ is thin. This leaves only $n-1$ homotopy classes that may be thick, as required. $\square$

We are now in position to complete this subcase and complete the entire proof. Let $G' := (G \cap \Delta_h) \cup \mathcal{U}(\mathcal{P})$. By Claim 14 we can regard G' as embedded in a disk with at most $l := (2n+1)^{2n(m-1)}$ strips

$$\Sigma' := \Delta_h \cup S'_1 \cup \dots S'_l.$$

We describe how to reduce the Π-linkage problem in G to a Π'-linkage problem in G'. Let $P \in \mathcal{P}$. Let x be the first vertex of P which is on $bd(\Delta_h)$, and let y be the last vertex of P which is on $bd(\Delta_h)$. Observe that since $V(\Pi) \subseteq bd(\Sigma)$, it is easy to see that both x and y can only intersect the new strips $S'_1, \dots, S'_l$ at corner points. Let $\alpha \in \Gamma$ be the group-value of Px, and let $\beta \in \Gamma$ be the group-value of yP. Place $(x, y, \gamma(P) - (\alpha + \beta))$ in Π'. If no such x exists, we do nothing. We then repeat for all paths of $\mathcal{P}$. At first glance it seems as if we have increased the complexity of our problem, since we have more strips than we began with. However, by Claim 15, at most $n-1$ of the strips $S'_1, \dots, S'_l$ are thick. By re-indexing, we may assume that $S'_n, \dots, S'_l$ are all thin. By deleting all the edges contained in $S'_n \cup \dots \cup S'_l$, and keeping track of how the paths in $\mathcal{P}$ pass through $S'_n \cup \dots \cup S'_l$, we reduce to a Π''-linkage in $\Delta_h \cup S'_1 \cup \dots \cup S'_{n-1}$, where $|V(\Pi'')| \leq 2k + m((2n+1)^{2n(m-1)} - n + 1)$. Since v is an h-insulated vertex with respect to Π'', and $h \geq N \geq g(k + m(2n+1)^{2n(m-1)}, \Gamma, n-1)$, it follows that v is redundant for Π'', and hence also for Π. This completes the second subcase.

The remaining subcase is handled by a combination of the previous two techniques and is omitted.

For the base of the induction, one can show using Theorem 7.8.2 that $g(k, \Gamma, 0) = k^2$ suffices for the disk. This completes the proof. $\square$

8.4 Sufficient Conditions

In this section, we describe conditions which are sufficient to guarantee that a Γ-labelled graph G has a Π-linkage. One example of a set of such conditions was given in Theorem 6.1.2, related to large $K(\Gamma, n)$-minors. Another example was given in Theorem 7.8.2, concerning graphs embedded on the cylinder. Here we generalize Theorem 7.8.2 and discuss sufficient conditions for Γ-labelled graphs embedded in a fixed surface Σ.

To better motivate the reader, we first consider the special case when Γ is trivial, that is, just graphs. We start with graphs embedded on the disk, where it is possible to give conditions which are necessary and sufficient.

Let G be a graph embedded on a disk Δ, and let Π be a pattern of G with $V(\Pi)$ contained on the boundary of Δ. Let δ denote the boundary of Δ, and let D be a dual δ-path with ends a and b on δ. For each $\{s, t\} \in \Pi$, we say that D *separates* $\{s, t\}$ if s and t are in different (topological) components of $\delta \backslash \{a, b\}$ or if $\{s, t\} \cap \{a, b\} \neq \emptyset$. If every dual δ-path D separates at most $|V(D)|$ elements of Π, we say that (G, Π) satisfies the *topological cut condition*.

Theorem 8.4.1. *Let G be a graph embedded on a disk Δ, and let Π be a pattern of G with $V(\Pi) \subseteq bd(\Delta)$. Then G has a Π-linkage if and only if Π is topologically realizable in Δ and (G, Π) satisfies the topological cut condition.*

The stated conditions are clearly necessary for G to have a Π-linkage. As we do not need this result, we omit the proof of the converse. See [36, Theorem 3.6].

Theorem 8.4.2. *For any $k \in \mathbb{N}$ and surface Σ, there exist constants $r(k, \Sigma)$ and $w(k, \Sigma)$ with the following property. Let G be a graph 2-cell embedded in $\widehat{\Sigma}$ and let Π be a k-pattern in G where $V(\Pi) \subseteq bd(\Sigma)$. If G has a respectful tangle $\mathcal{T}$ of order at least $r(k, \Sigma)$ such that*

- *Π is topologically realizable in Σ,*
- *if δ_1 and δ_2 are distinct holes of Σ, then $d_\Sigma(\delta_1, \delta_2) \geq w(k, \Sigma)$, and*
- *$V(\Pi)$ is $\mathcal{T}$-independent,*

then G has a Π-linkage.

Proof. The result actually follows quite easily from Theorem 8.1.1. Let μ be the function from Theorem 8.1.1. We will show that

- $r(k, \Sigma) := 4\mu(k, \{0\}, \Sigma) + 6$

- $w(k, \Sigma) := 2\mu(k, \{0\}, \Sigma) + 3$

suffice. Let G, Σ, Π, k, and $\mathcal{T}$ be given as above. The idea is to superfluously add some new edges to G so that the enlarged graph clearly has a Π-linkage. We then apply Theorem 8.1.1 to show that each of the newly added edges is actually redundant for Π, and hence G has a Π-linkage.

Let $\mathcal{H}$ be the set of holes of Σ. For notational convenience we let $n := \mu(k, \{0\}, \Sigma)$. By Corollary 7.5.5, for each $\delta \in \mathcal{H}$ there is a cycle C_δ of G such that

- C_δ bounds a disk Δ_δ in $\widehat{\Sigma}$ with $\delta \subseteq \Delta_\delta$, and
- $d_\Sigma(y, \delta) = n$, for all vertices y of C_δ.

Let X_δ be the vertices of Π on the hole δ. By Lemma 7.5.6, there is a family $\mathcal{P}_\delta$ of $|X_\delta|$ disjoint X_δ-C_δ paths in G. For each $x \in X_\delta$, let $x' \in V(C_\delta)$ be the other endpoint of the path in $\mathcal{P}_\delta$ that contains x. Repeat this for all holes of Σ and define

$$X' := \{x' : x \in V(\Pi)\}.$$

Let Π' be the pattern in G, with $V(\Pi') = X'$ that is naturally induced by Π. Observe that $d_\Sigma(\delta_1, \delta_2) \geq 2n + 3$, for any two distinct holes δ_1 and δ_2 of Σ. Therefore, the family of disks $\mathcal{D} := \{\Delta_\delta : \delta \in \mathcal{H}\}$ is a disjoint family. Let Σ' be the surface obtained from $\widehat{\Sigma}$ by removing the interiors of each disk in $\mathcal{D}$. Since Π is topologically realizable in Σ, evidently Π' is topologically realizable in Σ'. Furthermore, if $G \cap \Sigma'$ has a Π'-linkage, then G has a Π-linkage.

Let $\mathcal{L}$ be a topological realization of Π' in Σ'. *A priori*, we cannot assume that $\mathcal{L}$ intersects $G \cap \Sigma'$ only at vertices, but we may assume that $\mathcal{L}$ intersects $G \cap \Sigma'$ finitely often. Define G' with vertex set $V(G) \cup (\mathcal{L} \cap G)$ and with edge set the arcs in $G \cup \mathcal{L}$ with both endpoints in $V(G')$. Define

$$E_{\mathcal{L}} := \{e \in E(G') : \ e \subseteq L \text{ for some } L \in \mathcal{L}\}.$$

Observe that $G' \backslash E_{\mathcal{L}}$ is a subdivision of G.

We now regard Π as a pattern in G' and show that every edge in $E_{\mathcal{L}}$ is redundant for Π in G'. Let $e \in E_{\mathcal{L}}$. Since $\mathcal{T}$ is a respectful tangle of order at least $4n + 6$, by Corollary 7.5.5 there are vertex disjoint cycles $C_1, \ldots, C_n$ in G' bounding disks $\Delta_1, \ldots, \Delta_n$ in Σ such that $e \subset \Delta_1 \subset \cdots \subset \Delta_n$. Moreover, since $\mathcal{L}$ is disjoint from the interior of each of the disks in $\mathcal{D}$, we may assume that $V(\Pi)$ is disjoint from $int(\Delta_n)$. That is, e is an n-protected edge in Σ with respect to Π. By Theorem 8.1.1, e is redundant for Π.

We now choose another edge in $E_{\mathcal{L}}$ and repeat the same argument with G' replaced by $G' \backslash e$. Proceeding sequentially through $E_{\mathcal{L}}$ we conclude that $G' \backslash E_{\mathcal{L}}$ has a Π-linkage if and only if G' has a Π-linkage. But G' manifestly does have a Π-linkage, by its construction. Therefore, $G' \backslash E_{\mathcal{L}}$ has a Π-linkage. But, $G' \backslash E_{\mathcal{L}}$ is a subdivision of G. Therefore, G also has a Π-linkage, as required. $\square$

Before moving on to the sufficient conditions for group-labelled graphs, it is necessary to take a brief topological interlude.

Let Σ be a surface, and let p be a fixed point of Σ. Let $\pi(\Sigma)$ be the set of homotopy classes of closed curves in Σ with basepoint p. It is not hard to show that there is a natural group structure on $\pi(\Sigma)$ defined via composition of curves. We call $\pi(\Sigma)$ the *fundamental group of* Σ. The *(first) homology group of* Σ, denoted $H_1(\Sigma)$, is the abelianization of $\pi(\Sigma)$. We remark that $\pi(\Sigma)$ is independent of the choice of p and only depends on the homeomorphism class of Σ. See [51] for an agreeable introduction to algebraic topology.

