Łukasz Bożyk

New Results in (Ob)structural Graph Theory

PhD dissertation

University of Warsaw

Faculty of Mathematics, Informatics and Mechanics

Supervisor:
dr hab. Michał Pilipczuk

Institute of Informatics
University of Warsaw

September 2023

Contents

Abstract

We present a collection of results in structural graph theory and the theory of parameterized algorithms, all of which are based upon understanding the structure of graphs disallowed to contain certain objects (called *obstructions*). The results involve the settings of both undirected and directed graphs, and different notions of graph containment and graph decomposition.

First, we focus on the class of tournaments. We establish the so-called Erdős-Pósa property for (directed) immersions and topological minors, i.e. we prove that for every simple digraph H, $k \in \mathbb{N}$, and tournament T, the following statements hold:

- If in T one cannot find k arc-disjoint immersion copies of H, then there exists a set of $\mathcal{O}_H(k^3)$ arcs that intersects all immersion copies of H in T.
- If in T one cannot find k vertex-disjoint topological minor copies of H, then there exists a set of $\mathcal{O}_H(k \log k)$ vertices that intersects all topological minor copies of H in T.

Moreover, for a fixed simple digraph H without isolated vertices, we consider the problem of deleting arcs from a given tournament to get a digraph which does not contain H as an immersion. We prove that for every H, this problem admits a polynomial kernel when parameterized by the number of deleted arcs.

Next, we turn our attention to width notions other than the well-established treewidth and we make an attempt of identification of the related obstructions. In particular we focus on tree-cut width — a graph parameter corresponding to immersion containment in the following sense: the tree-cut width of a graph is functionally equivalent to the largest size of a wall contained in it as an immersion. We propose a new definition of tree-cut width, functionally equivalent to the ones previously existing, but for which we can prove a tight duality theorem relating it to naturally defined dual objects: appropriately defined brambles and tangles. Using this result we also design a game characterization of tree-cut width.

We aim towards understanding the width notions and obstructions in the setting of general digraphs (not necessarily tournaments) as well. We identify an object, which we call an *onion star*, and explain the structure of digraphs excluding a fixed onion star as an immersion. The main discovery is that in such digraphs we can prove directed analogues of some duality statements true in the undirected setting. Specifically, we show the following two theorems.

- There is a function $f\colon \mathbb{N} \to \mathbb{N}$ satisfying the following: If a digraph D contains a set X of $2t+1$ vertices such that for any $x, y \in X$ there are $f(t)$ arc-disjoint paths from x to y, then D contains the t-onion star as an immersion.
- There is a function $g\colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ satisfying the following: If x and y is a pair of vertices in a digraph D such that there are at least $g(t,k)$ arc-disjoint paths from x to y and there are at least $g(t,k)$ arc-disjoint paths from y to x, then either D contains the t-onion star as an immersion, or there is a family of $2k$ pairwise arc-disjoint paths with k paths from x to y and k paths from y to x.

Finally we consider the class of bipartite permutation graphs and the parameterized complexity of the (NP-complete) bipartite permutation vertex deletion problem, asking, for a given n-vertex graph, whether one can remove at most k vertices to obtain a bipartite permutation graph. We analyze the structure of what we call almost bipartite permutation graphs (the difference being allowing them to contain long induced cycles) and investigate the structural properties of the shortest hole in such a graph. We use it to obtain an algorithm for the bipartite permutation vertex deletion problem with running time $\mathcal{O}(9^k \cdot n^9)$, and also give a polynomial-time 9-approximation algorithm.

Acknowledgements

First and foremost, I want to thank my supervisor, Michał Pilipczuk, for being my guiding light over the past four years. For continuous support, motivation, and willingness to share knowledge, intuition and ideas which have been invaluable. He also involved me in various projects, giving me a taste of different aspects of research.

I am deeply grateful to Marcin Pilipczuk for hosting me under the CUTACOMBS grant for most of my time in the Doctoral School at the University of Warsaw. Without this support, I would not have had the opportunity to undisturbedly develop my research skills and connect with amazing people from all over the world.

I want to express my appreciation to all the wonderful colleagues I have had the pleasure of working with. Special thanks to Karolina Okrasa and Bartłomiej Kielak for numerous memorable and enjoyable math-and-smile-filled meetings. To all my other co-authors: Marthe Bonamy, Oscar Defrain, Jan Derbisz, Andrzej Grzesik, Meike Hatzel, Tomasz Krawczyk, Tomáš Masařík, and Jana Masaříková.

To my friends and my family, it is you who have been making this time meaningful and complete. To my Piotrek, ∀.

PS I am grateful that the following inspiring and motivating concepts, objects and phenomena exist; thanks for shaping this work and keeping me who I am (in no particular order): badminton, SET, OMJ, Fundusz, MBL, vegetables, σ-fields, escape rooms, Italian cuisine and Tuza's conjecture.

Attribution of support. Author's work leading to the results presented herein was supported by funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement No. 714704.

Attribution of graphics. Most of the illustrations are directly taken from joint publications. The pictures presented in Figures 5.2 to 5.5, 6.2 and 6.3 were created by Oscar Defrain, the pictures presented in Figures 1.4 and 7.1 to 7.4 were created by Tomasz Krawczyk, the diagram presented in Figure 1.1 was created by Karolina Okrasa. The remaining pictures were created by the author.

Notation

Here we present the common notions and definitions used throughout the thesis. Note that some additional chapter-specific objects and concepts will be defined later.

General conventions. Denote by $\mathbb{N}$ the set of positive integers. For integers m, n with $m \leqslant n$ we set:

$$\begin{aligned}
[m,n] &:= \{m, m+1, \ldots, n\}, \\
[\![n]\!] &:= [0,n] = \{0,1,2,\ldots,n\}, \\
[n] &:= [1,n] = \{1,2,\ldots,n\}, \\
[-n] &:= [-n,-1] = \{-1,-2,\ldots,-n\}, \\
\log n &:= \log_2 n, \\
f^{\langle n \rangle}(x) &:= \underbrace{f(f(\ldots f}_{n \text{ times}}(x)\ldots)), \text{ where } f \text{ is a function whose domain and codomain are equal.}
\end{aligned}$$

A *near-partition* of a set Ω is a family $\mathcal{X}$ of pairwise-disjoint subsets of Ω such that $\Omega = \bigcup \mathcal{X}$. If moreover all elements of $\mathcal{X}$ are nonempty, we call $\mathcal{X}$ a *partition* of Ω.

Graphs. We follow the classical graph theory terminology an assume that the reader is already familiar with it. All graphs considered in this thesis are finite. Undirected graphs are considered to be simple (i.e. without loops or parallel edges) unless explicitly stated otherwise. Directed graphs are loopless but allowed to have parallel arcs as long as the arcs are oriented differently (i.e. there are no multiple arcs with same head and tail).

For a (possibly directed) graph G we denote:

$V(G)$		the set of all vertices of G,
$E(G)$		the set of all edges of G,
$A(G)$		the set of all arcs of G (if G is directed),
$\lvert G\rvert$	$:=$	$\lvert V(G)\rvert$,
$\lVert G\rVert$	$:=$	$\lvert V(G)\rvert + \lvert E(G)\rvert$ ($A(G)$ if G is directed),
$G[X]$		the subgraph of G *induced* by set $X \subseteq V(G)$, i.e. comprising of all vertices in X and all edges/arcs with both endpoints in X,
$G - X$	$:=$	$G[V(G) \setminus X]$, where $X \subseteq V(G)$,
$G - F$		the graph $(V(G), E(G) \setminus F)$, where $F \subseteq E(G)$ ($A(G)$ if G is directed),
$N(x)$	$:=$	$\{v \in V(G) \mid xv \in E(G)\}$ — the *neighborhood* of $x \in V(G)$,
$N(X)$	$:=$	$\bigcup_{x \in X} N(x) \setminus X$, where $X \subseteq V(G)$,
$\delta(X)$		the set of all edges with one endpoint in X and the other in $V(G) \setminus X$, where $X \subseteq V(G)$,
$\deg_G(x)$	$:=$	$\lvert\delta(\{x\})\rvert$ the *degree* of a vertex $x \in V(G)$ (equal to the number of neighbors of x if G is simple).

We say that two disjoint sets $A, B \subseteq V(G)$ are *complete* to each other if every $a \in A$ is adjacent to every $b \in B$. Similarly, A and B are *anti-complete* to each other if every $a \in A$ is nonadjacent to every $b \in B$.

Let G be an undirected graph. A *path* in G is a connected subgraph of G where every vertex has degree 2 apart from exactly two vertices of degree one, called the *endpoints* of the path. A *cycle* in G is a connected subgraph of G where every vertex has degree 2. The *length* of a path or a cycle is defined as its edge count. Two vertices connected by a pair of parallel edges (if we allow them) are considered a cycle of length 2.

An undirected graph G is *bipartite* if there exists a partition of $V(G)$ into sets X and Y such that each edge of G has one endpoint in X and another in Y. A bipartite graph is *balanced* if there exists a choice of the bipartition classes of G such that the numbers of vertices in the two classes are equal. By $K_{n,n}$ we denote the complete balanced graph on $2n$ vertices.

We denote a complete graph and a cycle on n vertices by K_n and C_n, respectively.

Let $u, v \in V(G)$. We say that u and v *are at distance k (in G)* if k is the length of a shortest path between u and v in G. For a set of edges $F \subseteq E(G)$, we say that the pair u, v is *disconnected* by F if u and v belong to different connected components of $G - F$.

Digraphs. For a directed graph (*digraph*) D we denote:

$$\begin{aligned}
\overrightarrow{A}(X,Y) &\coloneqq \{uv \in A(D)\colon\ u \in X \text{ and } v \in Y\}, \text{ where } X, Y \subseteq V(D),\\
A(X,Y) &\coloneqq \overrightarrow{A}(X,Y) \cup \overrightarrow{A}(Y,X),\\
\mathsf{tail}(a) &\coloneqq u, \text{ where } a = uv \in A(D),\\
\mathsf{head}(a) &\coloneqq v, \text{ where } a = uv \in A(D),\\
\delta^+(X) &\coloneqq \{a \in A(D) : \mathsf{tail}(a) \in X \text{ and } \mathsf{head}(a) \notin X\}, \text{ where } X \subseteq V(D),\\
\delta^-(X) &\coloneqq \{a \in A(D) : \mathsf{tail}(a) \notin X \text{ and } \mathsf{head}(a) \in X\}, \text{ where } X \subseteq V(D),\\
&\ (\text{if } X = \{x\}, \text{ we omit internal brackets and write } \delta^+(x) \text{ and } \delta^-(x), \text{ respectively}).
\end{aligned}$$

A *strong component* of D is an inclusion-wise maximal induced subgraph C of D that is *strongly connected*, that is, for every pair of vertices u and v of C, there are directed paths in C both from u to v and from v to u.

A *directed path* is a (not necessarily simple) path in which all edges are consistently oriented. We remark that in Chapter 6, we will introduce an unusual exact formalization of this definition (treating a directed path as a sequence of distinct arcs). Regardless of the definition, for a directed (and not necessarily simple) path P we denote:

$$\begin{aligned}
\mathsf{first}(P) &\quad \text{the first arc on } P,\\
\mathsf{last}(P) &\quad \text{the last arc on } P.
\end{aligned}$$

A *cut* in D is a partition (A, B) of $V(D)$. An (a, b)-*cut* in D, where $a, b \in V(D)$, is a cut (A, B) with $a \in A$ and $b \in B$. The *size* of the cut (A, B) is $|\delta^+(A)| = |\delta^-(B)|$.

The classical theorem of Menger describes the relation between the size of (a, b)-cuts and the number of arc-disjoint paths connecting a and b.

Theorem (Menger [Men27]). *Let D be a finite digraph and let $a, b \in V(D)$. The maximum number of arc-disjoint (a, b)-paths equals the minimum size of an (a, b)-cut.*

We remark that Menger's Theorem is true in the undirected setting as well, and moreover it is valid for *multi(di)graphs* (i.e. with parallel edges or arcs allowed).

Theorem (Menger, undirected version). *Let G be a finite multigraph and let $a, b \in V(G)$. The maximum number of mutually edge-disjoint paths from a to b equals the minimum number of edges whose removal disconnects a and b.*

Chapter 1

Introduction

Obstruction is one of the central notions in classical and modern structural graph theory, as numerous vital results have the following generic form: either a graph admits some *global structural property*, or it contains a *local obstruction* (certifying lack of the structure).

The celebrated theorem of Leonhard Euler, sometimes considered to be the inauguration of graph theory, can be stated in a way resembling the above description: a connected graph either has a circuit visiting each edge exactly once, or it has a vertex of odd degree [Hie73].

Other classical examples are the theorems characterizing planar graphs due to Kuratowski [Kur30] and Wagner [Wag37]: either a graph is planar (meaning it can be drawn in the Euclidean plane with no edge intersections), or it *contains* at least one of the two canonical obstructions, K_5 or $K_{3,3}$. Note that the notion of containment is imprecise and can be formalized in several different ways. In Kuratowski's theorem we have the so-called *topological containment*: a graph H is a *topological minor* of graph G if there is a map sending vertices of H to mutually different vertices of G and edges of H to mutually internally vertex-disjoint paths joining the images of respective vertices. In Wagner's theorem in turn, we are dealing with *minors*.

Graph H is a *minor* of graph G if a subset of $V(G)$ can be decomposed into pairwise disjoint non-empty sets indexed by $V(H)$ in such a way that each set induces a connected subgraph of G and whenever $uv \in E(H)$, there exists an edge in G joining some vertex from the set corresponding to u with some vertex from the set corresponding to v. Alternatively, and equivalently, H is a minor of G if it can be obtained from G by a sequence of edge deletions, vertex deletions, or edge contractions (replacing a pair of neighbors with a single vertex connected to all other neighbors of the pair).

The Graph Minor series of Robertson and Seymour has brought fundamental understanding to the minor order on graphs. In particular, by proving that it is a well-quasi-ordering [RS04], they have deeply generalized Wagner's result: every class of graphs closed under the operation of taking minors can be characterized by a finite list of minimal forbidden minors. On the way to that result, they have introduced and proved numerous notions and theorems still applied and inspiring new research directions in modern structural graph theory.

The following notion is crucial in the theory of graph minors. Suppose a graph G is represented as a union of its subgraphs X_i indexed by nodes of a tree T in a way that whenever $i, j, k \in V(T)$ with j lying on the path from i to k, we have $V(X_i) \cap V(X_k) \subseteq V(X_j)$. Then the family $(X_i)_{i \in V(T)}$ is called a *tree-decomposition* of G with *bags* X_i. The *treewidth* of G is the smallest integer k such that G admits a tree-decomposition with all bags of size at most $k+1$. Intuitively, treewidth is a parameter encapsulating tree-likeness of a graph.

The works of Robertson and Seymour have identified obstructions for a graph to be tree-like. In [RS86] they have proved that for every planar graph H there exists $k \in \mathbb{N}$ such that if a graph does not contain H as a minor, then its tree-width is at most k. Specifically, as forbidding a planar graph causes excluding a grid minor (of size dependent only on H), this boils down to the celebrated Grid Minor Theorem: for every grid H, every graph whose treewidth is large enough relative to $|H|$ contains H as a minor. In essence, large grid minors are obstructions to having low treewidth. The works of Robertson and Seymour reach far beyond grid minors: in particular their Structure Theorem can be used to describe graphs which exclude a fixed graph H as a minor [RS03].

Moreover, there is tight duality between treewidth and other obstructions, the so-called *brambles* and *tangles*. A bramble is a family of connected subgraphs of a given graph with the property that the union of vertex sets of any two of them induces a connected subgraph as well. The *order* of a bramble is the minimum size of a (vertex) set *hitting* (i.e. having nonempty intersection with) all of its elements. A tangle of *order* k is in turn an orientation of (vertex) separations of size smaller than k satisfying a number conditions (intuitively, each separation is to be oriented towards the part of the graph with more complicated structure and this orientation needs to be globally consistent).

The precise relationship between the aforementioned objects and treewidth, as proved in part by Robertson and Seymour in [RS91], and later by Seymour and Thomas in [ST93], is given by equivalence of the following condition for a graph G and a positive integer k:

- G has no tree decomposition of width $< k$;
- G has a bramble of order $\geqslant k$; and
- G has a tangle of order $\geqslant k$.

A corollary of the above result is an elegant characterization of treewidth in terms of the cops and robber game, described by Seymour and Thomas in [ST93]. In this game k cops, occupying vertices of a given graph, aim to catch a robber moving along edges of that graph (both cops and the robber take turns and move accordingly with a number of technical rules). The result states that k cops can catch the robber if and only if the treewidth of the graph is less than k.

Understanding graphs of bounded treewidth has also found numerous applications in algorithm design, especially in the field of parameterized algorithms — it turns out that many problems (NP-)hard in general become tractable when parameterized by treewidth of a graph. Or even more generally — for well-structured graphs, e.g. those avoiding certain obstructions.

Beyond minors. Research directions similar to those presented in Graph Minor series have been undertaken for different notions of graph embeddings than minors. Among them, perhaps most notably, is the immersion order on graphs. Immersibility is an embedding notion based on edge-disjointness and edge cuts, as opposed to the notions of minors and of topological minors, which are based on vertex-disjointness and (vertex) separations.

A graph H is an *immersion* of a graph G if there exists an injective mapping of vertices of H to vertices of G sending edges of H to pairwise edge-disjoint paths in G so that every such path connects images of the endpoints of its corresponding edge. Just as to the minor order, the immersion order is a well-quasi-ordering on all graphs [RS10], which suggests that a decent structural theory related to immersions should exist.

In [Wol15] Wollan introduced the graph parameter *tree-cut width*, which can be considered an analogue of treewidth for immersions, and proved similar duality theorems, mirroring the Grid Minor theorem of Robertson and Seymour. More precisely, Wollan proved that there is a functional dependence between the size of the largest wall that can be immersed in a graph and the tree-cut width of that graph. This implies that a class of graphs has a bounded tree-cut width if and only if all the graphs from exclude some fixed subcubic planar graph as an immersion.

Giannopoulou et al. [GPR+21] proposed another, roughly equivalent, way of measuring the graph complexity corresponding to immersibility. Their way of defining *tree-cut decomposition* is as follows. A tree-cut decomposition of a graph G is a tree T and a collection of *bags* $(X_i)_i \in V(T)$ forming a near-partition of $V(G)$. Thus, every edge e of T naturally corresponds to a partition of $V(G)$ into two parts, corresponding to unions of bags of the two connected components of $T - e$. The set of edges of G crossing this partition is called the *adhesion* of e. Roughly speaking (and skipping technicalities), a tree-cut decomposition as above has a bounded width if all of the following quantities are bounded: the size of each bag, the size of each adhesion, and the degree of each node in T. We remark that Wollan's original definition was also based on tree-cut decompositions, but defined in a more convoluted and unmanageable way.

Digraphs and tournaments. The notions of minors, subdivisions and immersions can be naturally lifted to the setting of directed graphs. For example a directed graph H can be immersed in a digraph D if one can find a mapping sending vertices of H to pairwise different vertices of D, and arcs of H to pairwise arc-disjoint oriented paths in H connecting the images of endpoints (topological minors are defined in the same way except the paths are required to be vertex-disjoint rather than arc-disjoint).

There has been significant interest in research towards understanding structural properties of those notions. In particular, Kawarabayashi and Kreutzer [KK15] proved the directed counterpart of the Grid Minor Theorem, while a series of recent papers [GKKK20, GKKK22] is working towards a directed analogue of the Structure Theorem. This brings a natural question about existence of a similar structure theory for directed immersions.

It seems that so far, not much is known for general digraphs. However, the situation changes when we restrict our attention to *tournaments*, i.e. oriented complete graphs. A meaningful structural theory for tournaments[1] with forbidden immersions was pioneered by Chudnovsky, Ovetsky Fradkin, Kim, and Seymour [COS12, CS11, OS13, OS15, KS15], see also [FP19].

In particular, as proved in the aforementioned works, there are two main width notions for tournaments: *cutwidth* and *pathwidth*. Intuitively, both of them can be used to impose a linear layout on the tournament. The first one is tightly relateed to (directed) immersions as follows: if a tournament T excludes a digraph H as an immersion, then the cutwidth of T is bounded by a constant depending only on H. Pathwidth is connected to *topological minors* and *strong minors* in the same sense. These structural results were used for the design of parameterized algorithms for containment problems in tournaments in [COS12, OS13, FP19].

Hitting-packing duality. A frequent and important kind of problems related to obstructions are those concerning hitting and packing in various combinatorial contexts.

The first result of this kind in graph theory was delivered by Erdős and Pósa [EP65], who proved that for every undirected graph G and $k \in \mathbb{N}$, one can find in G either k vertex-disjoint cycles, or a set of $\mathcal{O}(k \log k)$ vertices that meets all the cycles. This theorem may be viewed as relation between packing and hitting minor models of a triangle, which suggests the following generalization. We say that H has the *Erdős-Pósa property for minors* if there exists a function f such that for every graph G and $k \in N$, one can find in G either k vertex-disjoint minor models of H, or a set of at most $f(k)$ vertices that meets all minor models of H. Robertson and Seymour [RS86] proved that a graph H has the Erdős-Pósa property for minors if and only if H is planar.

Since the work of Erdős and Pósa, establishing the Erdős-Pósa property for different objects in graphs, as well as finding tight estimates on the best possible *bounding functions* f, became a recurrent topic in graph theory and there are still many open problems in this area (cf. a website maintained by Raymond [Ray] for an overview of the current state of knowledge on Erdős-Pósa problems). In particular, there is still much to be done in the directed setting. The analogue of the result of Erdős and Pósa for packing and hitting directed cycles was established by Reed et al. [RRST96], while a characterization of strongly connected digraphs H possessing the Erdős-Pósa property for topological minors was recently announced by Amiri et al. [AKKW16].

Paremeterized algorithms. In well-structured graphs many algorithmic problems which are hard in general turn out to be tractable. This can be formalized in the language of parameterized complexity, where the parameter is usually a small number corresponding to structural complexity of the graph, e.g. one of its width measures.

Formally, an input of a parameterized problem Π consists of an instance I and a *parameter* $k \in \mathbb{N}$. We say that Π is *fixed parameter tractable* (FPT) if there exists an algorithm deciding whether (I, k) is a **yes**-instance of Π in time $f(k) \cdot |I|^{\mathcal{O}(1)}$, where f is some computable function. Intuitively, the superpolynomial factor in algorithm's time complexity is a function of the parameter (which is small in well-structured graph).

For instance many problems which are NP-hard in general, e.g. Vertex Cover, Hamiltonian Cycle, Dominating Set, turn out to be FPT when parameterized by treewidth of a graph. In fact, the celebrated theorem of Courcelle [Cou90] states that every graph property expressible in Monadic Second Order logic (i.e. allowing quantification over sets of vertices and edges of the graph) is decidable in time $f(k) \cdot n$ for n-vertex graphs of treewidth k.

Hereditary classes. So far we have not mentioned the perhaps simplest notion of graph containment — the induced subgraph order. It turns out that classes of graphs closed under taking induced subgraphs, or — equivalently — closed under vertex deletion, behave particularly well in the context of many standard computational problems. Such classes are called *hereditary*.

Due to the practical and theoretical applications, some of such graph classes are particularly intensively studied. Among them are:

- *perfect graphs*: graphs in which the size of the largest complete subgraph is equal to the size of the smallest complete graph in which the given graph is homomorphically embeddable,

[1] In this line of work, most results concern the class of *semi-complete digraphs*, which differ from tournaments by allowing that a pair of vertices can be also connected by two oppositely-oriented arcs. We focus on the setting of tournaments for simplicity.

- *interval graphs*: intersection graphs of intervals on the real line,
- *proper interval graphs*: intersection graphs of intervals on the real line none of which is contained in another,
- *chordal graphs*: intersection graphs of subtrees of a tree,
- *function and permutation graphs*: intersection graphs of continuous and linear functions, respectively, defined on the interval $[0, 1]$,
- *comparability graphs*: graphs whose edges correspond to the pairs of vertices comparable in some fixed partial order $<$ on the vertex set (such an order is called a *transitive orientation* of the graph),
- *co-comparability graphs*: the complements of comparability graphs.

It is well known that the class of function graphs corresponds to the class of co-comparability graphs [GRU83], and the class of permutation graphs corresponds to the intersection of comparability and co-comparability graphs [PLE71] (see Figure 1.1 for the hierarchy of inclusions).

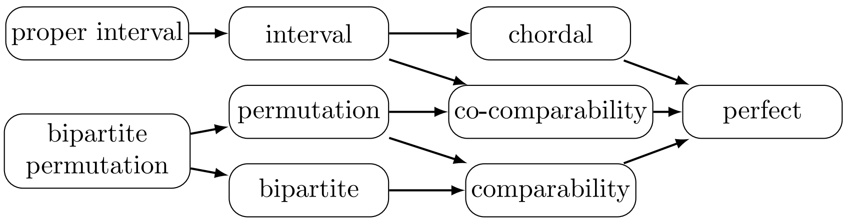

Figure 1.1: Hierarchy of inclusions between selected subclasses of perfect graphs. An arrow from graph class $\mathcal{A}$ to graph class $\mathcal{B}$ indicates that $\mathcal{A} \subset \mathcal{B}$.

Being hereditary is a very useful property in algorithmic design as every such class of graphs can be uniquely characterized in terms of minimal obstructions: a graph belongs to a class $\mathcal{G}$ if and only if it does not contain any graph from some family $\mathcal{F}$ as an induced subgraph. For every graph class introduced above, a characterization by forbidden subgraphs is known, see [CRST06] for perfect graphs, [LB62] for interval graphs, [Gal67] for comparability and permutation graphs. However, for all of them, the family of forbidden subgraphs is infinite and it may also be quite complex.

Polynomial-time algorithms devised for the above-mentioned graph classes can sometimes be adjusted to also work on graphs that are "close" to graphs from these classes. Usually, the "closeness" of a graph G to a graph class $\mathcal{G}$ is measured by the number of operations required to transform G into a graph from the class $\mathcal{G}$. In particular we have the following parameterized problem.

$\mathcal{G}$-VERTEX DELETION

Input: A graph G (typically not from $\mathcal{G}$) and a number k
Parameter: k
Output: Is there a set $X \subseteq V(G)$, such that $|X| \leqslant k$ and $G - X$ belongs to $\mathcal{G}$?

We remark that there are closely related variants of this problem, differing by the modifications allowed: edge deletion problem, edge completion problem, and edge edition problem (the latter allowing both deletions and additions of edges). For each class of graphs defined above, all four variants of the modification problem are NP-hard — see [Man08] for references to NP-hardness proofs. In particular, Lewis and Yannakakis [LY80] showed that the vertex deletion problem into any non-trivial hereditary class of graphs is NP-hard.

It turns out that characterizations by forbidden structures are sometimes useful to design FPT algorithms for graph modification problems. For example, Cai [Cai96] proposed an FPT algorithm for modification problems into classes of graphs characterized by a finite family of forbidden induced subgraphs $\mathcal{F}$. His algorithm identifies a forbidden structure in the input graph (which can be done in polynomial time when $\mathcal{F}$ is finite) and branches over all possible ways of modifying that structure. Since the families of forbidden structures are infinite for graph classes introduced above, modification algorithms for these classes have to be much more sophisticated. For several of them modification problems have satisfactory solutions:

- chordal graphs: all four versions of the modification problem are FPT [CM16, Mar10];
- interval graphs: edge completion and edge deletion are FPT [VHPT09, Cao16], vertex deletion is FPT [CM16], the complexity of edge edition remains open;
- proper interval graphs: all four versions of the modification problem are FPT [Cao17].

On the other hand, it is known that the vertex deletion to perfect graphs is W[2]-hard [HvtHJ+13]. It is unknown whether comparability graphs, co-comparability graphs, and permutation graphs have FPT modification algorithms.

Kernelization. *Kernelization* is a preprocessing routine that identifies and reduces "simple" parts of an instance of a hard problem. In the parameterized complexity theory we can formalize this concept as follows. A *kernelization procedure* (shortly a *kernel*) for a parameterized decision problem Π is a polynomial-time algorithm that given an instance (I, k) of Π (k being the parameter), outputs an equivalent instance (I', k') with both k' and size of I' bounded by some computable function of k. If this function is a polynomial in k, we say that the kernel is *polynomial*. The search for (polynomial) kernels is already an established research aera in the field of parameterized algorithms, cf. the monograph of Fomin et al. [FLSZ19].

One particularly important method in the design of kernelization algorithms is *protrusion replacement*. Intuitively, the idea is to find a large *protrusion*: a piece of graph that is "simple" (well structured and easily resolvable, e.g. of bounded treewidth) and has an easily comprehensible interaction with the rest of the graph. After finding a protrusion, we may quickly understand and encode its nature with respect to the problem and replace it with a smaller protrusion with essentially the same functionality (thus reaching an equivalent instance of the problem). Protrusion-based techniques were originally introduced by Bodlaender et al. [BFL+16], but by now have become a part of the standard toolbox of kernelization. We refer the interested reader to [FLSZ19, Part 2] for more information.

A particularly important achievement in the development of protrusion-based kernelization procedures is the result of Fomin et al. [FLMS12], who gave a polynomial kernel for the Planar $\mathcal{F}$-Deletion problem. This problem asks, for a fixed family $\mathcal{F}$ of graphs containing at least one planar graph, if from a graph G one can remove at most k vertices to get rid of all minor models of graphs from $\mathcal{F}$. Fomin et al. gave a polynomial kernel for this problem for every fixed family $\mathcal{F}$ as above. Giannopolou et al. [GPR+21] were able to apply the concept of protrusion in the setting of immersions. They considered the $\mathcal{F}$-Immersion Deletion problem, where one wishes to hit all immersion models of graphs from $\mathcal{F}$ using a hitting set of edges of size at most k. By loosely following the approach of [FLMS12], Giannopolou et al. [GPR+21] gave a *linear* kernel for $\mathcal{F}$-Immersion Deletion for every family $\mathcal{F}$ that contains a subcubic planar graph. Here, the main idea was to adjust the notions of protrusions to the graph parameter *tree-cutwidth* and corresponding *tree-cut decompositions*, which play the same role for immersions as treewidth and tree decompositions play for minors.

Again the natural question one could ask is whether the understanding of the undirected world can be used in order to design corresponding kernels for directed graphs.

Organization of the thesis

The dissertation is divided into three parts, each corresponding to a research project held in a different group, which led to a publication (or a series of publications). Each chapter corresponds to a different paper, most of which have already been published. The content of chapters is taken (almost) verbatim from the corresponding papers.

Part I is devoted to development of structural theory for tournaments with respect to embedding notions of immersions and topological minors, and is based on works:

[BP22a] Łukasz Bożyk and Michał Pilipczuk. On the Erődś-Pósa property for immersions and topological minors in tournaments. *Discrete Mathematics & Theoretical Computer Science*, 24(1), 2022;

[BP22b] Łukasz Bożyk and Michał Pilipczuk. Polynomial kernel for immersion hitting in tournaments. In *30th Annual European Symposium on Algorithms (ESA 2022)*, volume 244 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:17, 2022.

Part II concerns general (directed and undirected) graphs and studies obstructions to immersion-based properties. This is based on works:

[BDOP22a] Łukasz Bożyk, Oscar Defrain, Karolina Okrasa, and Michał Pilipczuk. On objects dual to tree-cut decompositions. *Journal of Combinatorial Theory, Series B*, 157:401–428, 2022;

[BDOP22b] Łukasz Bożyk, Oscar Defrain, Karolina Okrasa, and Michał Pilipczuk. On digraphs without onion star immersions, 2022 (preprint, submitted).

Finally Part III is devoted to results following from a joint work, in which we study the parameterized complexity of the bipartite permutation vertex deletion problem:

[BDK$^+$22] Łukasz Bożyk, Jan Derbisz, Tomasz Krawczyk, Jana Novotná, and Karolina Okrasa. Vertex deletion into bipartite permutation graphs. *Algorithmica*, 84(8):2271–2291, 2022.

The Erdős-Pósa Property for Immersions and Topological Minors in Tournaments

In Chapter 3 we consider the Erdős-Pósa property for immersions and topological minors in tournaments. The following two definitions formally introduce the properties we are interested in.

Definition. *A directed graph H has the* Erdős-Pósa property for immersions in tournaments *if there is a function $f\colon \mathbb{N} \to \mathbb{N}$, called further a* bounding function, *such that for every $k \in \mathbb{N}$ and every tournament T, at least one of the following holds:*
- *T contains k pairwise arc-disjoint immersion copies of H; or*
- *there exists a set of at most $f(k)$ arcs of T that intersects all immersion copies of H in T.*

Definition. *A directed graph H has the* Erdős-Pósa property for topological minors in tournaments *if there is a function $f\colon \mathbb{N} \to \mathbb{N}$, called further a* bounding function, *such that for every $k \in \mathbb{N}$ and every tournament T, at least one of the following holds:*
- *T contains k pairwise vertex-disjoint topological minor copies of H; or*
- *there exists a set of at most $f(k)$ vertices of T that intersects all topological minor copies of H in T.*

These two properties were investigated by Raymond [Ray18], who proved that as long as H is *simple* — there are no multiple arcs with the same head and tail — and *strongly connected* — for every pair of vertices u, v, there are directed paths both from u to v and from v to u — the considered Erdős-Pósa properties hold (Theorems 2 and 3 of [Ray18]).