Let G be a Γ-labelled graph 2-cell embedded in a surface Σ. If every face of G has group-value zero, then G induces a natural homomorphism $\phi_G : H_1(\Sigma) \to \Gamma$. We thus define a Γ*-labelled surface* to be a pair (Σ, ϕ), where Σ is a surface and ϕ is a homomorphism from $H_1(\Sigma) \to \Gamma$. Let $\mathcal{S} := (\Sigma, \phi)$ be a Γ-labelled surface, and let C be a curve in Σ. Observe that C is naturally equipped with a group-value $\phi(C)$ from Γ. The easiest way to determine $\phi(C)$ is to work with a convenient representation of $\mathcal{S}$. A natural representation of $\mathcal{S}$ are the disks with Γ-labelled strips introduced in Section 7.7. That is, if we decompose $\mathcal{S}$ into a disk with Γ-strips, then we can determine $\phi(C)$ simply by counting how many times C passes through each strip.

Two Γ-labelled surfaces $\mathcal{S}_1 := (\Sigma_1, \phi_1)$ and $\mathcal{S}_2 := (\Sigma_2, \phi_2)$, are *isomorphic*, if there is a homeomorphism from $\xi : \Sigma_1 \to \Sigma_2$ such that $\phi_2(\xi(C)) = \phi_1(C)$ for all curves C in Σ_1.

Let G be a Γ-labelled graph 2-cell embedded in a surface Σ such that every face of G has group-value zero. Recall that G induces a natural homomorphism $\phi_G : H_1(\Sigma) \to \Gamma$. We say that (Σ, ϕ_G) is the Γ-labelled surface *induced by* G. Let $\Pi := \{(s_i, t_i, \gamma_i) : i \in [k]\}$ be a pattern in G. Since

$\mathcal{S} := (\Sigma, \phi_G)$ is a Γ-labelled surface, it is not nonsense to ask whether Π has a *topological realization* in $\mathcal{S}$. A topological realization of Π in $\mathcal{S}$ is a family $\{L_i : i \in [k]\}$ of disjoint arcs in Σ such that for each $i \in [k]$, the tail of L_i is s_i, the head of L_i is t_i, and $\phi(L_i) = \gamma_i$.

We are almost ready to present the theorem, but first we introduce some group-theoretic notation. Let G be a Γ-labelled graph, and let Γ' be a subgroup of Γ. We let Γ/Γ' denote the factor group of cosets of Γ'. For $\gamma \in \Gamma$, $\gamma + \Gamma'$ denotes the coset $\{\gamma + \gamma' : \gamma' \in \Gamma'\}$. We abuse notation and let G/Γ' denote the (Γ/Γ')-labelled graph obtained from G by reducing all the edge-labels (*modulo* Γ'). Similarly, if Π is a pattern in G, we let Π/Γ' be the pattern obtained from Π by replacing each (s, t, γ) in Π by $(s, t, \gamma + \Gamma')$. Thus, if Π is a pattern in G, then Π/Γ' is a pattern in G/Γ'.

Let G be a Γ-labelled graph 2-cell embedded in a surface Σ, and let F be the set of faces of G. We define Γ_F to be the subgroup of Γ generated by $\{\gamma(f) : f \in F\}$. We call Γ_F the *face subgroup* of Γ. We have previously shown that if the face subgroup of Γ is trivial, then G naturally endows Σ as a Γ-labelled surface. We now extend this definition by considering G/Γ_F. Evidently, every face of G/Γ_F is zero-valued in Γ/Γ_F. Thus, G/Γ_F induces a (Γ/Γ_F)-labelled surface $\mathcal{S}$. We again say that $\mathcal{S}$ is the (Γ/Γ_F)-labelled surface *induced by G*.

Finally, if A is a multiset of elements of Γ we say that A *strongly generates* Γ if every $\gamma \in \Gamma$ is the sum of the members of some sub-multiset of A. Here is the main result of this section.

Theorem 8.4.3. *For any $k \in \mathbb{N}$, any finite abelian group Γ, and any surface Σ there exist constants $R(k, \Gamma, \Sigma)$ and $W(k, \Gamma, \Sigma)$ with the following property. Let G be a Γ-labelled graph 2-cell embedded in $\widehat{\Sigma}$, and let Π be a k-pattern in G where $V(\Pi) \subseteq bd(\Sigma)$. Let Γ_F be the face subgroup of Γ, let $\Gamma' := \Gamma/\Gamma_F$, and let $\mathcal{S}$ be the Γ'-labelled surface induced by G.*

If G has a respectful tangle $\mathcal{T}$ in Σ of order at least $R(k, \Gamma, \Sigma)$ and for each $i \in [k]$ there exists a family $\mathcal{F}_i$ of faces of G satisfying

- *Π/Γ_F is topologically realizable in $\mathcal{S}$,*
- *$V(\Pi)$ is $\mathcal{T}$-independent,*
- *if δ_1 and δ_2 are distinct holes of Σ, then $d_\Sigma(\delta_1, \delta_2) \geq W(k, \Gamma, \Sigma)$,*
- *if δ is a hole and $f \in \bigcup_{i\in[k]} \mathcal{F}_i$, then $d_\Sigma(\delta, f) \geq W(k, \Gamma, \Sigma)$,*

- *if f_1 and f_2 are distinct faces of $\bigcup_{i\in[k]} \mathcal{F}_i$, then $d_\Sigma(f_1, f_2) \geq W(k, \Gamma, \Sigma)$, and*
- *for each $i \in [k]$, the multiset $\{\gamma(f) : f \in \mathcal{F}_i\}$ strongly generates Γ_F,*

then G has a Π-linkage.

Proof. Let μ be the function from Theorem 8.1.1. We will prove that

- $R(k, \Gamma, \Sigma) := 4\mu(k, \Gamma, \Sigma) + 2k|\Gamma|^2 + 6$
- $W(k, \Gamma, \Sigma) := 2\mu(k, \Gamma, \Sigma) + 5$

suffice.

Let $G, \Sigma, \Pi, k, \Gamma_F, \mathcal{F}_1, \ldots, \mathcal{F}_k, \mathcal{S}$, and $\mathcal{T}$ be given as above. For notational convenience we let $n := \mu(k, \Gamma, \Sigma)$. Let $\mathcal{F} := \bigcup_{i\in[k]} \mathcal{F}_i$. Let $\Pi := \{(s_i, t_i, \gamma_i) : i \in [k]\}$.

The first step is to create a new linkage problem from Π and the faces in $\mathcal{F}$. Observe that any multiset that strongly generates Γ contains a multiset of size at most $|\Gamma|^2$ that strongly generates Γ. Therefore, we may assume that $|\mathcal{F}| \leq k|\Gamma|^2$. Let X denote the set of vertices of G incident to a face in $\mathcal{F}$. By Theorem 7.5.3, there is a 2-cell subdrawing G_0 of $G\backslash X$ such that G_0 has a respectful tangle $\mathcal{T}_0$ of order at least $R(k, \Gamma, \Sigma) - 2k|\Gamma|^2$. For each $f \in \mathcal{F}$, let f^+ be the unique face in G_0 containing f. Let $\mathcal{R}_f$ be a set of two disjoint $V(f)$-$V(f^+)$ paths in G. We let u_f and v_f be the ends of $\mathcal{R}_f$ on $V(f)$ and x_f and y_f be the corresponding ends of $\mathcal{R}_f$ on $V(f^+)$. Let $\mathcal{J}_f$ be the set of two u_f-v_f paths through f. Perform shifts in G so that for all $f \in \mathcal{F}$,

- the two paths in $\mathcal{R}_f$ are both zero-labelled,
- the counter-clockwise path from u_f to v_f through f is zero-labelled.

Observe that this implies that the clockwise path from u_f to v_f in $\mathcal{J}_f$ has value $\gamma_G(f)$. We create a pattern $\Pi_{\mathcal{F}}$ in G as follows. Let $f_1, \ldots, f_m$ be the faces in $\mathcal{F}_1$. Place each of

$$(s_1, x_{f_1}, 0), (y_{f_1}, x_{f_2}, 0), \ldots, (y_{f_{m-1}}, x_{f_m}, 0), (y_{f_m}, t_1, \gamma_1)$$

into $\Pi_{\mathcal{F}}$. Repeat for $\mathcal{F}_2, \ldots, \mathcal{F}_k$.

Claim 1. *If G_0/Γ_F has a $(\Pi_{\mathcal{F}}/\Gamma_F)$-linkage, then G has a Π-linkage.*

Proof. Let $\mathcal{Q}$ be a $(\Pi_{\mathcal{F}}/\Gamma_F)$-linkage in G_0/Γ_F. We may also view $\mathcal{Q}$ as a set of paths in G_0, and hence also in G. We will show how to combine the paths in $\mathcal{Q}$, $\bigcup_{f\in\mathcal{F}}\mathcal{R}_f$, and $\bigcup_{f\in\mathcal{F}}\mathcal{J}_f$ into a Π-linkage of G. For each $i\in[k]$ let $\mathcal{Q}_i$ consist of the paths in $\mathcal{Q}$ incident with f^+, for some $f\in\mathcal{F}_i$. Since $\mathcal{Q}$ is a realization of $\Pi_{\mathcal{F}}/\Gamma_F$ in G_0/Γ_F, it follows that for each $i\in[k]$,

$$\sum_{Q\in\mathcal{Q}_i}\gamma_G(Q)=\gamma_i+\alpha_i,$$

for some $\alpha_i\in\Gamma_F$. However, by hypothesis the multiset $\{\gamma(f) : f\in\mathcal{F}_i\}$ strongly generates Γ_F for each $i\in[k]$. Therefore, by choosing an appropriate path from each $\mathcal{J}_f$ and then joining these paths to the paths in $\mathcal{Q}$, and $\bigcup_{f\in\mathcal{F}}\mathcal{R}_f$ in the obvious way, we obtain a realization of Π in G. $\square$

Let $\mathcal{H}$ denote the set of holes of Σ, and let $\mathcal{F}^+ := \{f^+ : f\in\mathcal{F}\}$. By Corollary 7.5.4, for each $h\in\mathcal{H}\cup\mathcal{F}^+$ there is a cycle C_h of G such that

- C_h bounds a disk Δ_h with $h\subseteq\Delta_h$, and
- $d_\Sigma(y,h)=n$, for all vertices y of C_h.

For each $h\in\mathcal{H}\cup\mathcal{F}^+$ let Z_h be the vertices of $\Pi_{\mathcal{F}}$ on h.