Raymond asked in [Ray18] whether the assumption that the digraph in question is strongly connected can be dropped, as it was important in his proof. We answer this question in affirmative by proving the following.

Theorem A. *Every simple directed graph H has the Erdős-Pósa property for immersions in tournaments with bounding function $f(k) \in \mathcal{O}_H(k^3)$.*

Theorem B. *Every simple directed graph H has the Erdős-Pósa property for topological minors in tournaments with bounding function $f(k) \in \mathcal{O}_H(k \log k)$.*

Compared to the results reported by Raymond in [Ray18], we also give explicit upper bounds on the bounding function that are polynomial in k: cubic for immersions and near-linear for topological minors. The presentation of [Ray18] relies on qualitative results of Chudnovsky et al. [COS12] and of Fradkin and Seymour [OS13]. These results respectively say the following: If a tournament T excludes a fixed digraph H as a immersion (respectively, as a topological minor), then the *cutwidth* of T (respectively, *pathwidth*) of T is bounded by a constant c_H that depends only on H. Instead of relying on the results of [COS12, OS13], we point out that we can use their quantitative improvements of Fomin and Pilipczuk [FP19], and thus obtain concrete bounds on the bounding function that are polynomial in k.

However, the bulk of our work concerns treating directed graphs H that are possibly not strongly connected. Similarly to Raymond [Ray18], using the results of [FP19] we may restrict attention to tournaments of bounded cutwidth or pathwidth, which in both cases provides us with a suitable linear "layout" of the tournament. Then we analyze how an immersion or a topological minor copy of H can look in this layout, and in particular how the strongly connected components of H are ordered by it. The main point is to focus on every topological ordering of the strongly connected components of H separately. Namely, we show that for a given topological ordering π, we can either find k disjoint copies of H respecting this ordering in the layout, or uncover a small hitting set for all copies respecting π. Then taking the union of the hitting sets for all topological orderings π finishes the proof.

We do not expect the estimates on the bounding function given by Theorems A and B to be optimal. In fact, on the way to proving Theorem A we establish an improved bound of $\mathcal{O}_H(k^2)$ under the assumption that H is strongly connected, which suggests that the same asymptotic bound (i.e. quadratic instead of cubic) should also hold without this assumption. However, to the best of our knowledge, in both cases it could even be that the optimal bounding function is linear in k. Finding tighter estimates is an interesting open question.

Polynomial Kernel for Immersion Hitting in Tournaments

The goal of Chapter 4 is to explore the applicability of the structural theory of tournaments for kernelization, with a particular focus of developing a sound protrusion-based methodology. For a simple directed graph H without isolated vertices we define the following parameterized problem:

H-HITTING IMMERSIONS IN TOURNAMENTS

Input: A tournament T and a positive integer k.
Parameter: k
Output: Is there a set $F \subseteq A(T)$, such that $|F| \leqslant k$ and $T - F$ is H-immersion-free?

That this problem is fixed-parameter tractable is proved in [FP19]. Our main result states that for every fixed H as above, H-HITTING IMMERSIONS IN TOURNAMENTS admits a polynomial kernel, of degree dependent on H. Formally, we prove the following theorem.

Theorem C. *For every simple digraph H without isolated vertices there exists a constant c and an algorithm that given an instance (T, k) of H-*HITTING IMMERSIONS IN TOURNAMENTS*, runs in polynomial time and returns an equivalent instance (T', k) with $|T'| \leqslant c \cdot k^c$.*

On a very high conceptual level, the proof of Theorem C follows the classic blueprint of protrusion-based kernelization, like in e.g. [FLMS12, GPR+21]. That is, if (T, k) is a given instance of H-HITTING IMMERSIONS IN TOURNAMENTS, we perform the following steps.

- We may assume that the cutwidth of T is bounded polynomially in k, for otherwise in T one can find $k + 1$ arc-disjoint immersion models of H; these witness a negative answer to the instance.
- Assuming that T is large — of size superpolynomial in k — but has cutwidth bounded polynomially in k, we may find in T a large *protrusion*. Here, a protrusion is an interval I in the vertex ordering σ witnessing small cutwidth such that σ restricted to I witnesses that $T[I]$ has constant cutwidth, and there is only a constant number of σ-backward arcs with one endpoint in I and the other outside of I. These are instantiations of the two desired properties of a protrusion: it has to have bounded width and communicates with the rest of the graph through a boundary of bounded size.

- We can replace the obtained protrusion with a strictly smaller one of the same "type", thus obtaining a strictly smaller equivalent instance. Applying this strategy exhaustively eventually yields a kernel of polynomial size.

Compared to the previous works, the main difficulty is to tame the interaction between a protrusion and the remainder of the instance. Namely, this interaction is not restricted to a set of vertices or arcs of constant size: as we work with tournaments, every vertex of a protrusion will necessarily have an arc connecting it to every vertex outside of the protrusion. The idea is that all but a constant number of those arcs will be forward arcs in the fixed vertex ordering σ with bounded cutwidth. We call those well-behaved forward arcs *generic*, while the remaining constantly many backward arcs are *singular*. Understanding the interaction between a protrusion and the rest of the tournament as being governed by few singular arcs and a large number of well-behaved generic arcs is the crux of our approach.

In particular, while looking for a large replaceable protrusion, we have to be extremely careful when arguing about how such a protrusion may interact with optimum solutions. Here, a key step is to find several protrusions that appear consecutively in σ (recall that our protrusions are intervals in σ), have the same *type* (in the sense of admitting partial immersions of H), and such that their union is a protrusion of again the same type. This step is done using Simon Factorization, a tool commonly used in the theory of automata and formal languages. Simon Factorization was recently used a few times in structural graph theory [BP20, NOP+21, NORS21, JMPW23], but we are not aware of any previous application in the context of kernelization.

The application of Simon Factorization is also the only step in the reasoning that causes the degree of the polynomial bounding the size of our kernel to depend on H. It is an interesting open question whether this can be improved, or in other words, whether there is a kernel of size at most $c \cdot k^d$, where c may depend on H but d does not. Judging by the results on hitting immersions in undirected graphs [GPR+21], we expect that this might be the case.

Objects Dual to Tree-cut Decompositions

In Chapter 5 we are in fact working with *multigraphs*, (so in the statements of our results one allows graphs to have parallel edges, but not loops).

We investigate *tree-cut width* — a graph parameter introduced by Wollan [Wol15] which is an analogue of treewidth for the immersion order on graphs in the following sense: the tree-cut width of a graph is functionally equivalent to the largest size of a wall that can be found in it as an immersion.

We believe that despite extensive work on tree-cut width and tree-cut decompositions, the currently used definitions of width — due to Wollan [Wol15] and to Giannopoulou et al. [GPR+21] — still seem somewhat unnatural, and hence not completely understood. The main motivation is to find a definition for tree-cut width that would feel more "right".

In this chapter we propose a different (but functionally equivalent) measure of the width of a tree-cut decomposition, and present natural analogues of brambles and tangles (corresponding to edge separations rather than vertex separations) for which we can state and prove the following result:

Theorem D. *For a graph G and a positive integer k, the following conditions are equivalent:*
- *G has no tree-cut decomposition of width $< k$;*
- *G has a bramble of order $\geqslant k$; and*
- *G has a tangle of order $\geqslant k$.*

An important aspect that emerged during our work is that tree-cut decompositions naturally have not one, but two orthogonal width measures. The *bag-width* is defined as the maximum size of a bag, and the *adhesion-width* (roughly) governs the sizes of adhesions. Mirroring this, brambles and tangles have two measures of order: *bag-order* and *adhesion-order*. We can prove the tight duality result also in this biparametric setting:

Theorem E. *For a graph G and positive integers a and b, the following conditions are equivalent:*
- *G has no tree-cut decomposition of adhesion-width $< a$ and bag-width $< b$;*
- *G has a bramble of adhesion-order $\geqslant a$ and bag-order $\geqslant b$; and*
- *G has a tangle of adhesion-order $\geqslant a$ and bag-order $\geqslant b$.*

In the hindsight, the biparametric aspect of tree-cut decompositions is visible in the previous definitions [Wol15, GPR+21], but the two parameters were combined into a single width measure in a somewhat arbitrary way. We believe that handling the two parameters separately actually clarifies the situation.

Similarly to the setting of treewidth, we can use our duality theorem to give a characterization of tree-cut width in terms of a search game that we call *cops, dogs, and robber game*. The game is played similarly to the standard cops and robber game, with the exception that the a cops are now always placed on edges of the graph, instead of vertices. Up to technicalities, the robber is always placed on a vertex and may move freely along paths that avoid edges occupied by cops. Obviously, cops placed on edges cannot directly catch the robber placed on a vertex, but they also have b dogs for this purpose. Namely, once the size of the set of vertices to which the robber can possibly move is limited to $\leqslant b$, the cops can unleash the dogs on those vertices and immediately catch the robber. We prove the following.

Theorem F. *A graph has a tree-cut decomposition with adhesion-width $\leqslant a$ and bag-width $\leqslant b$ if and only if there is a strategy to catch the robber using a cops and b dogs.*

Our proof strategy closely follows the standard way used for the duality theorem for treewidth, but adjusted to tree-cut decompositions [ST93, Maz13, Die16]. Once all the definitions are rightly set, the standard and well-understood proof strategy applies without any problems, providing a reasoning that is conceptually even simpler than that for the treewidth.

Digraphs without Onion Star Immersions

In this chapter (and in all the statements) we are working with *multidigraphs* — we are allowing parallel arcs of the same head and tail (but disallowing loops).

The aim of Chapter 6 is to make the first modest steps towards a structure theory for digraphs excluding a fixed digraph as an immersion. A more specific motivation is to provide opening moves towards a statement linking directed wall immersions with suitable width measures for digraphs.

We consider the t*-onion star*: the digraph depicted in Figure 1.2. It is obtained from the star with $2t$ leaves by replacing every edge with a triple of arcs; for t edges two of the arcs are oriented towards the center and one away from the center, and for the remaining t edges we use the reverse orientation. Onion star has the following property.

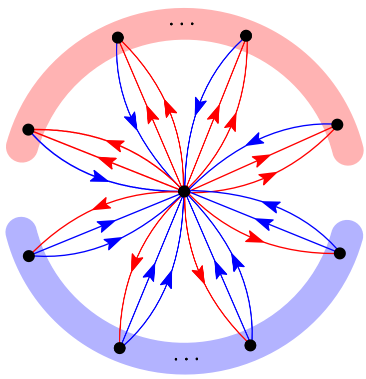

Figure 1.2: The onion star.

Observation 1.1. *Suppose D is a digraph on t vertices where every vertex has outdegree at most 2 and indegree at most 1, or vice versa. Then the t-onion star contains D as an immersion.*

Therefore, excluding the t-onion star as an immersion, for some fixed $t \in \mathbb{N}$, is a weaker condition than excluding any fixed digraph D satisfying the premise of Observation 1.1; this in particular applies to any reasonable notion of a directed wall. The t-onion star is a directed analogue of the graph $S_{3,t}$ that was used by Wollan in [Wol15] as an obstruction commonly found in various avenues of his proof of the (undirected) Wall Immersion Theorem.

For two vertices x, y in a digraph D, let $\mu(x, y)$ be the maximum number of arc-disjoint paths from x to y that one can find in D. As an opening step of his proof, Wollan proved the following statement.

Theorem 1.2 (Wollan, [Wol15, Lemma 1]). *Suppose a graph G contains a set X consisting of $t+1$ vertices such that for all distinct $x, y \in X$, we have $\mu(x, y) \geqslant t^2$. Then G contains the complete graph K_t as an immersion.*

The proof is a relatively easy application of flow-cut duality, which nonetheless crucially relies on the undirectedness of the graph. The main result of this chapter is the following weak directed analogue of Theorem 1.2.

Theorem G. *There exists a function $f\colon \mathbb{N} \to \mathbb{N}$ such that the following holds. Suppose a digraph D contains a set X consisting of $2t+1$ vertices such that for all distinct $x, y \in X$, we have $\mu(x, y) > f(t)$. Then D contains the t-onion star as an immersion.*

Note that in Theorem G we obtain an obstruction that is weaker than that of Wollan's: the t-onion star instead of a complete digraph.

The proof of Theorem G applies the same basic flow-cut strategy as the proof of Theorem 1.2 due to Wollan, but there is a major issue. For any $x, y \in X$, the premise of Theorem G provides the existence of a large family $\mathcal{P}_{x\to y}$ of arc-disjoint paths from x to y, and of a large family $\mathcal{P}_{y\to x}$ of arc-disjoint paths from y to x. However, in principle every path of $\mathcal{P}_{x\to y}$ could intersect every path of $\mathcal{P}_{y\to x}$, while for the aforementioned strategy to work, we need a family containing many arc-disjoint paths from x to y and many from y to x that are also arc-disjoint between each other. In general digraphs, there is no hope for exposing such a family: in Figure 1.3 we give an example where $\mu(x, y)$ and $\mu(y, x)$ can be arbitrarily large, but one cannot find even two arc-disjoint paths: one from x to y and one from y to x. However, we prove that the desired statement holds under the assumption of excluding an onion star.

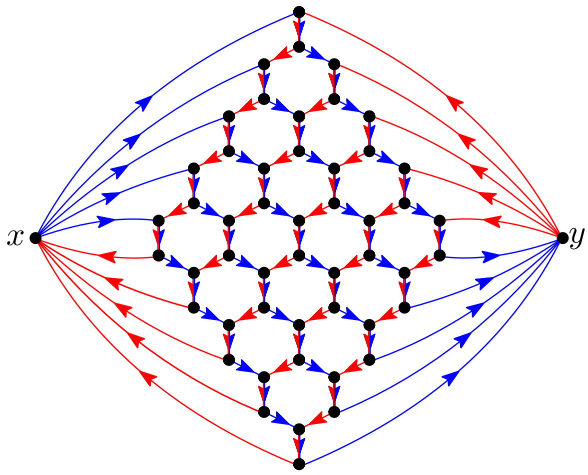

Figure 1.3: An example of a digraph that contains a large family of arc-disjoint paths from x to y (in blue), a large family of arc-disjoint paths from y to x (in red), but no two arc disjoint paths such that one goes from x to y, and the other one from y to x.

Theorem H. *There exists a function $g\colon \mathbb{N}\times\mathbb{N} \to \mathbb{N}$ such that the following holds. Suppose D is a digraph and x, y are two distinct vertices in D such that $\mu(x, y) > g(t, k)$ and $\mu(y, x) > g(t, k)$, for some $t, k \in \mathbb{N}$. Then at least one of the following holds:*

- *D contains the t-onion star as an immersion;*
- *in D there is a family of $2k$ arc-disjoint paths consisting of k paths from x to y and k paths from y to x.*

Theorem H is the main technical component in the proof of Theorem G, and its proof spans most of the chapter.

Vertex Deletion into Bipartite Permutation Graphs

In Chapter 7 we study the parameterized complexity of the bipartite permutation vertex deletion problem, which asks, for a given n-vertex graph, whether we can remove at most k vertices to obtain a bipartite permutation graph. This problem is NP-complete by the classical result of Lewis and Yannakakis [LY80]. We prove the following.

Theorem I. *There is an $\mathcal{O}(9^k \cdot |G|^9)$-time algorithm for instances (G, k) of the vertex deletion into bipartite permutation graphs problem.*

Our algorithm is based on the characterization of bipartite permutation graphs by forbidden subgraphs. Using the characterization, at first, we get rid of forbidden induced subgraphs on at most 9 vertices by branching, which is a standard technique in modification problems on hereditary graph classes [vtHV13, VHPT09]. We call graphs without these forbidden subgraphs *almost bipartite permutation graphs*.

Our main contribution is in the structural analysis of almost bipartite permutation graphs which may contain holes (on more than ten vertices) in contrast to bipartite permutation graphs. This approach is partially inspired by the ideas of van 't Hof and Villanger [vtHV13] who used similar tools in their work on proper interval vertex deletion problem. We use the result of Spinrad, Brandstädt, and Stewart [SBS87], who showed that the vertices of every connected bipartite permutation graph $G = (U, W, E)$ can be embedded into a strip in such a way that the vertices from U are on the bottom edge of the strip, the vertices from W are on the top edge of the strip, the neighbors $N(u)$ of u occur consecutively on the top edge of the strip for every $u \in U$ (*adjacency property*), the vertices from $N(u) - N(u')$ occur consecutively on the top edge of the strip for every $u, u' \in U$ (*enclosure property*), and the analogous properties are satisfied by the vertices in W (see Figure 1.4).

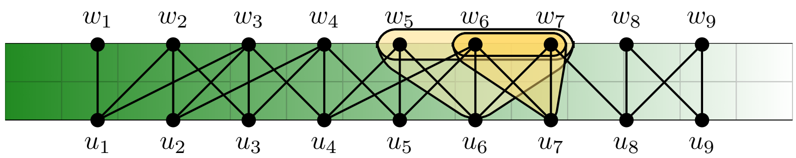

Figure 1.4: Embedding of a bipartite permutation graph (U, W, E) into a strip satisfying the adjacency and the enclosure properties.

Our structural result asserts that, depending on the parity of the length of the shortest hole, a connected almost bipartite permutation graph may be naturally embedded in either a cylinder, or a Möbius strip, locally satisfying adjacency and enclosure properties (see Figure 1.5).

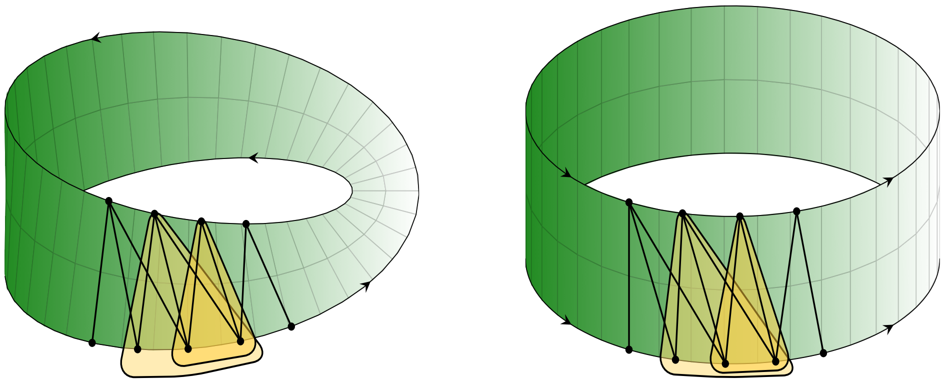

Figure 1.5: An embedding of a connected almost bipartite permutation graph in a cylinder or a Möbius strip that locally satisfies the adjacency and enclosure properties.

Once we obtain such structure, we show that every minimal vertex cut that destroys all holes lies near a few consecutive vertices from the shortest hole. This allows us to check all the possibilities where we can find a minimum cut. Finally, we use a polynomial algorithm for finding maximum flow (and thus a minimum cut).

The approach used to prove Theorem I can be slightly modified to obtain a 9-approximation algorithm for the bipartite permutation vertex deletion problem. We show the following.

Theorem J. *There exists a polynomial-time 9-approximation algorithm for vertex deletion into bipartite permutation graphs problem.*

Contribution statements

As already stated, the conceptual content presented herein is composed of parts of joint publications with no significant changes, and all the main results are fruits of joint research.

In most of the cases it is very difficult to clearly separate the contributions of individual team members, but in all of them the author has played an active role in proving main theorems. Below follow the declarations on the author's roles and activities in the research projects underlying respective publications:

- *Erdős-Pósa Property for Immersions and Topological Minors in Tournaments* (Chapter 3/[BP22a]): developement of the main conceptual idea behind proofs of Theorems A and B (i.e. based on topological orderings), figuring out the way of lifting the proof from undirected minor setting to the setting of directed embeddings, preparation of most of the write-up.
- *Polynomial Kernel for Immersion Hitting in Tournaments* (Chapter 4/[BP22b]): active participation in discussions leading to design of new tools allowing for lifting the protrusion replacement technique to the setting of tournaments (especially Section 4.1), preparation of most of the write-up, including design and development of the definitions and statements used to formalize the main result (Theorem C), delivery of the talk presenting the results at ESA 2022.
- *Objects Dual to Tree-cut Decompositions* (Chapter 5/[BDOP22a]): active participation in regular meetings and discussions, especially about the nature and role of tangles, contribution to the write-up and reviewing the final version of the paper.
- *Digraphs without Onion Star Immersions* (Chapter 6/[BDOP22b]): active participation in regular meetings and discussions, participation in the process of formalization of the results, writing (in particular in the preparation of results presented in Section 6.2 an working out the use of appropriate Ramsey-type theorems) and reviewing (in particular identification of numerous flaws in former versions of the main Theorem H).
- *Vertex Deletion into Bipartite Permutation Graphs* (Chapter 7/[BDK$^+$22]): participation in working out the conceptually crucial structure around a shortest hole (Section 7.2), contribution to the write-up and reviewing, delivery of the talk presenting the results at IPEC 2020.

Part I

Immersions in Tournaments

Chapter 2

Preliminaries on Tournaments

In this chapter we introduce basic notions and definitions required in Chapters 3 and 4, in particular concerning tournaments, their width measures and oriented types of graph containment.

Tournaments. A simple digraph $T = (V, A)$ is called a *tournament* if for every pair of distinct vertices $u, v \in V$, either $(u, v) \in A$, or $(v, u) \in A$ (but not both). Alternatively, one can represent the tournament T by providing a pair $(\sigma, \overleftarrow{A}_\sigma(T))$, where $\sigma\colon V \to [|V|]$ is an *ordering* of the set V and $\overleftarrow{A}_\sigma(T)$ is the set of *σ-backward arcs*, that is,

$$\overleftarrow{A}_\sigma(T) := \{\, (u, v) \in A \;\mid\; \sigma(u) > \sigma(v) \,\}.$$

All the remaining arcs are called *σ-forward.* If the choice of ordering σ is clear from the context, we will call the arcs simply *backward* or *forward.* For $\alpha, \beta \in \{0, 1, \ldots, |V|\}$, $\alpha \leqslant \beta$, we define

$$\sigma(\alpha, \beta] := \{v \in V \mid \alpha < \sigma(v) \leqslant \beta\}.$$

Sets $\sigma(\alpha, \beta]$ as defined above will be called *σ-intervals* (or simply *intervals*, if σ is clear from the context).

Cutwidth. Let $T = (V, A)$ be a tournament and σ be an ordering of V. If $I = \sigma[\alpha, \beta]$, we denote

$$\mathsf{start}_\sigma(I) := \alpha \qquad \text{and} \qquad \mathsf{end}_\sigma(I) := \beta.$$

Moreover, let $\sigma[\alpha] := \sigma(0, \alpha]$ and call this interval an *α-prefix* of σ. The set

$$\mathsf{cut}_\sigma[\alpha] = \{(u, v) \in A \mid \sigma(u) > \alpha \geqslant \sigma(v)\} \subseteq \overleftarrow{A}_\sigma(T)$$

is called the *α-cut* of σ. The *width* of the ordering σ is equal to $\max_{0 \leqslant \alpha \leqslant |V|} |\mathsf{cut}[\alpha]|$, and the *cutwidth* of T, denoted $\mathbf{ctw}(T)$, is the minimum width among all orderings of V.

It is perhaps a bit surprising that in tournaments, there is a very simple algorithm to compute an ordering of optimum width: just sort the vertices by outdegrees.

Lemma 2.1 (see [BPP18, Ove11]). *Let T be a tournament and σ be an ordering of T satisfying the following for every pair of vertices u and v: if u appears before v in σ, then the outdegree of u is not smaller than the outdegree of v. Then the width of σ is equal to $\mathbf{ctw}(T)$.*

If $I = \sigma(\alpha, \beta]$, then we denote

$$\partial^+(I) := \overrightarrow{A}(I, \sigma(0, \alpha]) \quad \text{and} \quad \partial^-(I) := \overrightarrow{A}(\sigma(\beta, |V|], I).$$

Note that $\partial^+(I) \subseteq \mathsf{cut}_\sigma[\alpha]$ and $\partial^-(I) \subseteq \mathsf{cut}_\sigma[\beta]$ and therefore $|\partial^+(I)| \leqslant c$ and $|\partial^-(I)| \leqslant c$, where c is the width of σ. These inclusions may be proper, as the arcs from the set $\widehat{\partial}(I) := \overrightarrow{A}(\sigma(\beta, |V|], \sigma(0, \alpha])$ contribute to the cuts but are not incident with I. We define $\partial(I) := \partial^+(I) \cup \partial^-(I)$ and call the elements of $\partial(I)$ *I-singular* (or simply *singular*) arcs. Moreover, we define

$$\Gamma^+(I) := \overrightarrow{A}(I, \sigma(\beta, |V|]), \qquad \Gamma^-(I) := \overrightarrow{A}(\sigma(0, \alpha], I), \quad \text{and} \quad \Gamma(I) = \Gamma^+(I) \cup \Gamma^-(I),$$

and call the elements of $\Gamma(I)$ *I-generic* (or simply *generic*) arcs.

If $I' = V - I$ where $I = \sigma(\alpha, \beta]$, then we call the set I' a *co-interval* and define *I'-singular* and *I'-generic* arcs as follows

$$\partial^-(I') := \partial^+(I), \quad \partial^+(I') := \partial^-(I), \quad \partial(I') := \partial^+(I') \cup \partial^-(I') = \partial(I),$$
$$\Gamma^-(I') := \Gamma^+(I), \quad \Gamma^+(I') := \Gamma^-(I), \quad \Gamma(I') := \Gamma^-(I') \cup \Gamma^+(I') = \Gamma(I).$$

Immersions. A digraph $\widehat{H}$ is an *immersion model* (or briefly a *copy*) of a digraph H if there exists a mapping ϕ, called an *immersion embedding*, such that:

- vertices of H are mapped to pairwise different vertices of $\widehat{H}$;
- each arc $(u, v) \in A(H)$ is mapped to a directed path in $\widehat{H}$ starting at $\phi(u)$ and ending at $\phi(v)$; and
- each arc of $\widehat{H}$ belongs to exactly one of the paths $\{\phi(a) : a \in A(H)\}$.

If it does not lead to misunderstanding, we will sometimes slightly abuse the above notation by identifying ϕ and $\widehat{H}$ and calling ϕ the immersion model of H. If the immersion embedding ϕ is clear for the context, then for a subgraph C of H we define $\widehat{H}|_C$ to be the subgraph of $\widehat{H}$ consisting of all the vertices and arcs participating in the image of C under ϕ. Note that thus, $\widehat{H}|_C$ is an immersion model of C.

Let H be a digraph. We say that a digraph G *contains H as an immersion* (or H can be *immersed* in G) if G has a subgraph that is an immersion model of H. It is clear that G admits H as an immersion if and only if there exists an immersion model of H in G in which all arcs are mapped to simple paths. Digraph G is called *H-immersion-free* (or *H-free* for brevity) if it does not contain H as an immersion.

We will use the following result of Fomin and Pilipczuk [FP19].

Theorem 2.2 (Theorem 7.3 of [FP19]). *Let T be a tournament which does not contain a digraph H as an immersion. Then $\mathbf{ctw}(T) \in \mathcal{O}(\|H\|^2)$.*

From Theorem 2.2 we can derive the following statements.

Corollary 2.3. *For every digraph H there exists a constant c_H such that for every H-free tournament T, we have $\mathbf{ctw}(T) \leqslant c_H$.*

Corollary 2.4. *Let T be a tournament which does not contain k arc-disjoint immersion copies of a digraph H. Then $\mathbf{ctw}(T) \in \mathcal{O}(\|H\|^2 k^2)$.*

Proof. Let D be the digraph obtained by taking k vertex-disjoint copies of H. Clearly, T does not contain D as an immersion, hence from Theorem 2.2 we conclude that $\mathbf{ctw}(T) \in \mathcal{O}(\|D\|^2) = \mathcal{O}(\|H\|^2 k^2)$. $\square$

Pathwidth. Denote by $\mathcal{I}$ the set of all nonempty intervals $[\alpha, \beta] \subseteq \mathbb{R}$ such that $\alpha, \beta \in \mathbb{Z}$. If $I = [\alpha, \beta]$, denote $\mathsf{first}(I) := \alpha$ and $\mathsf{last}(I) := \beta$. For $I, J \in \mathcal{I}$ we will write $I < J$ if and only if $\mathsf{last}(I) < \mathsf{first}(J)$.

For a tournament $T = (V, A)$, a function $I : V \to \mathcal{I}$ is called an *interval decomposition* of T if for every pair of vertices $u, v \in V$ such that $I(u) < I(v)$, we have $(u, v) \in A$. In other words, every arc joining disjoint intervals is *forward*. For $\alpha \in \mathbb{Z}$, the set

$$\mathsf{vcut}[\alpha] := \{v \in V \mid \alpha \in I(v)\}$$

is called the *α-cut* of I. The *width* of the decomposition I is equal to $\max_{\alpha \in \mathbb{Z}} |\mathsf{vcut}[\alpha]|$, and the *pathwidth* of T, denoted $\mathbf{pw}(T)$, is the minimum width among all interval decompositions of T.

Let us remark here that the definition of pathwidth used in [FP19] is seemingly somewhat different to the one delivered above: it is based on a notion of a *path decomposition*, which is a sequence of *bags* that correspond to sets $\{\mathsf{vcut}[\alpha] : \alpha \in \mathbb{Z}\}$ in an interval decomposition. However, it is straightforward to verify that the definitions are in fact equivalent.

Also, it is easy to see that given an interval decomposition I of a tournament T, one can adjust I to an interval decomposition I' of the same width where no two intervals share an endpoint and no interval has length 0. Indeed, whenever a subset of intervals all have endpoints at $\alpha \in \mathbb{Z}$, then one can shift those endpoints by pairwise different small reals — positive for the intervals ending at α and negative for those starting at α — so that they all become different, and then re-enumerate all the endpoints so that they stay integral. Similarly one can stretch an interval of length 0 which does not share endpoints with any other interval to an interval of positive length. Therefore, we will assume this property for all the considered interval decompositions: $\{\mathsf{first}(I(u)), \mathsf{last}(I(u))\} \cap \{\mathsf{first}(I(v)), \mathsf{last}(I(v))\} = \varnothing$ for all $u \neq v$ and $\mathsf{first}(I(u)) \neq \mathsf{last}(I(u))$ for all u. Moreover, by shifting all the intervals if necessary, we may (and will) assume that all endpoints correspond to non-negative integers.

If I is an interval decomposition of a tournament $T = (V, A)$, then for $\alpha, \beta \in \mathbb{Z}$ we define

$$I[\alpha, \beta] := \{v \in V \mid I(v) \subseteq [\alpha, \beta]\}.$$

In other words, $I[\alpha, \beta]$ is the set of all vertices of T corresponding to intervals entirely contained in $[\alpha, \beta]$. Note that if $\alpha_1 < \beta_1 \leqslant \alpha_2 < \beta_2$, then $I[\alpha_1, \beta_1] \cap I[\alpha_2, \beta_2] = \varnothing$. Also, let $I[\alpha] := I[0, \alpha]$.

Topological minors. Digraph $\widehat{H}$ is a *topological minor model* (or a *topological minor copy*) of a digraph H if there exists a mapping ϕ, called a *topological minor embedding*, such that:

- vertices of H are mapped to pairwise different vertices of $\widehat{H}$;
- each arc $(u, v) \in A(H)$ is mapped to a directed path in $\widehat{H}$ starting at $\phi(u)$ and ending at $\phi(v)$; and
- these paths are internally vertex-disjoint, do not contain any $\phi(u)$, $u \in V(H)$, as an internal vertex, and saturate the whole vertex set and arc set of $\widehat{H}$. In other words, every arc of $\widehat{H}$ and every vertex of $\widehat{H}$ that is not an image of a vertex of H participates in the image $\phi(a)$ of exactly one arc $a \in A(H)$.

If the topological minor embedding ϕ is clear for the context, then for a subgraph C of H we define $\widehat{H}|_C$ to be the subgraph of $\widehat{H}$ consisting of all the vertices and arcs participating in the image of C under ϕ. Note that thus, $\widehat{H}|_C$ is a topological minor model of C.

Let H be a digraph. We say that a digraph G *contains* H *as a topological minor* if G has a subgraph that is a topological minor model of H. Digraph G is called H*-topological-minor-free* if it does not contain H as a topological minor.

We will use another result of Fomin and Pilipczuk.

Theorem 2.5 (Theorem 7.1 of [FP19]). *Let T be a tournament which does not contain a digraph H as a topological minor. Then $\mathbf{pw}(T) \in \mathcal{O}(\|H\|)$.*

Applying Theorem 2.5 directly to the graph that is the disjoint union of k copies of a fixed digraph, we can derive the following statement.

Corollary 2.6. *Let T be a tournament that does not contain k vertex-disjoint topological minor copies of a digraph H. Then $\mathbf{pw}(T) \in \mathcal{O}(\|H\|k)$.*

Monoids and Simon factorization. Simon factorization was originally developed by Simon in [Sim90] and the currently best bounds are due to Kufleitner [Kuf08]. See also the work of Bojańczyk [Boj09] for a nice exposition; we mostly follow the notation from that source.