Claim 2. *If $h\in\mathcal{H}\cup\mathcal{F}^+$, then there is a family of $|Z_h|$ disjoint Z_h-C_h paths in G_0.*

SUBPROOF. If $h\in\mathcal{H}$, then Z_h is $\mathcal{T}$-independent, so the claim follows from Lemma 7.5.6. If $h\in\mathcal{F}$, then $|Z_h|=2$, so the claim follows since G is 2-cell embedded in Σ. $\square$

For each $h\in\mathcal{H}\cup\mathcal{F}^+$, let $\mathcal{P}_h$ be such a collection of paths. For each $z\in Z_h$ let $z'\in V(C_h)$ be the other endpoint of the path in $\mathcal{P}_h$ that contains z. Repeat this for all members of $\mathcal{H}\cup\mathcal{F}^+$, and define

$$Z' := \bigcup_{h\in\mathcal{H}\cup\mathcal{F}^+}\{z' : z\in Z_h\}.$$

Let $\Pi'_{\mathcal{F}}$ be the pattern obtained from $\Pi_{\mathcal{F}}$ by replacing each $(s,t,\gamma)\in\Pi_{\mathcal{F}}$ by (s',t',γ). Let $\mathcal{D}_{\mathcal{F}} := \{\Delta_f : f\in\mathcal{F}^+\}$ and $\mathcal{D}_{\mathcal{H}} := \{\Delta_\delta : \delta\in\mathcal{H}\}$. Since $d_\Sigma(h_1,h_2)\geq 2n+3$ for any two distinct members $h_1,h_2\in\mathcal{H}\cup\mathcal{F}^+$, the family $\mathcal{D} := \mathcal{D}_{\mathcal{H}}\cup\mathcal{D}_{\mathcal{F}}$ is a disjoint family of disks. Let Σ' be the surface obtained from $\widehat{\Sigma}$ be removing the interiors of those disks in $\mathcal{D}_{\mathcal{H}}$ (but not $\mathcal{D}_{\mathcal{F}}$). Define $G' := G_0 \cap \Sigma'$. We naturally regard Σ' as a (Γ/Γ_F)-labelled surface $\mathcal{S}'$ that is isomorphic to $\mathcal{S}$.

Claim 3. *$\Pi'_{\mathcal{F}}/\Gamma_F$ is topologically realizable in $\mathcal{S}'$.*

SUBPROOF. Decompose $\mathcal{S}'$ into a disk with Γ-strips such that all disks in $\mathcal{D}_{\mathcal{F}}$ are contained in the disk $\Delta(\mathcal{S}')$. Since by hypothesis Π/Γ_F is topologically realizable in $\mathcal{S}$, it follows readily that $\Pi'_{\mathcal{F}}/\Gamma_F$ is topologically realizable in $\mathcal{S}'$. $\square$

The remainder of the argument conforms to the proof of Theorem 8.4.2. Let $\mathcal{L}$ be a topological realization of $\Pi'_{\mathcal{F}}/\Gamma_F$ in $\mathcal{S}'$ and let $H := G_0/\Gamma_F$. We may assume that $\mathcal{L}$ intersects H finitely often. We will define a (Γ/Γ_F)-labelled graph $H + \mathcal{L}$ from H and $\mathcal{L}$ as follows. We define $V(H + \mathcal{L})$ to be $V(G_0) \cup (\mathcal{L} \cap G_0)$ and $E(H + \mathcal{L})$ to be the set of arcs in $G_0 \cup \mathcal{L}$ with both ends in $V(H + \mathcal{L})$. Each edge of $H + \mathcal{L}$ inherits an orientation from H or from $\mathcal{L}$. Let

$$E_{\mathcal{L}} := \{e \in E(H + \mathcal{L}) : \ e \subseteq L \text{ for some } L \in \mathcal{L}\}.$$

It remains to define the edge-labels of $H + \mathcal{L}$. For each $e \in E(H + \mathcal{L})$ we choose an edge-label $\gamma_{H+\mathcal{L}}(e)$ such that

- Every face of $H + \mathcal{L}$ has group-value zero (in Γ/Γ_F).
- The edges in $E_{\mathcal{L}}$ are the edges of a $(\Pi'_{\mathcal{F}}/\Gamma_F)$-linkage in the obvious way.
- $(H + \mathcal{L})\backslash E_{\mathcal{L}}$ is a topological minor of H in the obvious way.

The easiest way to do this is to take a (Γ/Γ_F)-labelled triangulation T of Σ' such that

- T and H both induce the same (Γ/Γ_F)-labelled surface, and
- H is a topological minor of T.

If T is sufficiently fine, it follows that we can actually choose $\mathcal{L}$ so that it only passes through edges of T. We then recover the edge labels of $H + \mathcal{L}$ from the edge labels in T.

We now regard $\Pi_{\mathcal{F}}/\Gamma_F$ as a pattern in $H + \mathcal{L}$ and show that every edge in $E_{\mathcal{L}}$ is redundant for $\Pi_{\mathcal{F}}/\Gamma_F$ in $H + \mathcal{L}$. Let $e \in E_{\mathcal{L}}$. Since $\mathcal{T}$ is a respectful tangle of order at least $2n+1$, by Corollary 7.5.5 there are vertex disjoint cycles $C_1, \ldots, C_n$ in $H + \mathcal{L}$ bounding disks $\Delta_1, \ldots, \Delta_n$ in Σ such that $e \subset \Delta_1 \subset \cdots \subset \Delta_n$. Moreover, since $\mathcal{L}$ is disjoint from the interior of each of the disks in $\mathcal{D}_{\mathcal{H}} \cup \mathcal{D}_{\mathcal{F}}$, we may assume that $V(\Pi_{\mathcal{F}}/\Gamma_F)$ is disjoint from $int(\Delta_n)$. That is, e is an n-protected edge in Σ with respect to $\Pi_{\mathcal{F}}/\Gamma_F$. By Theorem 8.1.1, e is redundant for $\Pi_{\mathcal{F}}/\Gamma_F$.

We proceed sequentially through $E_{\mathcal{L}}$ and conclude that $(H + \mathcal{L})\backslash E_{\mathcal{L}}$ has a ($\Pi_{\mathcal{F}}/\Gamma_F$)-linkage if and only if $H + \mathcal{L}$ has a ($\Pi_{\mathcal{F}}/\Gamma_F$)-linkage. But $H + \mathcal{L}$ clearly does have a ($\Pi_{\mathcal{F}}/\Gamma_F$)-linkage by its construction. Therefore, $(H + \mathcal{L})\backslash E_{\mathcal{L}}$ has a $\Pi_{\mathcal{F}}/\Gamma_F$-linkage. It follows that H must also have a $\Pi_{\mathcal{F}}/\Gamma_F$-linkage. By Claim 1, G has a Π-linkage, as required. $\square$

Chapter 9

Taming a Vortex

We previously encountered vortices when discussing the Graph Minors Structure Theorem in Chapter 5. In this chapter, we prove some important results concerning the structure of vortices. First we prove that if G is embedded in a surface Σ and H is a subgraph of G which is "clustered" in Σ, then we may view H as a vortex. Next we show that under certain conditions, we can slightly enlarge a vortex so that it has a special type of vortex decomposition that is "linked". This is quite crucial, as our final theorem asserts that it is always possible to reroute a linkage so that it only passes through a "linked" vortex a few times.

9.1 Vortices and Distance

We recall the definitions from Section 5.2. A *society* is a finite set of points S that are cyclically ordered. An *interval of* S is a (non-empty) set of consecutive vertices of S in the ordering. A *halving* of S is a partition of S into two intervals. For $u, v \in S$, we let $S(u, v)$ denote those vertices that occur after u but before v in S. We define $S[u, v] := S(u, v) \cup \{u, v\}$. So, $\{S(u, v), S[v, u]\}$ is a halving of S. If G is a graph and $S \subseteq V(G)$ is a society, we call the pair (G, S) a *vortex*. A vortex (G, S) has *adhesion at most* n if for any halving of S, there do not exist n vertex disjoint paths in G between the two halves.

Theorem 9.1.1. *Let $n \in \mathbb{N}$, let G be a graph embedded in a surface Σ, let d_Σ be the surface metric, and let $x \in V(G)$. If there is a cycle C of G bounding a disk Δ*

in Σ such that $d_\Sigma(x, y) = n$ for all $y \in V(C)$, then $(G \cap \Delta, V(C))$ is a vortex of adhesion at most $2n + 4$.

Proof. Suppose not. Let $\{X, Y\}$ be a halving of $V(C)$ such that there is a family $\mathcal{P}$ of $2n + 4$ disjoint X-Y paths of G contained in Δ. Let H be the subgraph of G induced by $E(C) \cup E(\mathcal{P})$. We regard H as embedded in Δ, and as such H has at least $2n+3$ faces in Δ. Note that the dual graph H^* of H is a path. We label the faces of H as $F_1, \ldots, F_m$ according to their order in H^*. Let i be the minimum index in $[m]$ such that $x \in \overline{F_i}$.

If $i \leq n + 1$, then we let z be any vertex in $V(C) \cap V(F_m)$ and D be any dual curve in Σ connecting x to z. If D is contained in Δ, then clearly D has length at least $n+2$. On the other hand, if D is not contained in Δ, then D must use a vertex y of C before reaching z. Since, $d_\Sigma(x, y) = n$, it follows that D has length at least $n + 2$ in this case as well. Thus, $d_\Sigma(x, z) \geq n + 1$, which is a contradiction.

The remaining possibility is if $i \geq n+2$, and $x \notin V(F_{n+1})$. The previous argument shows that $d_\Sigma(x, z) \geq n+1$ where z is any vertex in $V(C) \cap V(F_1)$. This completes the proof. □

9.2 Linked Vortex Decompositions

Let G be a graph and let L be a linearly ordered subset of $V(G)$. Recall that a vortex decomposition of (G, L) is a collection $\{G_v : v \in L\}$ of subgraphs of G such that for all $x, y \in L$, with $x \leq y$

(V1) $E(G_x \cap G_y) = \emptyset$, and $\bigcup_{v \in L} G_v = G$.

(V2) $G_x \cap G_y \subseteq \bigcap_{z \in L[x,y]} G_z$.

(V3) If $x \in V(G_y) \cap L$, then $y = x$ or y is the successor of x in L.