Let S be a finite monoid (i.e., a finite set equipped with an associative binary operation $\cdot$ and a neutral element 1). An element $e \in S$ is called *idempotent* if $e \cdot e = e$. For a finite alphabet A, by $A^\star$ we denote the set of all finite words over A, and a *morphism* $\alpha\colon A^\star \to S$ is a function satisfying $\alpha(\varepsilon) = 1$ (ε being the empty word) and $\alpha(w_1w_2) = \alpha(w_1) \cdot \alpha(w_2)$ for every $w_1, w_2 \in A^\star$. Note that a morphism is uniquely defined by the images of single symbols from A.

Given an alphabet A and a word $w \in A^\star$, we say that w can be *decomposed* into words w_1, w_2, ..., w_ℓ if $w = w_1w_2 \dots w_\ell$. An unranked rooted tree T labeled with elements of $A^\star$ is a *decomposition tree* of w if:

- the root is labeled with w;
- the label of every non-leaf node can be decomposed into labels of all children of this node (from left to right);
- the label of every leaf node is an element of $A \cup \{\varepsilon\}$ (here ε is the empty word).

The *degree* of a node of a decomposition tree is the number of its sons. Given a monoid S and a morphism $\alpha\colon A^\star \to S$, we say that decomposition $w = w_1w_2 \dots w_\ell$ satisfies the α*-idempotent rule* if there exists an idempotent $e \in S$ such that

$$\alpha(w_1) = \alpha(w_2) = \ldots = \alpha(w_\ell) = e.$$

We define an α*-factorization forest* for a word $w \in A^\star$ as a decomposition tree of w, whose each non-leaf node with degree greater than 2 corresponds to an α-idempotent rule.

Theorem 2.7 (Factorization Forest theorem of Simon). *Let S be a finite monoid, A be a finite alphabet, and $\alpha\colon A^\star \to S$ be a morphism. Then every $w \in A^\star$ has an α-factorization forest of height at most $3|S|$.*

The following lemma is a direct consequence of Simon Factorization.

Lemma 2.8. *Let S be a finite monoid, A be a finite alphabet, and $\alpha\colon A^\star \to S$ be a morphism. Suppose $w \in A^\star$ is a word of length at least $\ell^{3|S|}$. Then there exists a subword w' of w and an idempotent $e \in S$ such that $w' = w_1w_2 \dots w_\ell$, where $w_i \in A^\star$ are nonempty subwords of w and*

$$\alpha(w_1) = \alpha(w_2) = \ldots = \alpha(w_\ell) = e.$$

Proof (of Lemma 2.8). Fix $w \in A^\star$ of length at least $\ell^{3|S|}$. By Theorem 2.7 there exists an α-factorization forest T for w of height at most $3|S|$. Note that if T admits a node of degree at least ℓ, then the idempotent rule corresponding to this node provides the desired subword w' and its decomposition into ℓ non-empty subwords (in the case of degree strictly larger than ℓ, we may concatenate all subwords whose indices are not smaller than ℓ into a single subword).

We will prove that such node always exists. Suppose for the sake of contradiction that all non-leaf nodes of T have degree strictly smaller than ℓ. As T has height at most $3|S|$, the number of leaves of T is therefore strictly smaller than $\ell^{3|S|}$. On the other hand, each letter of w has a different corresponding leaf node of T, so there are at least $|w| \geqslant \ell^{3|S|}$ leaf nodes. The contradiction finishes the proof. $\square$

The intuition behind Lemma 2.8 is as follows. If word w is appropriately long then existence of a constant-height factorization forest forces existence of a node of large degree (hence corresponding to an idempotent rule) in this forest.

Note that in the setting of Lemma 2.8, given a word $w \in A^\star$ of length $n \geqslant \ell^{3|S|}$, one can easily find w' and a suitable decomposition $w' = w_1 w_2 \dots w_\ell$ in time $\mathcal{O}(|S| \cdot n^3)$ assuming unit cost of operations in S. Indeed, one can guess e (by trying at most $|S|$ possibilities) and the first position of w' within w (by trying n possibilities), and then for every subword w'' starting at this position compute the longest possible decomposition of the form $w'' = w_1 w_2 \dots w_{\ell'}$ such that $\alpha(w_1) = \ldots = \alpha(w_{\ell'}) = e$, if existent. The latter can be done by a standard left-to-right dynamic programming in time $\mathcal{O}(n^2)$.

Chapter 3

The Erdős-Pósa Property for Immersions and Topological Minors in Tournaments

This chapter is devoted to proving Theorem A and Theorem B, i.e. that all simple digraphs admit the Erdős-Pósa property for immersions and topological minors in tournaments.

Theorem A (restated). *Every simple directed graph H has the Erdős-Pósa property for immersions in tournaments with bounding function $f(k) \in \mathcal{O}_H(k^3)$.*

Theorem B (restated). *Every simple directed graph H has the Erdős-Pósa property for topological minors in tournaments with bounding function $f(k) \in \mathcal{O}_H(k \log k)$.*

Recall that the above theorems strengthen and improve the results of Raymond [Ray18] by giving better bounding functions and dropping the assumption that H is strongly connected.

3.1 Erdős-Pósa property for immersions

In this section we prove Theorem A. In the following, a subset of arcs F in a digraph D is *H-hitting* if the digraph $D - F$ is H-immersion-free. We also fix the constant d_{ctw} hidden in the $\mathcal{O}(\cdot)$-notation in Theorem 2.2; that is, if a tournament T does not contain H as an immersion then $\mathbf{ctw}(T) \leqslant d_{\mathrm{ctw}} \|H\|^2$. Note that the constant hidden in the $\mathcal{O}(\cdot)$-notation in Corollary 2.4 is also equal to d_{ctw}. Without loss of generality we assume that d_{ctw} is the square of an even integer.

We start with two straightforward observations which will be used several times later on.

Observation 3.1. *Suppose $\widehat{H}$ is an immersion model of a digraph H in a digraph G, and C is a strong component of H. Then there exists a strong component D of G such that $\widehat{H}|_C$ is a subgraph of D.*

Observation 3.2. *Let T be a tournament and σ be an ordering of $V(T)$. Let H be a strongly connected simple digraph with at least one arc and let $\widehat{H}$ be an immersion model of H in T. Let v be the vertex of $V(\widehat{H})$ that is last in the ordering σ. Then $A(\widehat{H})$ contains a σ-backward arc with tail v.*

We now consider two special cases: when H is acyclic and when H is strongly connected. For the acyclic case, we will use the following corollary of the classic results of Erdős and Hanani [EH63].

Lemma 3.3 (follows from [EH63]). *There exists a universal constant d_{eh} such that for all positive integers q, k, in a complete graph on at least $d_{\mathrm{eh}} \cdot q\sqrt{k}$ vertices one can find k pairwise arc-disjoint complete subgraphs, each on q vertices.*

From now on, we adopt the constant d_{eh} in the notation.

Lemma 3.4. *Let H be an acyclic simple digraph and let T be a tournament such that $|T| \geqslant d_{\mathrm{eh}} \cdot 2^{|H|}\sqrt{k}$. Then T contains k arc-disjoint subgraphs isomorphic to H.*

Proof. By Lemma 3.3, in T one can find k arc-disjoint subtournaments $T_1, \ldots, T_k$, each on $2^{|H|}$ vertices. It is well-known that a tournament on $2^{|H|}$ vertices contains a transitive (i.e. acyclic) subtournament on $|H|$ vertices. As H is acyclic, it is a subgraph of a transitive tournament on $|H|$ vertices. Hence, each of $T_1, \ldots, T_k$ contains a subgraph isomorphic to H, and these subgraphs are arc-disjoint. $\square$

Corollary 3.5. *Let H be a simple digraph that is acyclic and let k be a positive integer. Let T be a tournament that does not contain k arc-disjoint immersion copies of H. Then one can find in T a set of at most $d_{\mathrm{eh}}^2 \cdot 4^{|H|}k$ arcs that is H-hitting.*

Proof. We first consider the corner case when H does not contain any arc. Then T must have less than $|H|$ vertices, for otherwise repeating any set of $|H|$ vertices k times would yield k arc-disjoint immersion copies of H. Therefore, T in fact does not contain any immersion copy of H, due to having less vertices, and the empty set is H-hitting in T.

Hence, let us assume that H contains at least one arc. Observe that $|T| < d_{\mathrm{eh}} \cdot 2^{|H|}\sqrt{k}$, for otherwise, by Lemma 3.4, there would exist k arc-disjoint immersion copies of H in T. Since H has at least one arc, the set $A(T)$ of all the arcs of T is H-hitting, and this set has size at most $\binom{d_{\mathrm{eh}} \cdot 2^{|H|}\sqrt{k}}{2} \leqslant d_{\mathrm{eh}}^2 \cdot 4^{|H|}k$, as requested. $\square$

We now move to the case when H is strongly connected. Recall that this case was already considered by Raymond [Ray18], but we give a more refined argument that gives precise upper bounds on the bounding function. The proof relies on an strategy of finding a cut that separates the immersion copies of H in a roughly balanced way, and applying induction to each side of the cut. This strategy has been applied before in the context of Erdős-Pósa properties, see e.g. [FST11].

Lemma 3.6. *Let H be a simple digraph that is strongly connected and contains at least one arc, and let k be a positive integer. Let T be a tournament that does not contain k arc-disjoint immersion copies of H. Then one can find in T a set of at most $6d_{\mathrm{ctw}}\|H\|^2k^2$ arcs that is H-hitting.*

Proof. We prove the lemma by induction on k. For the base case $k = 1$, T does not contain any immersion copy of H, hence the empty set is H-hitting.

Assume then that $k \geqslant 2$. By Corollary 2.4, there is an ordering σ of $V(T)$ of width at most $d_{\mathrm{ctw}}\|H\|^2k^2$. Let $\alpha \in \{0, 1, \ldots, |V(T)|\}$ be the largest index such that the tournament $T[\sigma[\alpha]]$ does not contain $\lceil k/2 \rceil$ arc-disjoint immersion copies of H. Since $\lceil k/2 \rceil < k$, by induction there exists a set of arcs F_1 of size at most $6d_{\mathrm{ctw}}\|H\|^2\lceil k/2 \rceil^2$ that is H-hitting in $T[\sigma[\alpha]]$.

If $\alpha = |V(T)|$, or equivalently $T[\sigma[\alpha]] = T$, then F_1 is in fact H-hitting in T and we are done. Hence, we assume from now on that $\alpha < |V(T)|$. By the maximality of α we know that $T[\sigma[\alpha+1]]$ contains $\lceil k/2 \rceil$ arc-disjoint immersion copies of H. It follows that the tournament $T - \sigma[\alpha + 1]$ does not contain $\lfloor k/2 \rfloor$ arc-disjoint immersion copies of H, for otherwise together we would expose $\lceil k/2 \rceil + \lfloor k/2 \rfloor = k$ arc-disjoint immersion copies of H in T. By induction, there exists a set of arcs F_2 of size at most $6d_{\mathrm{ctw}}\|H\|^2\lfloor k/2 \rfloor^2$ that is H-hitting in $T - \sigma[\alpha + 1]$.

Let now

$$F \coloneqq F_1 \cup F_2 \cup \mathsf{cut}[\alpha] \cup \mathsf{cut}[\alpha + 1].$$

Observe that F is H-hitting in T. Indeed, since H is strongly connected and has at least one arc, every immersion copy of H in T that in not entirely contained in $T[\sigma[\alpha]]$ or $T - \sigma[\alpha + 1]$ is hit by $\mathsf{cut}[\alpha] \cup \mathsf{cut}[\alpha + 1]$, whereas immersion copies entirely contained in $T[\sigma[\alpha]]$ and in $T - \sigma[\alpha + 1]$ are hit by F_1 and F_2, respectively. It remains to estimate the size of F:

$$\begin{aligned}
|F| &\leqslant |F_1| + |F_2| + |\,\mathsf{cut}[\alpha]| + |\,\mathsf{cut}[\alpha + 1]| \\
&\leqslant 6d_{\mathrm{ctw}}\|H\|^2 \left(\lceil k/2 \rceil^2 + \lfloor k/2 \rfloor^2\right) + 2d_{\mathrm{ctw}}\|H\|^2k^2 \\
&\leqslant 6d_{\mathrm{ctw}}\|H\|^2 \left(\left(\frac{k+1}{2}\right)^2 + \left(\frac{k-1}{2}\right)^2\right) + 2d_{\mathrm{ctw}}\|H\|^2k^2 \\
&= d_{\mathrm{ctw}}\|H\|^2 \left(3(k^2+1) + 2k^2\right) \leqslant 6d_{\mathrm{ctw}}\|H\|^2k^2.
\end{aligned}$$

This concludes the inductive proof. $\square$

Actually, in our later proof we will not be able to rely on Lemma 3.6 for the following reason: we will need the copies to be vertex-disjoint, rather than arc-disjoint. The following statement is tailored to vertex-disjointness.

Lemma 3.7. *Let H be a simple digraph that is strongly connected and contains at least one arc, and let k be a positive integer. Further, let T be a tournament with* $\mathbf{ctw}(T) \leqslant c$ *that does not contain k vertex-disjoint immersion copies of H. Then one can find in T a set of at most $2(k-1)c$ arcs that is H-hitting.*

Proof. We proceed by induction on k. Let σ be an ordering of T of width at most c. If T does not contain any copy of H, then the empty set is H-hitting. This proves the base case $k = 1$, so from now on we may assume that $k \geqslant 2$ and that T contains at least one immersion copy of H.

Let α be the minimum integer satisfying the following: $T[\sigma[\alpha]]$ contains an immersion copy $\widehat{H}$ of H. Let $B_1 \coloneqq \{(u,v) \in \overleftarrow{A}_\sigma(T) \mid \sigma(u) = \alpha\}$ be the set of backward arcs with tail α and let $B \coloneqq B_1 \cup \mathsf{cut}[\alpha]$. As $B_1 \subseteq \mathsf{cut}[\alpha - 1]$, we have $|B| \leqslant 2c$.

Observe that in $T' \coloneqq T[V(T) \setminus \sigma[\alpha]]$ one cannot find a family of $k-1$ vertex-disjoint immersion copies of H. Indeed, if there was such a family, then adding $\widehat{H}$ to it would yield a family of k vertex-disjoint copies of H in T, a contradiction. Hence, by induction hypothesis, in T' there is a set S of at most $2(k-2)c$ arcs that is H-hitting. We claim that the set $B \cup S$ is H-hitting in T. Note that since $|B \cup S| \leqslant |B| + |S| \leqslant 2(k-1)c$, this will conclude the proof.

Indeed, suppose that $\widehat{H}'$ is an immersion copy of H in $T - (B \cup S)$. By Observation 3.1, either $V(\widehat{H}') \subseteq V \setminus \sigma[\alpha]$, or $V(\widehat{H}') \subseteq \sigma[\alpha]$. The first case is impossible, because every immersion copy of H in T' contains an arc from S. On the other hand, if $V(\widehat{H}') \subseteq \sigma[\alpha]$, then by the minimality of α we infer that $\sigma^{-1}(\alpha) \in V(\widehat{H}')$. Then Observation 3.2 implies that $\widehat{H}'$ needs to contain an arc of B_1, again a contradiction. $\square$

Note that by combining Lemma 3.7 with Corollary 2.4, we obtain a statement analogous to Lemma 3.6, however with a bound of $\mathcal{O}(k^3)$ instead of $\mathcal{O}(k^2)$. This drawback will accordingly affect the final dependency on k in Theorem A.

We now proceed to the main part of the proof, which concerns digraphs that are not acyclic and that are not necessarily strongly connected.

Lemma 3.8. *Let H be a simple digraph that is not acyclic and let k be a positive integer. Let T be a tournament that does not contain k arc-disjoint immersion copies of H. Then one can find in T a set consisting of at most $2d_{\mathrm{ctw}}^{3/2} \cdot |H|! \cdot |H| \cdot \|H\|^3 \cdot k^3$ arcs that is H-hitting.*

Proof. Let Comps be the family of all strong components of H and let $h \coloneqq |\mathsf{Comps}|$. Since H is not acyclic, Comps contains at least one strong component C that is *non-trivial*, that is, $|C| > 1$. In particular, $h \leqslant |H| - 1$. Further, let Π be the set of all *topological orderings* of the strong components of H; that is, the elements of Π are orderings $\pi\colon \mathsf{Comps} \to [|\mathsf{Comps}|]$ such that for every arc of H with tail in $C \in \mathsf{Comps}$ and head in $D \in \mathsf{Comps}$, we have $\pi(C) \leqslant \pi(D)$. It is well-known that $\Pi \neq \varnothing$. Also, note that $|\Pi| \leqslant h! \leqslant (|H|-1)!$.

Let $T = (V, E)$. By Corollary 2.4, there is an ordering σ of vertices of T of width at most c, where

$$c \coloneqq d_{\mathrm{ctw}} \|H\|^2 k^2.$$

We also define

$$s \coloneqq \sqrt{d_{\mathrm{ctw}}} \cdot h \|H\| k.$$

Note that thus, $s = h\sqrt{c}$ and s is an even integer, because we assume d_{ctw} to be a square of an even integer.

Let $\mathcal{I}$ be the set of all σ-intervals. For $I \in \mathcal{I}$, we define

$$\mathsf{cut}^-(I) \coloneqq \mathsf{cut}[\mathsf{start}_\sigma(I)] \qquad \text{and} \qquad \mathsf{cut}^+(I) \coloneqq \mathsf{cut}[\mathsf{end}_\sigma(I)].$$

We define functions

$$I\colon \mathsf{Comps} \times [|V|] \to \mathcal{I} \qquad \text{and} \qquad A, B\colon \mathsf{Comps} \times [|V|] \to \mathsf{Pow}\left(\overleftarrow{A}_\sigma(T)\right),$$

where $\mathsf{Pow}(X)$ denotes the power set of X, as follows:

- $I(C, \alpha)$ is the inclusion-wise minimal σ-interval I such that $\mathsf{start}_\sigma(I) = \alpha$ and $T[I]$ contains at least s vertex-disjoint immersion copies of C. If no such interval exists, we set $I(C, \alpha) \coloneqq \sigma(\alpha, |V|]$. Note that either way, $T[I]$ does not contain $s + 1$ vertex-disjoint immersion copies of C.

- If C is trivial, then $A(C,\alpha)$ is the set of all backward arcs contained in $T[I(C,\alpha)]$. If C is non-trivial, then $A(C,\alpha)$ is a set of arcs that is C-hitting in $T[I(C,\alpha)]$ and is of size at most $2sc$, whose existence follows from Lemma 3.7.
- $B(C,\alpha) \coloneqq \mathsf{cut}^+(I(C,\alpha))$.

Note that if C is trivial, then $|I(C,\alpha)| \leqslant s$. This implies that $|A(C,\alpha)| \leqslant \binom{s}{2} \leqslant 2sc$. Hence, in all cases we have

$$|A(C,\alpha)| \leqslant 2sc \qquad \text{and} \qquad |B(C,\alpha)| \leqslant c.$$

Consider an arbitrary topological ordering $\pi \in \Pi$. We define indices $\alpha_0, \alpha_1, \ldots, \alpha_h$ and intervals $I_{\pi,1}, I_{\pi,2}, \ldots, I_{\pi,h}$ by induction as follows: $\alpha_0 \coloneqq 0$ and, for $i = 1, 2, \ldots, h$, we set

$$I_{\pi,i} \coloneqq I(\pi^{-1}(i), \alpha_{i-1}) \qquad \text{and} \qquad \alpha_i \coloneqq \mathsf{end}_\sigma(I_{\pi,i}),$$

where if needed we put $\mathsf{end}_\sigma(\varnothing) = |V|$. Moreover, for $i \in [h]$ we define

$$A_{\pi,i} \coloneqq A(\pi^{-1}(i), \alpha_{i-1}) \qquad \text{and} \qquad B_{\pi,i} \coloneqq B(\pi^{-1}(i), \alpha_{i-1}).$$

Our next step is to show that if for some $\pi \in \Pi$, each interval $I_{\pi,i}$ contains s vertex-disjoint immersion copies of H, then we get a contradiction: there are k vertex-disjoint immersion copies of H in T. For this, we will use the following auxiliary statement.

Claim 3.1. *Let G be a graph with vertex set partitioned into disjoint sets $V_1, \ldots, V_h$, each of size s. Suppose that for each pair of indices $1 \leqslant i < j \leqslant h$, there are at most $\frac{s^2}{h^2}$ edges with one endpoint in V_i and second in V_j. Then one can find $s/2$ pairwise disjoint independent sets $I_1, \ldots, I_{s/2}$ in G such that each independent set I_t, $t \in [s/2]$, contains exactly one vertex from each set V_i, $i \in [h]$.*

Proof. For each $i \in [h]$ let us arbitrarily enumerate the vertices of V_i as $v_i[0], \ldots, v_i[s-1]$. Consider the following random experiment: draw independently and uniformly at random numbers $t, a_1, \ldots, a_h$ from $\{0, 1, \ldots, s-1\}$, and let

$$I \coloneqq \{v_i[(t + a_i) \bmod s] \colon i \in [h]\}.$$

Note that for each fixed pair of indices $1 \leqslant i < j \leqslant h$, the probability that there is an edge between vertices $v_i[(t+a_i) \bmod s]$ and $v_j[(t+a_j) \bmod s]$ is bounded by $\frac{1}{h^2}$. By the union bound we infer that I is an independent set with probability at least $\frac{1}{2}$. Hence, there is a choice of $\hat{a}_1, \ldots, \hat{a}_h \in \{0, 1, \ldots, s-1\}$ such that conditioned on $a_1 = \hat{a}_1, \ldots, a_h = \hat{a}_h$, the probability (over the choice of t) that I is an independent set is at least $\frac{1}{2}$. In other words, for at least $s/2$ choices of t, the set $\{v_i[(t + \hat{a}_i) \bmod s] \colon i \in [h]\}$ is independent. This gives us the desired family of $s/2$ pairwise disjoint independent sets. ⌟

Claim 3.2. *Suppose that there exists $\pi \in \Pi$ such that for every $i \in [h]$, the tournament $T[I_{\pi,i}]$ contains s vertex-disjoint immersion copies of $\pi^{-1}(i)$. Then T contains k vertex-disjoint immersion copies of H.*

Proof. Denote $C_{\pi,i} \coloneqq \pi^{-1}(i)$. For each $i \in [h]$, let $\mathcal{C}_i$ be the family of s vertex-disjoint immersion copies of $C_{\pi,i}$ contained in $T[I_{\pi,i}]$.

Let G be a graph on vertex set $\mathcal{C}_1 \cup \ldots \cup \mathcal{C}_p$ where for each pair of indices $1 \leqslant i < j \leqslant h$ and pair of immersion copies $Q \in \mathcal{C}_i$ and $R \in \mathcal{C}_j$, we put an edge if and only if in T there is arc with tail in R and head in Q. Note that such an arc is backward in σ and belongs to $\mathsf{cut}[\alpha_i]$. Hence, for every pair of indices i, j as above, G contains at most c edges with one endpoint in $\mathcal{C}_i$ and second in $\mathcal{C}_j$.

Noting that $c = \frac{s^2}{h^2}$, we may apply Claim 3.1 to conclude that G contains $s/2$ pairwise independent sets, each consisting of one element from each of the families $\mathcal{C}_1, \ldots, \mathcal{C}_h$. As $s/2 \geqslant k$, let $I_1, \ldots, I_k$ be any k of those independent sets. Now, for each $t \in [k]$, we may construct an immersion copy of H contained in $T[\bigcup_{Q \in I_t} V(Q)]$ as follows: take the union of subgraphs $Q \in I_t$, which are immersion copies of $C_{\pi,1}, \ldots, C_{\pi,t}$, respectively, and for each arc (a,b) of H that is not contained in any of $C_{\pi,1}, \ldots, C_{\pi,t}$, say $a \in V(C_{\pi,i})$ and $b \in V(C_{\pi,j})$ where we necessarily have $i < j$, map (a,b) to the single edge between the corresponding two vertices from the copies of $C_{\pi,i}$ and $C_{\pi,j}$ in I_t. Note that this edge is oriented forward in σ, because I_t is an independent set in G (a backward arc would have generated an edge in G). Thus, we have constructed k vertex-disjoint copies of H in T. ⌟

If the assumption of Claim 3.2 holds, then we immediately obtain a contradiction and the proof is finished. Therefore, we may further assume that for every $\pi \in \Pi(H)$ there exists $j \in [h]$ such that $I_{\pi,j}$ contains less than s vertex-disjoint copies of $\pi^{-1}(j)$. Observe that this implies that $\mathsf{end}_\sigma(I_{\pi,j}) = |V|$, hence in particular we have

$$\bigcup_{i=1}^{h} I_{\pi,i} = V(T) \qquad \text{for each } \pi \in \Pi.$$

Let

$$S := \bigcup_{\pi\in\Pi}\bigcup_{i=1}^{h} A_{\pi,i}\cup B_{\pi,i}.$$

Observe that

$$\begin{aligned}|S| &\leqslant |\Pi|\cdot|H|\cdot(2sc+c)\\ &\leqslant (|H|-1)!\cdot|H|\cdot(2s+1)c\\ &= |H|!\cdot(2\sqrt{d_{\mathrm{ctw}}}\cdot h\|H\|k+1)\cdot d_{\mathrm{ctw}}\|H\|^2k^2\\ &\leqslant 2d_{\mathrm{ctw}}^{3/2}\cdot|H|!\cdot|H|\cdot\|H\|^3\cdot k^3,\end{aligned}$$

so to finish the proof it suffices to show that S is H-hitting in T. Let $T' := T - S$.

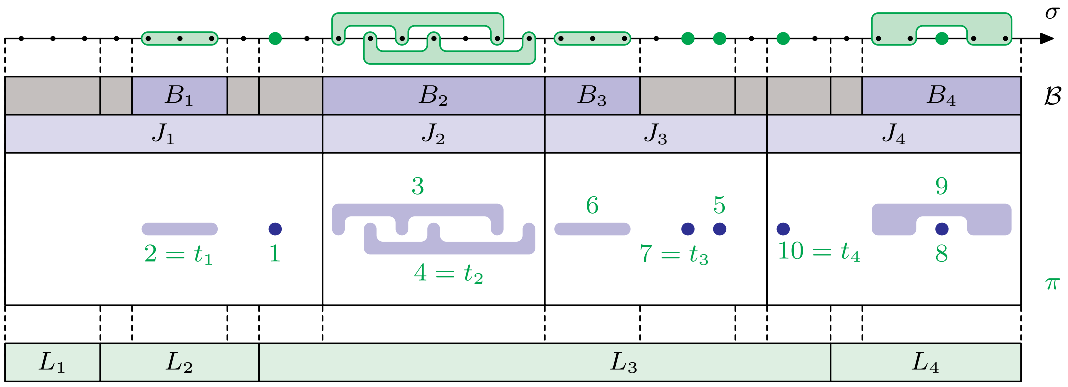

Figure 3.1: Objects defined in the proof of Claim 3.3 with $h = 10$, $m = 4$, $n = 3$.

Claim 3.3. *T' is H-immersion-free.*

Proof. Let $\mathcal{B}$ be the family of all inclusion-wise maximal σ-intervals B satisfying the following property: for every $\pi\in\Pi$ and $i\in[h]$, either $B\subseteq I_{\pi,i}$ or $B\cap I_{\pi,i}=\varnothing$. Call elements of $\mathcal{B}$ *base intervals* and observe that $\mathcal{B}$ is a partition of $V(T')$. Let $\mathcal{F}$ be the family of all σ-intervals which are disjoint unions of collections of base intervals. For two disjoint intervals $J, J'\in\mathcal{F}$, we write $J < J'$ if $\mathsf{end}_\sigma(J)\leqslant\mathsf{start}_\sigma(J')$.

Suppose for contradiction that T' contains an immersion model $\widehat{H}$ of H. We fix some immersion embedding of H in $\widehat{H}$, to which we will implicitly refer when considering subgraphs $\widehat{H}|_C$ for $C\in\mathsf{Comps}$.

Note that in T' there are no backward arcs with endpoints in different intervals from $\mathcal{B}$, as $B_{\pi,i}\subseteq S$ for every $\pi\in\Pi$ and $i\in[h]$. Hence, every non-trivial strongly connected subgraph of T' must have all vertices contained in a single base interval. In particular, from Observation 3.1 we infer that for every non-trivial strong component $C\in\mathsf{Comps}$, the subgraph $\widehat{H}|_C$ has all its vertices contained in a single base interval. Note that this conclusion also holds trivially when C is trivial.

Let $B_1 < B_2 < \ldots < B_m$ be all the base intervals containing subgraphs $\widehat{H}|_C$ for non-trivial components $C\in\mathsf{Comps}$. Consider any partition of $V(T)$ into intervals $J_1, J_2, \ldots, J_m\in\mathcal{F}$ such that $B_i\subseteq J_i$ for each $i\in[m]$. Note that this implies that $J_1 < J_2 < \ldots < J_m$. For each $i\in[m]$, let $\mathsf{Comps}_i\subseteq\mathsf{Comps}$ be the set of all (including trivial) components $C\in\mathsf{Comps}$ such that $V(\widehat{H}|_C)\subseteq J_i$. Note that $\{\mathsf{Comps}_i : i\in[m]\}$ is a partition of Comps and each family Comps_i contains at least one non-trivial component.

Observe that if $C\in\mathsf{Comps}_i$ and $C'\in\mathsf{Comps}_{i'}$, where $i\neq i'$, and in H there is an arc (u,v) with $u\in V(C)$ and $v\in V(C')$, then we necessarily have $i<i'$. Indeed, the image of (u,v) in the immersion embedding is a path in T' that starts in J_i and ends in $J_{i'}$, while in T' arcs with endpoints in different intervals among $\{J_1,\ldots,J_m\}$ always point from an interval with a smaller index to an interval with a higher index. Therefore, there exists a topological ordering $\pi\in\Pi$ such that for all $i,i'\in[m]$ satisfying $i<i'$, all the components of Comps_i appear in π before all the components of $\mathsf{Comps}_{i'}$. In other words, there exist integers $0=t_0<t_1<t_2<\ldots<t_m=h$ such that for each $i\in[m]$, we have $\pi(\mathsf{Comps}_i) = (t_{i-1},t_i]\cap\mathbb{Z}$ (cf. Figure 3.1). For every $i\in[m]$, we define

$$L_i := \bigcup_{C\in\mathsf{Comps}_i} I_{\pi,\pi(C)}.$$

Note that L_i is a σ-interval belonging to $\mathcal{F}$, because the set Comps_i is contiguous in the ordering π. Furthermore $\{L_i : i \in [m]\}$ is a partition of $V(T)$ and $L_1 < L_2 < \ldots < L_m$.

Recalling that both $\{J_i : i \in [m]\}$ and $\{L_i : i \in [m]\}$ are partitions of $V(T)$, we can define n to be the smallest positive integer satisfying $\bigcup_{i=1}^{n} J_i \subseteq \bigcup_{i=1}^{n} L_i$. By the minimality of n, we have $J_n \subseteq L_n$.

Recall that $B_n \subseteq J_n \subseteq L_n$ and B_n is a base interval. Therefore, there exists $C \in \mathsf{Comps}_n$ such that $B_n \subseteq I_{\pi,\pi(C)}$. If C is non-trivial, then the set of arcs $A_{\pi,\pi(C)}$ is C-hitting in $T[I_{\pi,\pi(C)}]$. This implies that $T'[I_{\pi,\pi(C)}]$ is C-immersion-free, and so is its subgraph $T'[B_n]$. However, B_n is the only interval among $\{B_1, \ldots, B_m\}$ that is contained in $I_{\pi,\pi(C)}$, hence C being a non-trivial component from Comps_n implies that $V(\widehat{H}|_C) \subseteq B_n$; a contradiction. If C is trivial, then $T[I_{\pi,\pi(C)}] - A_{\pi,\pi(C)}$ is acyclic, hence $T'[I_{\pi,\pi(C)}]$ is C'-immersion-free for every non-trivial component $C' \in \mathsf{Comps}_n$. Since there exists such a non-trivial component C' and it again satisfies $V(\widehat{H}_{C'}) \subseteq B_n \subseteq I_{\pi,\pi(C)}$, we again obtain a contradiction. ⌟

As argued, Claim 3.3 finishes the proof of Lemma 3.8. □

With Lemma 3.8 in place, we can finish the proof of Theorem A.

Proof (of Theorem A). If H has no arcs, then the statement holds trivially for bounding function $f(k) = 0$. Hence, from now on assume that H has at least one arc. Suppose T is a tournament that does not contain k arc-disjoint immersion copies of H. If H is acyclic, then, by Corollary 3.5, we may find in T a set of at most $d_{\mathrm{eh}}^2 \cdot 4^{|H|} k \in \mathcal{O}_H(k)$ arcs that is H-hitting. On the other hand, if H is not acyclic, then by Lemma 3.8 we may find in T an H-hitting set of arcs of size at most $2d_{\mathrm{ctw}}^{3/2} \cdot |H|! \cdot |H| \cdot \|H\|^3 \cdot k^3 \in \mathcal{O}_H(k^3)$. □

3.2 Erdős-Pósa property for topological minors

In this section we prove Theorem B. The proof follows similar ideas to the ones presented in the previous section, only adjusted to the setting of interval decompositions. Throughout this section, the notions of a *copy* and of *hitting* will refer to topological minor copies. Let us fix the constant d_{pw} hidden in the $\mathcal{O}(\cdot)$-notation in Theorem 2.5 and note that the constant hidden in the $\mathcal{O}(\cdot)$-notation in Corollary 2.6 is also equal to d_{pw}.