A vortex decomposition is *linked* if it additionally satisfies

(V4) For any three consecutive vertices x, y, z of L, there is a collection of disjoint paths in G_y linking $V(G_x \cap G_y)$ to $V(G_y \cap G_z)$.

The *linked depth* of a vortex is the minimum depth taken over all linked vortex decompositions.

We end this section by making some observations about linked vortex decompositions that will be needed later. Let $\{G_v : v \in L\}$ be a linked vortex decomposition of (G, L) of depth d. The first observation is that **(V4)** implies that $|V(G_x \cap G_y)| = d$ for all consecutive vertices $x, y \in L$. Next we observe that if $[x, y]$ is an interval of L, then it is clear that $(\bigcup_{v\in[x,y]} G_v, \bigcup_{v\notin[x,y]} G_v)$ is a separation of G of order at most $2d$. We denote this separation as $G[x, y]$ and call it the separation of G *induced by* $[x, y]$. It is clear that nested intervals of L induce nested separations of G.

Lemma 9.2.1. *Let $\{G_v : v \in L\}$ be a linked vortex decomposition of (G, L) and let $[x_1, y_1] \subseteq \ldots \subseteq [x_m, y_m]$ be intervals of L. Then $G[x_1, y_1], \ldots, G[x_m, y_m]$ is a nested sequence of separations of G.*

We now prove a strengthened form of **(V4)**.

Lemma 9.2.2. *Let $\{G_v : v \in L\}$ be a linked vortex decomposition of (G, L) of depth d. If $v_1, \ldots, v_n$ are consecutive vertices of L, then there are d vertex disjoint paths in $\bigcup_{i=2}^{n-1} G_{v_i}$ between $V(G_{v_1} \cap G_{v_2})$ and $V(G_{v_{n-1}} \cap G_{v_n})$.*

Proof. We proceed by induction on n. The base case $n = 3$ is handled by **(V4)**. Let $v_1, \ldots, v_n$ be consecutive vertices of L. By induction there is a collection $\mathcal{P}$ of d vertex disjoint paths in $\bigcup_{i=2}^{n-2} G_{v_i}$ between $V(G_{v_1} \cap G_{v_2})$ and $V(G_{v_{n-2}} \cap G_{v_{n-1}})$. By **(V4)** there is a collection $\mathcal{Q}$ of d vertex disjoint paths in $G_{v_{n-1}}$ between $V(G_{v_{n-2}} \cap G_{v_{n-1}})$ and $V(G_{v_{n-1}} \cap G_{v_n})$. By **(V2)**, we have that $V(\mathcal{Q}) \cap V(\mathcal{P}) = V(G_{v_{n-2}} \cap G_{v_{n-1}})$. Therefore, by combining the paths in $\mathcal{P}$ and $\mathcal{Q}$ appropriately, we get the desired set of paths. $\square$

A sequence $(A_1, B_1), \ldots, (A_m, B_m)$ of separations of G is *nested*, if $A_i \subseteq A_{i+1}$ and $B_{i+1} \subseteq B_i$ for all $i \in [m-1]$.

Lemma 9.2.3. *Let $\{G_v : v \in L\}$ be a linked vortex decomposition of (G, L) of depth d. Let $[x_1, y_1] \subseteq \ldots \subseteq [x_m, y_m]$ be intervals of L such that the separations $G[x_1, y_1], \ldots, G[x_m, y_m]$ all have the same order, say n. Let U be the vertex boundary of $G[x_1, y_1]$ and V be the vertex boundary of $G[x_m, y_m]$. Then there is a family of n disjoint U-V paths in $\bigcup_{v\in[x_m,x_1]\cup[y_1,y_m]} G_v$.*

Proof. By Lemma 9.2.2, there is a family $\mathcal{P}$ of d paths such that for each $P \in \mathcal{P}$ and each $v \in L$, G_v contains a vertex of P. Therefore, letting $G' := \bigcup_{v \in [x_m, x_1] \cup [y_1, y_m]} G_v$, we see that $\mathcal{P} \cap G'$ contains the required family of n disjoint U-V paths. $\square$

9.3 Linking a Vortex

The main result of this section is subject to the following assumptions. Let G be a graph and let (G_0, G_1) be a separation of G such that:

(A1) G_0 is 2-cell embedded in a surface Σ.

(A2) $V(G_0 \cap G_1)$ are the vertices of a face F of G_0.

(A3) If $V(F)$ is cyclically ordered via the boundary of F then $(G_1, V(F))$ is a vortex of adhesion at most n.

(A4) G_0 has a respectful tangle $\mathcal{T}$ of order at least $2n + 1$.

Under these conditions, it is possible to enlarge F to a disk Δ, so that the portion of G inside Δ has a *linked* vortex decomposition of depth at most n.

Theorem 9.3.1. *There is a cycle C_0 in G_0 bounding a disk Δ containing F such that the vortex $(G_1 \cup (\Delta \cap G_0), V(C_0))$ has a linked vortex decomposition of depth at most n.*

Proof. For each $i \in [n]$, let S_i be the set of vertices of G_0 at distance exactly $i - 1$ from F (with respect to d_Σ). By Corollary 7.5.5, for each $i \in [n]$, there exists a cycle C_i of G_0 which only passes through vertices of S_i. Furthermore, each C_i bounds a disk Δ_i in Σ such that $\Delta_1 \subset \cdots \subset \Delta_n$. By **(A1)**, we may assume $\Delta_1 = \overline{F}$.

Let H' be the subgraph of G induced by $V(G_1) \cup V(C_1) \cup \cdots \cup V(C_n)$. Observe that $C_1 \cup C_n$ bounds a cylinder Σ' in Σ. Define $J' := H' \cap \Sigma'$. By construction, there exists a dual curve R of length n in Σ' that intersects each C_i exactly once. Label $V(R) = \{r_1, \ldots, r_n\}$, where $r_i \in V(C_i)$ for each i. Cut Σ' open along R, splitting each r_i into two copies u_i and v_i. Let H

and J be the graphs obtained from H' and J' after splitting $V(R)$, and let $U := \{u_i : i \in [n]\}$, $V := \{v_i : i \in [n]\}$.

Every u_n-v_n path P in J corresponds to a cycle $C(P)$ in J' which bounds a disk $\Delta(P)$ in Σ. We define $H'_P := (G_0 \cap \Delta(P)) \cup G_1$, and H_P to be the graph obtained from H_P by cutting Σ' open along the dual curve R. Note that H_P is a subgraph of H. We now choose a u_n-v_n path P in J such that

(1) $\kappa_{H_P}(U, V) = n$,

(2) subject to (1), $\Delta(P)$ is minimal (with respect to inclusion).

Let $\mathcal{Q} := \{Q_i : i \in [n]\}$ be a collection of n disjoint U-V paths in H_P, labelled according to their endpoints in $U = \{u_i : i \in [n]\}$. By planarity, if the last path Q_n uses an edge of G_1, then all paths in $\mathcal{Q}$ must also use an edge of G_1. Therefore, Q_n does not use an edge of G_1, otherwise the vortex (G_1, S) would have adhesion more than n. It now follows that Q_n must in fact connect u_n to v_n. Observe that by choice of P, we have $Q_n = P$. Let $Q_n = b_1 \dots b_m$. By choice of P, there is a family $\mathcal{Y} := \{Y_1, \dots, Y_m\}$ of (vertex) separations such that for each $i \in [m]$,

(S1) Y_i separates U from V in H_P,

(S2) $|Y_i| = n$, and

(S3) $b_i \in Y_i$.

If $\mathcal{Y}$ is such a family and $Y_i, Y_j \in \mathcal{Y}$, we say that Y_i and Y_j *cross* if $i < j$ and there exist $a \in Y_i$ and $b \in Y_j$ such that b occurs before a on some path of $\mathcal{Q}$.

Claim. *There exists such a family $\mathcal{Y}'$ such that no two members of $\mathcal{Y}'$ cross.*

SUBPROOF. Let $\mathcal{Y} := \{Y_1, \dots, Y_m\}$ satisfy **(S1)**, **(S2)**, and **(S3)**. We will uncross the sets in $\mathcal{Y}$ one at a time. Let $Y_i, Y_j \in \mathcal{Y}$ with $i < j$. Observe that there are separations (A, B) and (C, D) of H_P such that

- $V(A \cap B) = Y_i$, and $V(C \cap D) = Y_j$.
- $U \subseteq V(A \cap C)$, and $V \subseteq V(B \cap D)$.

Since (A, B) and (C, D) are both minimum order separations separating U and V, so are $(A \cap C, B \cup D)$ and $(A \cup C, B \cap D)$, by submodularity. We thus set $Y_i := bd(A \cap C, B \cup D)$ and $Y_j := bd(A \cup C, B \cap D)$ and refer to this operation as *uncrossing* Y_i and Y_j.

Let $I = \{(i, j) : i, j \in [m], i < j\}$. For $(i, j) \in I$ and $(i, j) \neq (m-1, m)$ we let $(i, j)^+$ be the successor of (i, j) in I under the usual lexicographic ordering. It is now easy to state how to construct the required family $\mathcal{Y}'$. Starting with $(i, j) = (1, 2)$, uncross Y_i and Y_j. Set $(i, j) = (i, j)^+$ and recurse until $(i, j) = (m-1, m)$. $\square$

So, let $\mathcal{Y}' = \{Y_1, \ldots, Y_m\}$ be such a family. Therefore, for each $i \in [m]$ there exists a separation (A_i, B_i) of H_P of order n such that

- $bd(A_i, B_i) = Y_i$,
- $U \subseteq A_i$ and $V \subseteq B_i$,
- $A_i \subseteq A_{i+1}$, and $B_{i+1} \subseteq B_i$, for each $i \in [m-1]$.

For each $i \in [m-1]$ we define $G'_i := A_{i+1} \cap B_i$.

Now note that the last path $Q_n \in \mathcal{Q}$ actually corresponds to a cycle C_0 in G_0 since its endpoints become identified. Since C_0 does not use an edge of G_1, it follows that C_0 bounds a disk Δ in Σ containing F. Let $G' := G_1 \cup (\Delta \cap G_0)$. By construction, $G'_1 \cup \cdots \cup G'_{m-1}$ is a linked vortex decomposition of depth n of the vortex $(G', V(C_0))$, as required. $\square$

9.4 Avoiding Vortices

In this section G is a Γ-labelled graph that is l-near embedded in a surface Σ, where the near embedding has no apex vertices. We will analyze the structure of linkages in G with respect to such a near embedding. Let Π be a pattern in G. Our main result is that if $\mathcal{P}$ is a set of paths realizing Π, then it is always possible to reroute the paths in $\mathcal{P}$ so that they still realize Π, but only pass through the vortices a few times.

Let us be precise. Let $G_0 \cup G_1 \cup \cdots \cup G_l$ be a l-near embedding of G in Σ with no apex vertices, G_0 as the embedded subgraph and $(G_1, L_1), \ldots, (G_l, L_l)$ the vortices of the near embedding.

A path P in G is a Σ-*jump* if for some $i \in [l]$

- $P \subseteq G_0$,
- both ends of P are on $bd(\Sigma)$, and
- no internal vertex of P is on $bd(\Sigma)$.

The main result of this section is the following. Our proof is based on a proof of an analogous result for graphs in Graph Minors XXI [43, Theorem 6.3].

Theorem 9.4.1. *For every surface Σ, every finite abelian group Γ and every $k, l \in \mathbb{N}$ there exists $j := j(k, l, \Gamma, \Sigma) \in \mathbb{N}$ such that, if $G_0 \cup G_1 \cup \cdots \cup G_l$ is an l-near embedding of a Γ-labelled graph G in Σ with no apex vertices, G_0 as the embedded subgraph and $(G_1, L_1), \ldots, (G_l, L_l)$ as the vortices and Π is a realizable k-pattern in G, then there is a set of paths $\mathcal{P}$ in G realizing Π with at most j Σ-jumps.*

We will require the following two obvious lemmas.

Lemma 9.4.2. *Let $(\gamma_1, \ldots, \gamma_m)$ be a sequence of elements from a finite abelian group Γ. If $m > |\Gamma|$, then there exist indices $i < j$ such that $\sum_{n=i+1}^{j} \gamma_n = 0$.*

Proof. By the pigeonhole principle, there exist indices $i < j$ such that $\sum_{n=1}^{i} \gamma_n = \sum_{n=1}^{j} \gamma_n$. But then $\sum_{n=i+1}^{j} \gamma_n = 0$, as required. □

Lemma 9.4.3. *Let G be a Γ-labelled graph, Π be a pattern in G, and (A, B) be a separation of G. If $\mathcal{P}$ is a Π-linkage in G, then $\mathcal{P} \cap B$ is a Π'-linkage in B, where $V(\Pi') \subseteq (V(\Pi) \backslash V(A)) \cup V(A \cap B)$.*

The key idea of the proof is the next lemma, which allows us to control how a linkage passes through a sequence of nested separations.

Lemma 9.4.4. *For all $k, n \in \mathbb{N}$ and all finite abelian groups Γ, there exists $m := m(k, n, \Gamma) \in \mathbb{N}$ with the following property. Let G be a Γ-labelled graph and let Π be a realizable k-pattern in G. Let $(A_1, B_1), \ldots, (A_m, B_m)$ be a nested sequence of separations of G, each of order n. Finally, let $\mathcal{L}$ be a linkage in G satisfying*

- *$\{tail(L) : L \in \mathcal{L}\} = V(A_1 \cap B_1)$, and*
- *$\{head(L) : L \in \mathcal{L}\} = V(A_m \cap B_m)$.*

Then there is a realization $\mathcal{P}$ of Π, and indices $s < t$ such that $\mathcal{P} \cap (B_s \cap A_t) \subseteq \mathcal{L} \cap (B_s \cap A_t)$.

Proof. We show that $m(k, n, \Gamma) := |\Gamma|(2^{|\Gamma|(2k+n)^2}) + 1$ suffices. Fix k, n and Γ and let

$$G, \Pi, (A_1, B_1), \ldots, (A_m, B_m), \mathcal{L}$$

be a counterexample with $|E(G)|$ minimal. By shifting, we may assume that all edges in $E(\mathcal{L})$ are zero-labelled. Let $\mathcal{L} := \{L_1, \ldots, L_n\}$ and for each $i \in [m]$, label the vertices in $V(A_i \cap B_i)$ as $\{l_i^1, \ldots, l_i^n\}$ according to the (unique) L_j which passes through it. We remark that it is possible that say $l_i^1 = l_j^1$, for distinct i and j. Let $X := V(\Pi)$ and consider the set $X \cup [n]$. For each finite abelian group Γ, there are at most $2^{|\Gamma|(2k+n)^2}$ patterns Π', such that $V(\Pi') \subseteq X \cup [n]$. Let $\mathcal{Q}$ be a realization of Π. For any $i \in [m]$, let Π_i be the pattern of $\mathcal{Q} \cap B_i$ in B_i. By Lemma 9.4.3, $V(\Pi_i) \subseteq X \cup \{l_1, \ldots, l_n\}$. If we identify $\{l_1, \ldots, l_n\}$ with $[n]$, then we may regard each Π_i as a pattern with $V(\Pi_i) \subseteq X \cup [n]$. Let $g := |\Gamma| + 1$. Since $m = |\Gamma|(2^{|\Gamma|(2k+n)^2}) + 1$, there are indices $i_1 < i_2 < \cdots < i_g$ such that Π_{i_j} are the same for all $j \in [g]$. In particular, this implies that $X \cap A_{i_1} = X \cap A_{i_g}$ and $X \cap B_{i_1} = X \cap B_{i_g}$. We may assume that $A_{i_g} \cap B_{i_{g-1}}$ contains an edge in $E(\mathcal{L})$. Otherwise, we are done since we can take $\mathcal{P} := \mathcal{Q}$, $s = i_{g-1}$, and $t = i_g$.

There are two cases to consider. The first is if $V(\Pi_{i_g})$ is disjoint from $[n]$. But then, $E(\mathcal{Q}) \cap E(A_{i_g} \cap B_{i_{g-1}})$ is empty, a contradiction.

By symmetry, the remaining case is if $(x, i, \gamma) \in V(\Pi_{i_g})$, for some $x \in X$, $i \in [n]$, and $\gamma \in \Gamma$. Since $\Pi_{i_g} = \Pi_{i_{g-1}} = \cdots = \Pi_{i_1}$, it follows that there is a path $Q \in \mathcal{Q}$ that passes through x and each of l_j^i for all $j \in \{i_1, \ldots, i_g\}$. By switching i if necessary, we may assume that Q contains an edge of $A_{i_g} \cap B_{i_{g-1}}$. Now, for each $j \in [g]$ we let γ_j be the group-value of the subpath of Q from x to l_j^i. By Lemma 9.4.2, there exist $q < r$ such that $\sum_{j=q+1}^{r} \gamma_j = 0$. Let $Q' := Q \cap (A_{i_r} \cap B_{i_q})$ and $L' := L_i \cap (A_{i_r} \cap B_{i_q})$. We replace Q in $\mathcal{Q}$ by $(Q \backslash Q') \cup L'$. Letting e be an edge of L', we have that G/e is a smaller counterexample, a contradiction. □

We now proceed to prove Theorem 9.4.1. We illustrate the main idea by first considering the case when Σ is the disk.

Theorem 9.4.5. *For any $k, d \in \mathbb{N}$, and any finite abelian group Γ, there exists $j := j(k, d, \Gamma) \in \mathbb{N}$ with the following property. Let G be a Γ-labelled graph, Π be a realizable k-linkage in G and (G_0, G_1) be a separation of G such that*

- *G_0 is embedded in a disk Δ with $V(\Pi) \subseteq bd(\Delta)$, and*
- *$(G_1, V(G_1) \cap bd(\Delta))$ has a linked vortex decomposition of depth at most d.*

Then there is a realization of Π in G with at most j Δ-jumps.

Proof. Fix k, d, and Γ, and let $f := m(k, 2d, \Gamma)2d$, where m is the function from Lemma 9.4.4. We will show that $j := f^f$ suffices. Let $G := G_0 \cup G_1$, and Π be a counterexample with $|E(G)|$ minimal. Among all realizations of Π in G choose $\mathcal{P}$ with the minimum number of Δ-jumps. By choice of G, we note that $\mathcal{P}$ contains more than j Δ-jumps. Let $\delta := bd(\Delta)$, $L := V(G) \cap \delta$, and let $\{G_v : v \in L\}$ be a linked vortex decomposition of (G_1, L) of depth d. Note that we regard L as a sub-order of δ, see Remark 5.2.2.

Claim. $V(G_0) = V(\Pi)$.

SUBPROOF. Deleting any vertices of G_0 which do not appear in $V(\mathcal{P})$ would yield a smaller counterexample. Now, let e be an edge of G_0 with ends x and y such that x is not δ. By shifting at x so that $\gamma_G(e) = 0$, and then contracting e onto y, we get that G/e is a smaller counterexample. □

Therefore the Δ-jumps are simply the edges of G_0 and it suffices to prove that $|E(G_0)| \leq j$. In order to apply Lemma 9.4.4 we require a suitable collection of nested separations, which will be provided via certain dual curves. We call a dual curve b in Δ a *bite* if b only meets G_0 at its ends, which are on L. Let b be a bite with ends $x < y$ on L. We define $\delta[x, y]$ to be the clockwise arc of δ from x to y. We let $\Delta(b)$ be the disk in Δ bounded by $b \cup \delta[x, y]$. Now it is easy to see how to construct a separation of G from b. Namely, we define

$$A(b) := (G_0 \cap \Delta(b)) \cup \bigcup_{z \in L[x,y]} G_z, \text{ and}$$
$$B(b) := (G_0 \cap \overline{\Delta \backslash \Delta(b)}) \cup \bigcup_{z \notin L[x,y]} G_z.$$

It is clear that $(A(b), B(b))$ is a separation of G of order at most $2d$. Moreover,

Claim. *If $b_1, \dots, b_n$ is a sequence of bites such that $\Delta(b_1) \subseteq \dots \subseteq \Delta(b_n)$, then the sequence of separations $(A(b_1), B(b_1)), \dots, (A(b_n), B(b_n))$ is nested.*

SUBPROOF. Immediate from Lemma 9.2.1. □

Recall that we must prove that $|E(G_0)| \leq j$. We do this by bounding the number of edges in G_0^*, the dual graph of G_0 in Δ. Recall that the vertices of G_0^* are the faces of G_0 in Δ, and two faces $F_1, F_2 \in V(G_0^*)$ are adjacent if and only if $bd(F_1) \cap bd(F_2)$ contains an edge of G_0. Evidently, $|E(G_0^*)| = |E(G_0)|$, and G_0^* is a tree. We will show that G_0^* has maximum degree f and that every path in G_0^* has length at most f. It will thus follow that $|E(G_0^*)| \leq f^f$, as required.