Consider first the acyclic case. The following statements are analogues of Lemma 3.4 and Corollary 3.5.

Lemma 3.9. *Let H be an acyclic simple digraph let T be a tournament such that $|T| \geqslant 2^{|H|}k$. Then T contains k vertex-disjoint subgraphs isomorphic to H.*

Proof. Arbitrarily partition the vertex set of T into subsets $W_1, \ldots, W_k$ so that $|W_i| \geqslant 2^{|H|}$ for each $i \in [k]$. Since a tournament on $2^{|H|}$ vertices contains a transitive subtournament on $|H|$ vertices, which in turn contains H as a subgraph, we infer that each tournament $T[W_i]$, $i \in [k]$, contains a topological minor copy of H. This gives k vertex-disjoint topological minor copies of H in T. □

Corollary 3.10. *Let H be a simple digraph that is acyclic and let k be a positive integer. Let T be a tournament that does not contain k vertex-disjoint topological minor copies of H. Then one can find in T a set of at most $2^{|H|}k$ vertices that is H-hitting.*

Proof. By Lemma 3.9 we have $|T| < 2^{|H|}k$, so we can take the whole vertex set of T as the requested H-hitting set. □

We now proceed to the strongly connected case and prove an analogue of Lemma 3.6. Note that in this setting, we can use the strategy from the proof of Lemma 3.6 and directly achieve vertex-disjointness. Hence, we will need no counterpart of Lemma 3.7.

Lemma 3.11. *Let H be a strongly connected simple digraph and let T be a tournament that does not contain k vertex-disjoint topological minor copies of H. Then in T one can find a set of at most $2d_{\mathrm{pw}}\|H\| \cdot k \log k$ vertices that is H-hitting.*

Proof. We proceed by induction on k. In the base case $k = 1$ there are no copies of H in T, hence we can take the empty set as an H-hitting set. Let us then assume that $k \geqslant 2$.

If T does not contain $\lceil k/2 \rceil$ vertex-disjoint copies of H, then as $\lceil k/2 \rceil < k$, we may apply the induction assumption for $\lceil k/2 \rceil$. Hence, from now on assume that T contains $\lceil k/2 \rceil$ vertex-disjoint copies of T.

By Corollary 2.6, T admits an interval decomposition of width at most $d_{\mathrm{pw}}\|H\| \cdot k$. Recall that we may assume that the endpoints of intervals in I correspond to pairwise different nonnegative integers.

Let α be the largest integer such that $T[I[\alpha]]$ does not contain $\lceil k/2 \rceil$ vertex-disjoint copies of H. By the assumption from the previous paragraph, α is well defined and $T[\alpha+1]$ contains $\lceil k/2 \rceil$ vertex-disjoint copies of H. It follows that $T[I[\alpha+2,\infty]]$ does not contain $\lfloor k/2 \rfloor$ vertex-disjoint copies of H, for otherwise in total we would obtain $\lceil k/2 \rceil + \lfloor k/2 \rfloor = k$ vertex-disjoint copies of H.

By induction assumption, in $T[I[\alpha]]$ and in $T[I[\alpha+2,\infty]]$ we can find H-hitting sets S_1 and S_2 of sizes $2d_{\mathrm{pw}}\|H\| \cdot \lceil k/2 \rceil \log\lceil k/2 \rceil$ and $2d_{\mathrm{pw}}\|H\| \cdot \lfloor k/2 \rfloor \log\lfloor k/2 \rfloor$, respectively. Let

$$S \coloneqq S_1 \cup S_2 \cup \mathsf{vcut}[\alpha+1].$$

We claim that S is H-hitting in T. Indeed, since H is strongly connected, every copy of H in T that does not intersect $\mathsf{vcut}[\alpha]$ must be entirely contained either in $T[I[\alpha]]$ or in $T[I[\alpha+2,\infty]]$, and then it intersects S_1 or S_2, respectively.

We are left with bounding the size of S. Observe that

$$\begin{aligned}
|S| &\leqslant |S_1| + |S_2| + |\,\mathsf{vcut}[\alpha+1]| \\
&\leqslant 2d_{\mathrm{pw}}\|H\| \left(\lceil k/2 \rceil \log\lceil k/2 \rceil + \lfloor k/2 \rfloor \log\lfloor k/2 \rfloor\right) + d_{\mathrm{pw}}\|H\|k \\
&\leqslant 2d_{\mathrm{pw}}\|H\| \left(k/2 \left(\log\lceil k/2 \rceil + \log\lfloor k/2 \rfloor\right) + 1/2 \left(\log\lceil k/2 \rceil - \log\lfloor k/2 \rfloor\right)\right) + d_{\mathrm{pw}}\|H\|k \\
&\leqslant 2d_{\mathrm{pw}}\|H\| \left(k \log(k/2) + 1/2\right) + d_{\mathrm{pw}}\|H\|k \\
&= d_{\mathrm{pw}}\|H\| \left(2k \log k - 2k + 1 + k\right) \leqslant 2d_{\mathrm{pw}}\|H\| \cdot k \log k.
\end{aligned}$$

This concludes the proof. □

We proceed to the main part of the proof, which is is again conceptually very close to the one presented in the case of immersions. It is arguably simpler, as we work only with vertex-disjointness.

Lemma 3.12. *Let H be a simple digraph that is not acyclic and let k be a positive integer. Let T be a tournament that does not contain k vertex-disjoint topological minor copies of H. Then one can find in T a set consisting of at most $6d_{\mathrm{pw}} \cdot |H|! \cdot \|H\| \cdot k \log k$ vertices that is H-hitting.*

PROOF. Denote $T = (V, E)$. Define Comps, h, Π, topological ordering and (non-)trivial components as in the proof of Lemma 3.8. Note that the same assertions about these objects apply.

By Corollary 2.6, T admits an interval decomposition I of width at most

$$p \coloneqq d_{\mathrm{pw}}\|H\| \cdot k.$$

Recall that we may assume that the endpoints of the intervals of I are pairwise different nonnegative integers. Let N be the largest interval end, i.e. $N \coloneqq \max\{\mathsf{last}(I(v)) \colon v \in V\}$. Define functions

$$\beta\colon \mathsf{Comps} \times \mathbb{Z} \to \mathbb{Z} \qquad \text{and} \qquad A,\ B\colon \mathsf{Comps} \times \mathbb{Z} \to \mathsf{Pow}(V)$$

as follows:

- $\beta(C,\alpha)$ is the minimum integer β with the property that interval $T[I[\alpha,\beta]]$ contains at least k vertex-disjoint topological minor copies of C. If no such β exists, we set $\beta = N$. Note that since we assume that the endpoints of the intervals in I are pairwise different, in either case $T[I[\alpha,\beta]]$ does not contain $k+1$ vertex-disjoint topological minor copies of C.
- If C is trivial, then $A(C,\alpha) = I[\alpha, \beta(C,\alpha)]$. If C is non-trivial, then $A(C,\alpha)$ is a C-hitting set of vertices in $T[I[\alpha,\beta(C,\alpha)]]$ of size at most $2d_{\mathrm{pw}}\|H\| \cdot (k+1)\log(k+1) \leqslant 5d_{\mathrm{pw}}\|H\| \cdot k \log k$, whose existence follows from Lemma 3.11.
- $B(C,\alpha) \coloneqq \mathsf{vcut}[\beta(C,\alpha)]$.

Note that since $p \leqslant 5d_{\mathrm{pw}}\|H\| \cdot k \log k$, for all C and α we have

$$|A(C,\alpha)| \leqslant 5d_{\mathrm{pw}}\|H\| \cdot k \log k \qquad \text{and} \qquad |B(C,\alpha)| \leqslant p.$$

Consider an arbitrary $\pi \in \Pi$ and define indices $\alpha_{\pi,0}$, $\alpha_{\pi,1}$, $\ldots$, $\alpha_{\pi,h}$ by induction as follows: $\alpha_{\pi,0} \coloneqq 0$ and, for $i = 1, 2, \ldots, h$, set

$$\alpha_{\pi,i} \coloneqq \beta(\pi^{-1}(i), \alpha_{\pi,i-1}).$$

Moreover, for $i \in [h]$ we define

$$I_{\pi,i} \coloneqq I[\alpha_{\pi,i-1}, \alpha_{\pi,i}], \qquad A_{\pi,i} \coloneqq A(\pi^{-1}(i), \alpha_{\pi,i-1}) \qquad \text{and} \qquad B_{\pi,i} \coloneqq B(\pi^{-1}(i), \alpha_{\pi,i-1}).$$

Note that since no interval in the decomposition I has length 0, sets $I_{\pi,i}$ for $i \in [h]$ are pairwise disjoint. Moreover, since intervals in I have pairwise different endpoints, for all $1 \leqslant i < j \leqslant h$, all arcs with one endpoint in $I_{\pi,i}$ and second in $I_{\pi,j}$ have tail in $I_{\pi,i}$ and head in $I_{\pi,j}$.

The following statement can be proved using the same arguments as the corresponding claim in the proof of Lemma 3.8 (that is, Claim 3.2). We simply join the copies of strong components of H by single forward arcs between intervals $I_{\pi,i}$.

Claim 3.4. *Suppose that there exists $\pi \in \Pi$ such that for every $i \in [h]$, the tournament $T[I_{\pi,i}]$ contains k vertex-disjoint topological minor copies of $\pi^{-1}(i)$. Then T contains k vertex-disjoint topological minor copies of H.*

Just as in the proof of Lemma 3.8, due to Claim 3.4 we may now assume that

$$\bigcup_{i=1}^{h} I_{\pi,i} = V(T) \qquad \text{for each } \pi \in \Pi.$$

Consider

$$S \coloneqq \bigcup_{\pi\in\Pi} \bigcup_{i\in[h]} A_{\pi,i} \cup B_{\pi,i}.$$

Since $|\Pi| \leqslant (|H|-1)!$ due to H not being acyclic, we have

$$|S| \leqslant (|H|-1)! \cdot |H| \cdot (5d_{\mathrm{pw}}\|H\| \cdot k \log k + d_{\mathrm{pw}}\|H\| \cdot k) \leqslant 6d_{\mathrm{pw}} \cdot |H|! \cdot \|H\| \cdot k \log k.$$

So it is enough to prove that S is H-hitting in T. Let $T' \coloneqq T - S$.

Claim 3.5. *T' is H-topological-minor-free.*

Proof. The proof follows precisely the same steps as the one of Claim 3.3, with minor and straightforward adjustments (e.g. instead of σ-intervals we consider simply intervals). ⌟

Claim 3.5 finishes the proof of Lemma 3.12. □

Now we can finish the proof of Theorem B.

Proof (of Theorem B). Suppose T is a tournament that does not contain k vertex-disjoint topological minor copies of H. If H is acyclic, then, by Corollary 3.10, we may find in T a set of at most $2^{|H|}k \in \mathcal{O}_H(k)$ vertices that is H-hitting. On the other hand, if H is not acyclic, then by Lemma 3.12 we may find in T an H-hitting set of vertices of size at most $6d_{\mathrm{pw}} \cdot |H|! \cdot \|H\| \cdot k \log k \in \mathcal{O}_H(k \log k)$. □

Chapter 4

Polynomial Kernel for Immersion Hitting in Tournaments

This chapter is devoted to proving Theorem C, i.e. designing a polynomial-time kernelization procedure for H-hitting Immersions in Tournaments (with the degree dependent on H).

Theorem C (restated). *For every simple digraph H without isolated vertices there exists a constant c and an algorithm that given an instance (T, k) of H-*hitting Immersions in Tournaments*, runs in polynomial time and returns an equivalent instance (T', k) with $|T'| \leqslant c \cdot k^c$.*

The proof shows how the protrusion replacement technique can be adapted to work in the world of tournaments. We mimic and adjust the steps developed in [FLMS12, GPR+21]: first we restrict our attention to sufficiently large and well-structured tournament (of cutwidth bounded in terms of the parameter k), then we identify and appropriately define large protrusion — "simple part" whose internal structure and interaction with the rest of the graph can be efficiently understood and encoded, and finally we show how to replace it with a protrusion of the same "behavior" with respect to the problem (i.e. keeping the answer unchanged) and strictly smaller size.

Throughout this chapter we fix a simple digraph H without isolated vertices and an integer $k \in \mathbb{N}$. For a tournament T, a set $F \subseteq A(T)$ is called a *solution* if $T - F$ is H-free. Moreover, F is an *optimal solution* if it is a solution of the smallest possible size. So (T, k) is a **yes**-instance of H-hitting Immersions in Tournaments if and only if in T there exists an optimal solution of size at most k.

4.1 Partial immersions

Our goal in this section is to extend the notion of an immersion to *partial immersions*. These will be used to understand possible behaviors of immersion models in a tournament T with respect to different intervals in an ordering of the vertex set of T. Let then $T = (V, A)$ be a tournament and let σ be an ordering of V. For now, fix a σ-interval $I \coloneqq \sigma(\alpha, \beta]$.

Definition 4.1. *A* scattered path *in I of* size *$q \geqslant 0$ is a sequence $\widetilde{P} = (P_i)_{i=1}^q$ satisfying the following properties:*

- *for each $i \in [q]$, P_i is a directed (simple) path of length at least 1 consisting of arcs that belong to $A(T[I]) \cup \partial(I) \cup \Gamma(I)$;*
- *paths P_i for $i \in [q]$ are pairwise arc-disjoint;*
- *for every $i \in [q]$, $i \neq 1$, we have $\mathsf{first}(P_i) \in \Gamma^-(I) \cup \partial^-(I)$;*
- *for every $i \in [q]$, $i \neq q$, we have $\mathsf{last}(P_i) \in \Gamma^+(I) \cup \partial^+(I)$.*

Each term in the sequence $(P_i)_{i=1}^q$ will be called a piece *of $\widetilde{P}$. The set of arcs of all pieces of $\widetilde{P}$ is denoted $A(\widetilde{P})$. If $\mathsf{first}(P_i) \in \Gamma^-(I)$ and $\mathsf{last}(P_i) \in \Gamma^+$, then the piece P_i is called* generic.

Note that $\mathsf{first}(P_1)$ is allowed to be an arc in the set $A(T[I]) \cup \Gamma^+(I) \cup \partial^+(I)$. If it is such, we call the vertex $\mathsf{tail}(\mathsf{first}(P_1)) \in I$ the *beginning* of $\widetilde{P}$ and denote it by $\mathsf{start}(\widetilde{P})$. Similarly, if $\mathsf{last}(P_q) \in A(T[I]) \cup$

$\Gamma^-(I) \cup \partial^-(I)$, then we call the vertex $\mathsf{head}(\mathsf{last}(P_q))$ the *end* of $\widetilde{P}$ and denote it by $\mathsf{end}(\widetilde{P})$. Note that the empty sequence is a scattered path of size 0. Also, a scattered path with only one piece, whose both beginning and end exist, is just a path in $T[I]$. By $\mathcal{P}_I$ we denote the family of all scattered paths in I.

We say that a scattered path $\widetilde{P} = (P_i)_{i=1}^q$ in I can be *shortened* to a scattered path $\widetilde{P}'$ (or that $\widetilde{P}'$ is a *shortening* of $\widetilde{P}$) if:

- $\mathsf{start}(\widetilde{P}) = \mathsf{start}(\widetilde{P}')$ and $\mathsf{end}(\widetilde{P}) = \mathsf{end}(\widetilde{P}')$ (meaning either equal or simultaneously undefined);
- for each piece P'_s of $\widetilde{P}'$ there exist indices $i_s^- \leqslant i_s^+$ such that $\mathsf{tail}(\mathsf{first}(P'_s)) = \mathsf{tail}(\mathsf{first}(P_{i_s^-}))$ and $\mathsf{head}(\mathsf{last}(P'_s)) = \mathsf{head}(\mathsf{last}(P_{i_s^+}))$; and
- whenever $s < s'$, we have $i_s^+ < i_{s'}^-$ and P'_s appears before $P'_{s'}$ in $\widetilde{P}'$.

Intuitively, shortening of the path means removing some pieces and replacing several contiguous subsequences of the pieces with single pieces, keeping the tail of the beginning and the head of the end of the replaced subsequence. Note that in particular, some pieces of $\widetilde{P}$ can be simply omitted in $\widetilde{P}'$ (other than the initial and the terminal one).