If the dual graph G_0^* contains a vertex of degree greater than f or a path of length greater than f, then there is a sequence of bites $b_1, \dots, b_f$ in Δ and distinct edges $e_1, \dots, e_f$ of G_0 such that

- $\Delta(b_1) \subseteq \dots \subseteq \Delta(b_f)$, and
- $e_i \in \Delta(b_i) \backslash \Delta(b_{i-1})$ for each $i \in [f]$. (define $\Delta(b_0) = \emptyset$)

Since $f = m(k, 2d, \Gamma)2d$, by re-indexing we may assume that there is a subsequence $b_1, \dots, b_m$ such that the separations $(A(b_i), B(b_i))$ all have the same order, say $n \leq 2d$.

By Lemma 9.2.3 there is a collection $\mathcal{L}$ of n disjoint paths in G_1 such that

- $\{tail_G(L) : L \in \mathcal{L}\}$ is the vertex boundary of $(A(b_1), B(b_1))$, and
- $\{head_G(L) : L \in \mathcal{L}\}$ is the vertex boundary of $(A(b_m), B(b_m))$.

We are now in prime position to apply Lemma 9.4.4. We conclude that there are indices $s < t$ in $[m]$ and another realization $\mathcal{P}'$ of Π such that $\mathcal{P}' \cap (A(b_t) \cap B(b_s))$ is a subset of $\mathcal{L} \cap (A(b_t) \cap B(b_s))$. Since $e_t \in E(A(b_t) \cap B(b_s))$, but clearly $e_t \notin E(\mathcal{P}')$, it follows that $G \backslash e_t$ is a smaller counterexample, a contradiction. □

We now prove the general case.

Theorem 9.4.6. *For every surface Σ, every finite abelian group Γ and every $k, l \in \mathbb{N}$ there exists $j := j(k, l, \Gamma, \Sigma) \in \mathbb{N}$ such that, if $G_0 \cup G_1 \cup \cdots \cup G_l$ is a l-near embedding of a Γ-labelled graph G in Σ with no apex vertices, G_0 as the embedded subgraph and $(G_1, L_1), \ldots, (G_l, L_l)$ as the vortices and Π is a realizable k-pattern in G with $V(\Pi) \subseteq bd(\Sigma)$, then there is a realization of Π in G with at most j Σ-jumps.*

Proof. We will prove that

$$j(k, l, \Gamma, \Sigma) := (2lm(k, 4l, \Gamma))^{2lm(k,4l,\Gamma)+5} h(\Sigma)^2 \max\{\rho_1(\Sigma)\rho_2(\Sigma)\},$$

suffices, where $m(k, n, \Gamma)$ is the function from Lemma 9.4.4, $h(\Sigma)$ is the number of holes of Σ, $\rho_1(\Sigma)$ is the function from Lemma 7.1.3, and $\rho_2(\Sigma)$ is the function from Lemma 7.1.4.

Fix Σ, Γ, k, l and let $G := G_0 \cup G_1 \cup \cdots \cup G_l$ be a counterexample, with $|E(G)|$ minimal. For notational convenience we let $M = m(k, 4l, \Gamma)$. For each $i \in [l]$, let $\{G_v : v \in L_i\}$ be a linked vortex-decomposition of G_i of depth at most l, attached to the hole δ_i of Σ. Finally, let $\mathcal{P}$ be a realization of Π in G with the minimum number of Σ-jumps. As in the proof of Theorem 9.4.5 we have

Claim. $V(G_0) = V(G) \cap bd(\Sigma)$.

Therefore, the Σ-jumps of $\mathcal{P}$ are simply the edges of G_0. So, it suffices to prove that $|E(G_0)| \leq j$. We partition $E(G_0)$ as $E_1 \cup E_2 \cup E_3$, where E_1 consists of the edges of G_0 which connect up two different holes, E_2 consists of the non-contractible edges of G_0 which connect up a common hole, and E_3 consists of the contractible edges of G_0 which connect up a common hole. We will show that none of E_1 nor E_2 nor E_3 can be very large.

Let us start with E_1. For any two holes δ_i and δ_j of Σ, we let $E(\delta_i, \delta_j)$ be the edges of G_0 with one end on δ_i and the other on δ_j. By definition of ρ_2, there are at most $\rho_2(\Sigma)$ homotopy classes (with respect to δ_1 and δ_2) of edges in $E(\delta_1, \delta_2)$.

Therefore, if there are at least $8lM\rho_2(\Sigma)\binom{h(\Sigma)}{2}$ edges of G_0 which connect different holes, then for some two holes, say δ_1 and δ_2, there are $8lM\rho_2(\Sigma)$ edges in $E(\delta_1, \delta_2)$. Thus, there are $8lM$ pairwise homotopic edges $e_1, \ldots, e_{8lM}$ in $E(\delta_1, \delta_2)$. For each $i \in [8lM]$ let x_i be the end of e_i on δ_1 and y_i be the end of e_i on δ_2. By choosing an appropriate half of the edges and re-indexing, we may assume that either

- $x_1 < \cdots < x_{4lM}$ in L_1,
- $y_{4lM} < \cdots < y_1$ in L_2, and
- $e_1 \cup e_{4lM} \cup \delta_1[x_1, x_{4lM}] \cup \delta_2[y_{4lM}, y_1]$ bounds a disk in Σ,

or

- $x_1 < \cdots < x_{4lM}$ in L_1,
- $y_1 < \cdots < y_{4lM}$ in L_2, and
- $e_1 \cup e_{4lM} \cup \delta_1[x_1, x_{4lM}] \cup \delta_2[y_1, y_{4lM}]$ bounds a disk in Σ.

In either case, just as in the proof of Lemma 9.4.5, these $4lM$ edges induce a sequence $(A_1, B_1), \ldots, (A_{4lM}, B_{4lM})$ of nested separations of G such that

- the order of (A_i, B_i) is at most $4l$ for each $i \in [4lM]$, and
- $e_i \in E(A_i)\backslash E(A_{i-1})$. (define $E(A_0) = \emptyset$)

By taking an appropriate subsequence, we may assume that $(A_1, B_1), \ldots, (A_M, B_M)$ all have the same order, say $n \leq 4l$. Now since both (G_1, L_1) and (G_2, L_2) have linked vortex decompositions of width at most l, it follows that there is a linkage $\mathcal{L}$ of size n in G such that $\{tail_G(L) : L \in \mathcal{L}\} = bd(A_1, B_1)$ and $\{head_G(L) : L \in \mathcal{L}\} = bd(A_M, B_M)$. As in the proof of Lemma 9.4.5 we conclude that there is another realization $\mathcal{P}'$ of Π in G such that $e_t \notin E(\mathcal{P}')$ for some $t \in [M]$. Thus, $G\backslash e_t$ is a smaller counterexample, a contradiction. Therefore, $|E_1| < 8lM\binom{h(\Sigma)}{2}\rho_2(\Sigma)$, else we are done.

We now deal with E_2. If there are at least $4lM\rho_1(\Sigma)h(\Sigma)$ edges in E_2, then for some hole δ_1 of Σ, at least $4lM\rho_1(\Sigma)$ edges in E_2 are δ_1-edges. By definition of ρ_1, we conclude that at least $4lM$ of these δ_1-edges are pairwise homotopic. As in the previous case, we find M pairwise homotopic δ_1-edges $e_1, \ldots, e_M$ which induce a nested sequence of separations $(A_1, B_1), \ldots, (A_M, B_M)$ of G such that for each $i \in [M]$

- each (A_i, B_i) has the same order, say $n \leq 4l$, and
- $e_i \in E(A_i)\backslash E(A_{i-1})$. (define $E(A_0) = \emptyset$)

Since (G_1, L_1) has a linked vortex decomposition of width at most l, we again conclude that there is a linkage $\mathcal{L}$ of size n in G such that $\{tail_G(L) : L \in \mathcal{L}\} = bd(A_1, B_1)$ and $\{head_G(L) : L \in \mathcal{L}\} = bd(A_M, B_M)$. Thus, there is another realization $\mathcal{P}'$ of Π in G such that $e_t \notin E(\mathcal{P}')$ for some $t \in [M]$. So, $G \backslash e_t$ is a smaller counterexample, a contradiction. Thus $|E_2| < 4lM\rho_1(\Sigma)h(\Sigma)$, else we are done.

By choice of j, since there are not many edges in E_1 or E_2, it is clear that

$$|E_3| \geq 13lm(k, 4l, \Gamma)(2lm(k, 4l, \Gamma))^{2lm(k,4l,\Gamma)}h(\Sigma)^2 \max\{\rho_1(\Sigma)\rho_2(\Sigma)\}.$$

It follows that, for some hole δ_1 of Σ, there are at least

$$13lm(k, 4l, \Gamma)(2lm(k, 4l, \Gamma))^{2lm(k,4l,\Gamma)}h(\Sigma) \max\{\rho_1(\Sigma)\rho_2(\Sigma)\}$$

contractible edges of G_0 with both ends on δ_1. We denote this set as $E_3^{\delta_1}$. Recall that L_1 is the linear ordering of $V(\delta_1)$ given by the vortex (G_1, L_1). We again stress that we regard δ_1 as a linearly ordered set and L_1 is a sub-order of δ_1 (see Remark 5.2.2). Let u and v be the first and last vertices of L_1 respectively. Let e be an edge in $E_3^{\delta_1}$. Note that if Σ is the disk, we have *two* subarcs A_1 and A_2 of δ_1, such that both $e \cup A_1$ and $e \cup A_2$ bound a disk in Σ. If Σ is not the disk there is only *one* such choice. The rest of the proof is dedicated to overcoming this difficulty. For each $e \in E_3^{\delta_1}$, we let $a(e)$ be a subarc of δ_1 such that $e \cup a(e)$ bounds a disk in Σ. We say that an edge e of $E_3^{\delta_1}$ is *bad* if $a(e)$ contains both u and v, the minimal and maximal elements of L_1, and is *good* otherwise.

Claim. *$E_3^{\delta_1}$ contains fewer than $4lM$ bad edges.*

SUBPROOF. Suppose not. Observe that any two bad edges are homotopic. Hence, we have at least $4lM$ pairwise homotopic edges, and we can proceed exactly as we did in the previous case. □

Observe that there are at most $8lM(h(\Sigma) - 1)\rho_2(\Sigma)$ edges of E_1 with one endpoint on δ_1. Similarly, there are at most $4lM\rho_1(\Sigma)$ edges of E_2 with both ends on δ_1. Also, by the above claim there are at most $4lM$ bad edges in E_3. Let X be the set of vertices of G on δ_1 which are an endpoint of an edge in E_1, E_2, or a bad edge of $E_3^{\delta_1}$.

It follows that $|X| \leq 12lMh(\Sigma) \max\{\rho_1(\Sigma), \rho_2(\Sigma)\}$. Note that X partitions δ_1 into $|X|$ intervals, and that no point of X is strictly between the endpoints of a good edge of $E_3^{\delta_1}$ (with respect to L_1). It is easy to see that $E_3^{\delta_1}$ contains at least $(2lM)^{2lM}|X|$ good edges. Hence, one of these intervals contains the endpoints of at least $(2lM)^{2lM}$ good edges. We can thus finish the proof by proceeding exactly as in the proof of Lemma 9.4.5. $\square$

Chapter 10

The Algorithm

We end by giving a global overview of the algorithm.

Let G be a Γ-labelled graph and let Π be a k-pattern in G. We wish to determine whether G has a Π-linkage. We begin by testing if G has small branch-width. This can be done in linear-time by Theorem 3.1.4. If G does have small branch-width, then we can solve the problem directly by Corollary 3.6.2 (or by dynamic programming).

So, that leaves us with the case that G has huge branch-width. Here, "huge" is an enormous constant w that allows us to find the structures we require, but only depends on k and Γ. By the Grid Theorem (Theorem 5.1.2), $\widetilde{G}$ contains a large grid-minor H. We can efficiently find H by Remark 5.1.3. By Theorem 4.4.1, H induces a high order tangle $\mathcal{T}_H$ in G. By Lemma 4.4.2, the diagonal $D(H)$ of H is a large $\mathcal{T}_H$-independent subset of $V(G)$. Therefore, by Lemma 4.9.1, we can test whether any separation (A, B) is in $\mathcal{T}_H$, as long as $ord(A, B)$ is less than half the size of $D(H)$. Thus, $D(H)$ exhibits a tangle that is still of high order. Our algorithm will

- certify that G has a Π-linkage, or
- delete a redundant vertex for Π, or
- find a new tangle.

We now apply Theorem 5.6.5 to determine the structure of G relative to our current tangle in question. Either

- $\mathcal{T}_H$ controls a $K(n, \Gamma)$-minor in G, where n is still big.

- There exists a small set of vertices $X \subseteq V(G)$, such that the $\mathcal{T}$-large block of $G\backslash X$ is Γ'-balanced for some proper subgroup Γ' of Γ, and contains a big $K(\Gamma', m)$-minor.

- $\mathcal{T}_H$ does not control a big K_l minor in $\widetilde{G}$.

If $\mathcal{T}_H$ controls a $K(n, \Gamma)$-minor in G, where n is still big, then we can easily find a redundant vertex in the $K(n, \Gamma)$-minor (or certify that G has a Π-linkage), by Theorem 6.1.1.

For the second outcome we proceed as follows. We may assume that $V(\Pi) \subseteq X$, since k is fixed. Now if $G\backslash X$ only contains *one* block, then we can apply Theorem 6.2.1 to find a redundant vertex within the $K(\Gamma', m)$-minor. Suppose $G\backslash X$ has more than one block, and let B be the $\mathcal{T}_H$-large block of $G\backslash X$. A *piece* P of G is a subgraph of $G\backslash X$ which is maximal with respect to the following two properties.

- P is the union of blocks of $G\backslash X$, and

- P intersects B at a single vertex.

For each piece P of G we test if $G[V(P) \cup X]$ has branch-width at most w, where w is the same constant from the start of the algorithm. If $G[V(P) \cup X]$ has low branch-width for each piece P, then we can find the set of all realizable patterns Π' in $G[V(P) \cup X]$, with $V(\Pi') = X \cup \{y\}$, where y is the unique vertex of $V(P)$ in B. We can therefore reduce to the case where $G\backslash X$ consists of a single block, and use Theorem 6.2.1. If $G[V(P) \cup X]$ has high branch-width for some piece P, we attempt to find a new tangle. We find a large grid-minor H' in $G[V(P) \cup X]$, and let $D(H')$ be the diagonal of H'. It is possible that $\mathcal{T}_H$ and $\mathcal{T}_{H'}$ are the same tangle (up to the order that we care about). We can determine this by applying Theorem 4.9.6. If $\mathcal{T}_{H'}$ is the same tangle as $\mathcal{T}_H$, then we move to another piece and test its branch-width. If $\mathcal{T}_{H'}$ is a new tangle, then we start anew by determining the structure of G relative to $\mathcal{T}_{H'}$. If we cannot find a new tangle in this way, then $\mathcal{T}_H$ must be a leaf in the "tree of tangles". Recall that such tangles are called peripheral tangles. We can separate a peripheral tangle from all other tangles of the same order by a low order separation. So, we again reduce to the single block case, at the cost of introducing a few more linkage vertices.

The remaining outcome is if $\mathcal{T}_H$ does not control a big K_l-minor in $\widetilde{G}$. In this case, we use a constructive version of the Graph Minors Structure

Theorem. Up to 3-separations, G embeds in a surface Σ, with a bounded number of vortices of bounded adhesion, a bounded number of apex vertices A, and a large part of the grid-minor H lying in a disk of Σ. Let G_0 be the part of G embedded in Σ. We may assume that $V(\Pi) \subseteq A$, since k is fixed.

For this instance, we abuse terminology and define a *piece* of G to be a subgraph of G that is glued onto G_0 along an edge or reduction triangle. For the moment, assume that G does not contain any pieces. First we analyze how the apex vertices attach to G_0. As in the proof of Theorem 6.2.1, by flipping edges, we may assume that the edges from the apex set to G_0 are all directed toward G_0. For each $x \in A$ and $\gamma \in \Gamma$, we let

$$N_{x,\gamma} := \{u \in V(G_0 \backslash A) : e = xu \in E(G), \gamma_e = \gamma\}.$$

We let $\mathcal{N}$ denote the family of all $N_{x,\gamma}$. For the moment, assume that each member of $\mathcal{N}$ is of low rank in the tangle matroid $M_{\mathcal{T}_H}$. Thus, all the neighbours of A in G_0 are contained in a disk Δ, and there are not very many vertices of G_0 on the boundary of Δ. We may assume that $bd(\Delta) \cap G_0$ is $\mathcal{T}_H$-independent. We remove the interior of Δ from Σ, and let δ be the corresponding hole. Note that if G has a Π-linkage, then $G_0 \cap \overline{\Sigma \backslash \Delta}$ must have a Π'-linkage where $V(\Pi')$ is a subset of $bd(\Delta)$. We do not know what Π' is, but we can find a redundant vertex for *all* possible Π' via the sufficient conditions in Theorem 8.4.3. Let Γ_F be the face subgroup of Γ. We attempt to find a collection of faces of G that are pairwise far apart and strongly generate Γ_F. We then recurse, but always find faces that are far apart from the ones we have already found, and far apart from δ. If we can find a family $\{\mathcal{F}_1, \ldots, \mathcal{F}_l\}$ such that

- each $\mathcal{F}_i$ is a collection of pairwise far apart faces that strongly generate Γ_F,
- any two faces in $\mathcal{F} := \bigcup_{i\in[l]} \mathcal{F}_i$ are also far apart, and
- each face in $\mathcal{F}$ is far from δ,

then Theorem 8.4.3 implies that any Π' with its vertex set contained on δ is realizable in G_0. Thus, any vertex that is far from each face in $\mathcal{F}$ and also far from δ is redundant for Π', and hence also for Π.

If we cannot find such a collection $\mathcal{F}$, then there must be a proper subgroup Γ'_F of Γ_F and a small number of "clusters", such that every face of G_0 which is outside the union of all the clusters, has its group-value in Γ'_F. Note that we may regard each of these clusters as a vortex of bounded adhesion, by Theorem 9.1.1. We can therefore replace Γ by Γ/Γ'_F and then recurse. We have added a few more vortices, but have moved to a smaller group.

Eventually, we will either find a redundant vertex, or reduce to the case that $\Gamma = \{0\}$. By enlarging each of the vortices slightly, we may assume that they each have a linked vortex decomposition of bounded depth, by Theorem 9.3. Now we can appeal to Theorem 9.4.6, which asserts that if Π' is realizable in G_0, then it has a realization which does not pass through the vortices many times. We have therefore reduced to a problem in a surface, at the cost of introducing a few more linkage vertices. But now, any vertex that is far from δ and each of the linked vortices (including the new ones we created), is redundant for Π' in G_0 and hence also for Π in G.

In the case that some members of $\mathcal{N}$ are of high rank in the tangle matroid, we proceed exactly as in the proof of Theorem 6.2.1. Namely, we divide $\mathcal{N}$ into the low rank sets and the high rank sets. The low rank sets are contained in a disk in Σ, with not many vertices on the boundary, and the high rank sets are dispersed all over the surface and do not cause us problems.

Finally, we need to deal with the fact that G is only embedded in Σ up to 3-separations. Again, we handle this in the same way as we did for cliques. We test the branch-width of each piece of G in an attempt to find a new tangle. If we discover a new tangle, we again start over and find the structure of G relative to our new tangle. Note that this new tangle may correspond to either the clique case or the surface case. Either way, we keep a list of all the tangles that we have found so far and check potential new tangles against all tangles in our list. Since G has at most $(|E(G)|-2)/2$ maximal tangles by Corollary 4.8.3, we are guaranteed after at most $(|E|-2)/2$ steps to

- delete a redundant vertex, or
- certify that Π is realizable, or
- reach a peripheral tangle.

Once we reach a peripheral tangle we will certainly find a redundant vertex or certify that Π is realizable in the next step. Note that we only need to keep a list of tangles until we delete a redundant vertex. Once we delete a redundant vertex v, the algorithm begins again in earnest and tests whether $G\backslash v$ has branch-width at most w.

As stated, the running-time of the algorithm is certainly polynomial, but it is far from optimal since we are constructing the tree of tangles along the way. If we only care about finding a peripheral tangle, the running-time of the algorithm can be improved to $O(|V(G)|^6)$.

Bibliography

[1] Stefan Arnborg, Jens Lagergren, and Detlef Seese. Easy problems for tree-decomposable graphs. *J. Algorithms*, 12(2):308–340, 1991. 3.6.1

[2] Hans L. Bodlaender and Dimitrios M. Thilikos. Constructive linear time algorithms for branchwidth. In *Automata, languages and programming (Bologna, 1997)*, volume 1256 of *Lecture Notes in Comput. Sci.*, pages 627–637. Springer, Berlin, 1997. 3.1

[3] John M. Boyer and Wendy J. Myrvold. On the cutting edge: simplified $O(n)$ planarity by edge addition. *J. Graph Algorithms Appl.*, 8(3):241–273 (electronic), 2004. 1.4

[4] Maria Chudnovsky, William H. Cunningham, and Jim Geelen. An algorithm for packing non-zero A-paths in group-labelled graphs. *Combinatorica*, 28(2):145–161, 2008. 5.6.3

[5] Bruno Courcelle. Graph rewriting: an algebraic and logic approach. In *Handbook of theoretical computer science, Vol. B*, pages 193–242. Elsevier, Amsterdam, 1990. 3, 3.6

[6] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, third edition, 2005. 5.2

[7] Reinhard Diestel, Tommy R. Jensen, Konstantin Yu. Gorbunov, and Carsten Thomassen. Highly connected sets and the excluded grid theorem. *J. Combin. Theory Ser. B*, 75(1):61–73, 1999. 5.1

[8] T. A. Dowling. A class of geometric lattices based on finite groups. *J. Combinatorial Theory Ser. B*, 14:61–86, 1973. 2.3

[9] H.-D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer, 2nd edition, 1994. 3.2

[10] Jack Edmonds. Submodular functions, matroids, and certain polyhedra. In *Combinatorial Structures and their Applications (Proc. Calgary Internat. Conf., Calgary, Alta., 1969)*, pages 69–87. Gordon and Breach, New York, 1970. 2.6, 2.6

[11] Jack Edmonds and Rick Giles. A min-max relation for submodular functions on graphs. In *Studies in integer programming (Proc. Workshop, Bonn, 1975)*, pages 185–204. Ann. of Discrete Math., Vol. 1. North-Holland, Amsterdam, 1977. 4.7

[12] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multi-commodity flow problems. In *16th Annual Symposium on Foundations of Computer Science (Berkeley, Calif., 1975)*, pages 184–193. IEEE Computer Society, Long Beach, Calif., 1975. 1.4

[13] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM J. Comput.*, 5(4):691–703, 1976. 1.4

[14] Jim Geelen and Bert Gerards. Excluding a group-labelled graph. *J. Comb. Theory Ser. B*, 99(1):247–253, 2009. 5.6, 5.6, 5.6

[15] Jim Geelen, Bert Gerards, and Geoff Whittle. Towards a matroid-minor structure theory. *Manuscript*. 1.1

[16] Jim Geelen, Bert Gerards, and Geoff Whittle. Tangles, tree-decompositions and grids in matroids. *J. Comb. Theory Ser. B*, 99(4):657–667, 2009. 4.6, 4.6, 4.7, 4.8, 4.8

[17] Jim Geelen, Bert Gerards, and Geoff Whittle. Well-quasi-ordering binary matroids. *In preparation*, 2009. 1.1, 2.5

[18] Jim Geelen and Bruce Richter. On graph minors in surfaces. *Preprint*, 2008. 7.2

[19] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981. 4.9

[20] Satoru Iwata. A faster scaling algorithm for minimizing submodular functions. *SIAM J. Comput.*, 32(4):833–840 (electronic), 2003.

[21] Satoru Iwata, Lisa Fleischer, and Satoru Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *J. ACM*, 48(4):761–777 (electronic), 2001. 4.9

[22] H. A. Jung. Eine Verallgemeinerung des n-fachen Zusammenhangs für Graphen. *Math. Ann.*, 187:95–103, 1970. 1.4

[23] M. Juvan, A. Malnič, and B. Mohar. Systems of curves on surfaces. *J. Combin. Theory Ser. B*, 68(1):7–22, 1996. 7.1

[24] R.M. Karp. On the computational complexity of combinatorial problems. *Networks*, 5, 1975. 1.4

[25] Ken-ichi Kawarabayashi, Zhentao Li, and Bruce Reed. Faster minor recognition. *In preparation*, 2009. 1.4, 5.5

[26] Ken-ichi Kawarabayashi and Bojan Mohar. Some recent progress and applications in graph minor theory. *Graphs Combin.*, 23(1):1–46, 2007. 5.3

[27] D. G. Larman and P. Mani. On the existence of certain configurations within graphs and the 1-skeletons of polytopes. *Proc. London Math. Soc. (3)*, 20:144–160, 1970. 1.4

[28] W. Mader. Homomorphieeigenschaften und mittlere Kantendichte von Graphen. *Math. Ann.*, 174:265–268, 1967. 5.5

[29] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001. 5.2

[30] James G. Oxley. *Matroid theory*. Oxford Science Publications. The Clarendon Press Oxford University Press, New York, 1992. 2, 2.6

[31] Maurice Queyranne. Minimizing symmetric submodular functions. *Math. Programming*, 82(1-2, Ser. B):3–12, 1998. Networks and matroids; Sequencing and scheduling. 4.9

[32] Bruce Reed. Tree width and tangles: A new connectivity measure and some applications. *Survey in Combinatorics*, 241:87–158, 1997. 4.8

[33] Bruce Richter. Graph minors: Important first steps and generalizating Kuratowski's theorem. *Manuscript*. 3.1

[34] Gerhard Ringel and J. W. T. Youngs. Solution of the Heawood map-coloring problem. *Proc. Nat. Acad. Sci. U.S.A.*, 60:438–445, 1968. 1.1

[35] Neil Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *J. Combin. Theory Ser. B*, 41(1):92–114, 1986. 5.1

[36] Neil Robertson and P. D. Seymour. Graph minors. VI. Disjoint paths across a disc. *J. Combin. Theory Ser. B*, 41(1):115–138, 1986. 8.4

[37] Neil Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, 52(2):153–190, 1991. 3.1, 4, 4.1, 4.1, 4.4, 4.5, 4.6

[38] Neil Robertson and P. D. Seymour. Graph minors. XI. Circuits on a surface. *J. Combin. Theory Ser. B*, 60(1):72–106, 1994. 7.4, 7.5

[39] Neil Robertson and P. D. Seymour. Graph minors. XII. Distance on a surface. *J. Combin. Theory Ser. B*, 64(2):240–272, 1995. 7.5, 7.5

[40] Neil Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory Ser. B*, 63(1):65–110, 1995. 1.1, 1.4, 1.4

[41] Neil Robertson and P. D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Combin. Theory Ser. B*, 89(1):43–76, 2003. 5.3, 5.4

[42] Neil Robertson and P. D. Seymour. Graph minors. XX. Wagner's conjecture. *J. Combin. Theory Ser. B*, 92(2):325–357, 2004. 1.1, 2.5

[43] Neil Robertson and P. D. Seymour. Graph minors. XXI. Graphs with unique linkages. *J. Comb. Theory Ser. B*, 99(3):583–616, 2009. 1.4, 9.4

[44] Neil Robertson and P. D. Seymour. Graph minors. XXII. Irrelevant vertices in linkage problems. *Preprint*, 2009. 1.4

[45] Gian-Carlo Rota. Combinatorial theory, old and new. In *Actes du Congrès International des Mathématiciens (Nice, 1970), Tome 3*, pages 229–233. Gauthier-Villars, Paris, 1971. 2.5

[46] Alexander Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *J. Combin. Theory Ser. B*, 80(2):346–355, 2000. 4.9

[47] D. G. Seese and W. Wessel. Grids and their minors. *J. Combin. Theory Ser. B*, 47(3):349–360, 1989. 5.3

[48] P. D. Seymour. Disjoint paths in graphs. *Discrete Math.*, 29(3):293–309, 1980. 1.4.4

[49] Yossi Shiloach. A polynomial solution to the undirected two paths problem. *J. Assoc. Comput. Mach.*, 27(3):445–456, 1980. 1.4.4

[50] Richard P. Stanley. *Enumerative combinatorics. Vol. 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997. With a foreword by Gian-Carlo Rota, Corrected reprint of the 1986 original. 2.3

[51] John Stillwell. *Classical topology and combinatorial group theory*, volume 72 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1993. 8.4

[52] Robert E. Tarjan. Decomposition by clique separators. *Discrete Math.*, 55(2):221–232, 1985. 5.5

[53] Robin Thomas and Paul Wollan. An improved linear edge bound for graph linkages. *European J. Combin.*, 26(3-4):309–324, 2005. 1.4

[54] Carsten Thomassen. 2-linked graphs. *European J. Combin.*, 1(4):371–378, 1980. 1.4.4

[55] W. T. Tutte. A homotopy theorem for matroids. I, II. *Trans. Amer. Math. Soc.*, 88:144–174, 1958. 2.2

[56] Thomas Zaslavsky. Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory Ser. B*, 47(1):32–52, 1989. 1.1, 2

[57] Thomas Zaslavsky. Biased graphs. II. The three matroids. *J. Combin. Theory Ser. B*, 51(1):46–72, 1991. 1.1, 2