Definition 4.2. *A* partial immersion embedding *of H in I (or shortly, a* partial immersion *in I) is a mapping $\phi\colon A(H) \to \mathcal{P}_I$ such that*

- *all scattered paths $\phi(a)$ for $a \in A(H)$ are pairwise arc-disjoint;*
- *if $\mathsf{tail}(a) = \mathsf{tail}(a')$ then $\mathsf{start}(\phi(a))$ and $\mathsf{start}(\phi(a'))$ are either equal, or simultaneously undefined;*
- *if $\mathsf{start}(\phi(a))$ and $\mathsf{start}(\phi(a'))$ are defined and equal, then $\mathsf{tail}(a) = \mathsf{tail}(a')$;*
- *if $\mathsf{head}(a) = \mathsf{head}(a')$ then $\mathsf{end}(\phi(a))$ and $\mathsf{end}(\phi(a'))$ are either equal, or simultaneously undefined;*
- *if $\mathsf{end}(\phi(a))$ and $\mathsf{end}(\phi(a'))$ are defined and equal, then $\mathsf{head}(a) = \mathsf{head}(a')$.*

Intuitively, we can think of a partial immersion as of a "trace" which some immersion model $\widehat{H}$ of H in T "leaves" on the interval I. Some edges of H have images being paths in $\widehat{H}$ non-incident with I (these correspond to empty scattered paths in the partial immersion embedding). Some images of arcs of H come back and forth to I, intersecting with I along a non-empty scattered path (the ordering of paths on a single scattered path corresponds to the order of their appearance along the image of the respective arc of H). Finally, some arc images begin or end within I, which corresponds to the case when the beginning or the end of a scattered path is defined and is a vertex of I.

We call a partial immersion ϕ' in I a *shortening* of ϕ in I if for every $a \in A(H)$, the scattered path $\phi'(a)$ is a shortening of $\phi(a)$. We call ϕ *minimal* if there is no shortening of ϕ with at least one scattered path of strictly smaller size. Note that ϕ may be minimal even if some piece of some $\phi(a)$ can be replaced by a different single piece with equal first and last vertices. Shortening which does not decrease the size of any scattered path will be called *trivial*.

Note that each immersion model ϕ of H in T is a partial immersion in $V(T)$, in which all scattered paths $\phi(a)$, $a \in A(H)$, are paths in T beginning and ending at $\phi(\mathsf{tail}(a))$ and $\phi(\mathsf{head}(a))$, respectively. Moreover, each partial immersion ϕ in I gives rise to a natural partial immersion of H in $J \subseteq I$ in which all paths $\phi(a)$ where $a \in A(H)$ are "trimmed" to scattered paths consisting of precisely those arcs which are incident with J.

Formally, let I and J be σ-intervals such that $J \subseteq I$. If P is a path in I, then define the *trace* $P|_J$ of P on J to be the scattered path consisting of all arcs of P incident with J, arranged in the order of appearance along P. If $\widetilde{P} = (P_i)_{i=1}^q$ is a scattered path in I, then define the *trace* $\widetilde{P}|_J$ of $\widetilde{P}$ on J to be the concatenation of scattered paths $P_i|_J$. The *trace* $\phi|_J$ of a partial immersion ϕ of H in I on J is defined by setting $\phi|_J(a) = (\phi(a))|_J$ for every $a \in A(H)$.

Consider any σ-intervals I_1, I_2 with the property that there exist $\alpha, \beta, \gamma \in \{0, 1, \ldots, |V|\}$ such that $I_1 = \sigma(\alpha, \gamma]$ and $I_2 = \sigma(\gamma, \beta]$. We will call such a pair of intervals *consecutive*. Equivalently, two disjoint intervals are consecutive if their union $I = I_1 \cup I_2$ is a σ-interval as well. Let $\widetilde{P}_1$ be a scattered path in I_1 and $\widetilde{P}_2$ be a scattered path in I_2. We say that $\widetilde{P}_1$ and $\widetilde{P}_2$ are *compatible* if

- $A(\widetilde{P}_1) \cap A(I_1, I_2) = A(\widetilde{P}_2) \cap A(I_1, I_2)$;
- the set of pieces whose arc set is $A(\widetilde{P}_1) \cup A(\widetilde{P}_2)$ can be ordered to form a scattered path $\widetilde{P}$ in I with the property that all pieces of $\widetilde{P}_i$ appear in $\widetilde{P}$ in the same order as they do in $\widetilde{P}_i$, for $i = 1, 2$.

Every $\widetilde{P}$ described above will be called a *gluing* of $\widetilde{P}_1$ and $\widetilde{P}_2$. Note that a gluing is not necessarily uniquely defined, which can be seen particularly well when $A(\widetilde{P}_1) \cap A(I_1, I_2) = A(\widetilde{P}_2) \cap A(I_1, I_2) = \emptyset$ — in this case the pieces of both paths can be "shuffled" in any way only keeping the order of pieces originating from the same path.

Observation 4.1. *If $\widetilde{P}_1$ is compatible with $\widetilde{P}_2$ and $\widetilde{P}_2'$ is a shortening of $\widetilde{P}_2$, then there exists a shortening $\widetilde{P}_1'$ of $\widetilde{P}_1$ such that $\widetilde{P}_1'$ and $\widetilde{P}_2'$ are compatible and every gluing of them is a shortening of some gluing of $\widetilde{P}_1$ and $\widetilde{P}_2$.*

Proof. To construct $\widetilde{P}_1'$ from $\widetilde{P}_1$ it is enough to omit the pieces which do not share arcs with $\widetilde{P}_2'$. □

We will say that two partial immersions ϕ_1 in I_1 and ϕ_2 in I_2 are *compatible* if there exists a partial immersion ϕ in I such that $\phi_1 = \phi|_{I_1}$ and $\phi_2 = \phi|_{I_2}$, or — in other words — that for every $a \in A(H)$ the scattered path $\phi(a)$ is a gluing of $\phi_1(a)$ and $\phi_2(a)$. We will call every such ϕ a *gluing* of ϕ_1 and ϕ_2. Note that a gluing is not necessarily uniquely defined. Denote the set of all gluings of ϕ_1 and ϕ_2 by $\phi_1 \oplus \phi_2$.

Observation 4.2. *If $\phi \in \phi_1 \oplus \phi_2$ and ϕ is minimal, then ϕ_1 and ϕ_2 are minimal.*

Proof. Suppose that ϕ_1 is not minimal. Then there exists a shortening ϕ_1' of ϕ_1 which has a strictly smaller total number of pieces. For every $a \in A(H)$ the scattered path $\phi_1'(a)$ is compatible with some scattered path $\phi_2'(a)$ obtained by omitting several pieces in $\phi_2(a)$. In particular, the collection of all such $\phi_2'(a)$ — call it ϕ_2' — is a partial immersion compatible with ϕ_1' that is a shortening of ϕ_2. Any gluing belonging to $\phi_1' \oplus \phi_2'$ is a shortening of ϕ with a strictly smaller total number of pieces — contradiction. An analogous reasoning can be applied upon supposing that ϕ_2 is not minimal. □

The notions of a scattered path, partial immersion and trace can be naturally extended to co-intervals, by applying all definitions verbatim. If I is a σ-interval and $I' = V - I$ is the corresponding co-interval, then partial immersions ϕ_1 in I and ϕ_2 in I' are *compatible* if there exists an immersion ϕ in T such that $\phi_1 = \phi|_I$ and $\phi_2 = \phi|_{I'}$. Again, every such ϕ is called a *gluing* of ϕ_1 and ϕ_2.

Types of intervals. The key ingredient of our analysis is a constant-size encoding of the set of possible "behaviors" of partial immersions in intervals.

A σ-interval I shall be called *ℓ-long* if $|I| \geqslant \ell$. Further, we shall call I *c-flat* if $|\partial^+(I)| \leqslant c$, $|\partial^-(I)| \leqslant c$, and σ restricted to $T[I]$ has width at most c.

Note that if $I = \sigma(\alpha, \beta]$ is $2r$-long, then the intervals $I_r^- := \sigma(\alpha, \alpha + r]$ and $I_r^+ := \sigma(\beta - r, \beta]$ are disjoint. On the other hand, if I is c-flat, then we can color all backward arcs incident with I with at most $3c$ colors in such a way that each γ-cut of σ restricted to those arcs contains arcs of mutually different colors. This can be achieved e.g. by greedy coloring the γ-cuts for consecutive $\gamma = \alpha, \ldots, \beta + 1$. Formally, there exists a function $\xi\colon \overleftarrow{A}_\sigma(T) \cap A(I, V) \to [3c]$ such that for every $\gamma \in [|V| - 1]$ and every two distinct arcs $a_1, a_2 \in \mathsf{cut}_\sigma[\gamma] \cap A(I, V)$ we have $\xi(a_1) \neq \xi(a_2)$. In the following fix such a function.

Definition 4.3. *Let ϕ be a partial immersion in an interval I that is $2r$-long and c-flat. For each $a \in A(H)$ we define the (r, c)-type $\tau^{(r,c)}(\phi(a))$ of the scattered path $\phi(a) = (P_i)_{i=1}^q$ as the following sequence of length $2q$:*

$$(f_-(\mathsf{first}(P_1)), f_+(\mathsf{last}(P_1)), f_-(\mathsf{first}(P_2)), f_+(\mathsf{last}(P_2)), \ldots, f_-(\mathsf{first}(P_q)), f_+(\mathsf{last}(P_q))),$$

where the functions $f_\pm\colon A(T[I]) \cup \partial(I) \cup \Gamma(I) \to [-3c] \cup [r] \cup \{\mathsf{X}, \mathsf{H}\}$ are defined as follows

$$f_-(a) = \begin{cases} -\xi(a) & \text{if } a \in \partial^-(I), \\ \sigma(\mathsf{head}(a)) - \alpha & \text{if } a \in \Gamma^-(I) \text{ and } \mathsf{head}(a) \in I_r^-, \\ \mathsf{X} & \text{if } a \in \Gamma^-(I) \text{ and } \mathsf{head}(a) \notin I_r^-, \\ \mathsf{H} & \text{otherwise;} \end{cases}$$

$$f_+(a) = \begin{cases} -\xi(a) & \text{if } a \in \partial^+(I), \\ r + \sigma(\mathsf{tail}(a)) - \beta & \text{if } a \in \Gamma^+(I) \text{ and } \mathsf{tail}(a) \in I_r^+, \\ \mathsf{X} & \text{if } a \in \Gamma^+(I) \text{ and } \mathsf{tail}(a) \notin I_r^+, \\ \mathsf{H} & \text{otherwise.} \end{cases}$$

Let S be the set of all terms of the sequences $(\tau^{(r,c)}(\phi(a)))_{a \in A(H)}$, where by term we mean a value with an assigned position (i.e. equal values in different sequences are considered different terms). The (r, c)-type of ϕ is the collection of types $\tau^{(r,c)}(\phi) = (\tau^{(r,c)}(\phi(a)))_{a \in A(H)}$ equipped with a pair of equivalence relations (R_-, R_+) on the set $S \cup [3c]$ defined as follows.

- *If a piece* P_1 *of* $\phi(a_1)$ *and a piece* P_2 *of* $\phi(a_2)$ *satisfy* $\mathsf{head}(\mathsf{first}(P_1)) = \mathsf{head}(\mathsf{first}(P_2))$*, then the corresponding terms in* $\tau^{(r,c)}(\phi)$ *(elements of the set* $\{\mathsf{X}\} \cup [r]$ *with assigned positions in respective sequences) are in relation* R_-*.*
- *If a piece* P *of* $\phi(a)$ *is such that* $\mathsf{head}(\mathsf{first}(P))$ *is the tail of a singular arc of color* $x \in [3c]$*, then the corresponding term in* $\tau^{(r,c)}(\phi)$ *is in relation* R_- *with* x*.*
- *Analogously, if* $\mathsf{tail}(\mathsf{last}(P_1)) = \mathsf{tail}(\mathsf{last}(P_2))$*, then the terms* $f_+(\mathsf{last}(P_1))$ *and* $f_+(\mathsf{last}(P_2))$ *are in relation* R_+*. And* $x \in [3c]$ *is in relation with all terms corresponding to tails of ends of paths which are simultaneously the head of the singular arc of color* x*.*

Let us provide some intuition on what kind of information is stored in the type of ϕ defined above. First of all, the entire "singular interface" of this partial immersion is kept, i.e. in the type we remember precisely the singular arcs used to enter or exit I when traversing along each scattered path $\phi(a)$. Moreover, if we enter or exit I with a generic arc, we remember the precise vertex of entry/exit inside I, but only if it is close enough to the "border" of I (i.e. within the first or last r vertices in σ). Otherwise we remember the respective entry/exit as "generic", which is marked by the marker X. Moreover, we keep the information about whether the scattered path begins or ends inside I — the marker H represents that a vertex of H is mapped under the partial immersion embedding to a vertex within I. Finally if the extreme (first or last) arcs of some pieces are generic and have the same first/last vertex in I, we remember this fact in the equivalence relations $R_\pm$. We shall need this information to be able to "glue" two partial immersions without using the same generic gluing arc for different pieces. The incidence with singular arcs is also stored to avoid a situation when one attempts a generic gluing along a singular arc. The reason for colors being stored as negative integers is purely technical — it ensures that $[r] \cap [-3c] = \emptyset$.

An *(r,c)-type* is any collection of sequences of even lengths over the set $[-3c] \cup [r] \cup \{\mathsf{X}, \mathsf{H}\}$ indexed by $A(H)$ and equipped with a pair of equivalence relations (R_-, R_+) on the union of the set of all terms of these sequences and $[3c]$. An *I-admissible (r,c)-type* is every type τ such that there exists a partial immersion ϕ in I such that $\tau = \tau^{(r,c)}(\phi)$. The *size* of a type is the sum of lengths of its sequences (i.e. it does not depend on the equivalence relations).

Let γ_r be a function mapping each element from the set $[r]$ onto the single marker X, and identity otherwise. Intuitively, the function γ_r keeps only the information about generic nature of the end of a piece, and "forgets" about the closeness of this vertex to the boundary. We say that an (r,c)-type τ' is a *shortening* of an (r,c)-type τ if for every $a \in A(H)$ the sequence $\gamma_r(\tau'_a)$ is a subsequence of $\gamma_r(\tau_a)$ and the two sequences have the same first and last terms.

The following observation is a direct consequence of the definition of shortening of a partial immersion.

Observation 4.3. *If a partial immersion* ϕ' *of* (r,c)*-type* τ' *is a shortening of a partial immersion* ϕ *of* (r,c)*-type* τ*, then* τ' *is a shortening of* τ*.*

Definition 4.4. *An* I*-admissible* (r,c)*-type* τ *is called* minimal *(in* I*) if there is no* I*-admissible type* τ' *of strictly smaller size that would be a shortening of* τ*.*

Observation 4.4. *Let* ϕ *be a partial immersion in* I *of* (r,c)*-type* τ*. Then* ϕ *is minimal if and only if* τ *is minimal in* I*.*

Proof. If ϕ is not minimal, then for any nontrivial shortening ϕ' of ϕ, $\tau^{(r,c)}(\phi')$ is a shortening of τ of strictly smaller size. Conversely, if τ is not minimal, then there exists an I-admissible type τ' of strictly smaller size. Every partial immersion ϕ' of type τ' is a nontrivial shortening of ϕ. □

We now observe that in a minimal partial immersion on an interval I, all scattered paths will be relatively small, that is, will visit I only a bounded number of times.

Lemma 4.5. *Suppose that* $F \subseteq A(T)$ *has the property that* $|F| \leqslant f$*,* I *is* $4\|H\|(c+f+1)$*-long and* c*-flat, and a partial immersion* ϕ *in* I *is minimal and disjoint with* F*. Then for every* $a \in A(H)$ *the scattered path* $\phi(a)$ *has size at most* $2c+3$*.*

Proof. We will show that every partial immersion ϕ disjoint with F can be shortened to a partial immersion ϕ' disjoint with F in which every path $\phi(a)$ has at most one generic piece. This will yield the desired result as every nongeneric piece of $\phi(a)$ which is neither the initial, nor the terminal one, contains at least one arc from the set $\partial(I)$, which is of size not exceeding $2c$ (as I is c-flat).

Take an arbitrary ϕ disjoint with F. For every $a \in A(H)$, let $\phi(a) = (P^a_i)^{q_a}_{i=1}$. Let i^-_a be the smallest index i such that $\mathsf{first}(P^a_i) \in \Gamma^-(I)$ (if there is no such index, we put $i^-_a = \infty$). Similarly, let i^+_a be the greatest index i such that $\mathsf{last}(P_i^a) \in \Gamma^+(I)$ (or $i_a^+ = -\infty$ if no such index exists). Note that if $i < i_a^-$ or $i > i_a^+$, then P_i^a is not generic, so

$$\left| \bigcup_{a \in A(H)} \bigcup_{i \notin [i_a^-, i_a^+]} \{P_i^a\} \right| \leqslant \|H\| \cdot (2c+2).$$

The set on the left-hand side of the above inequality comprises of pieces that we will keep unchanged in ϕ; we will show a way to shorten ϕ by replacing the subsequence $(P_i^a)_{i=i_a^-}^{i_a^+}$ with a single generic piece for each $\phi(a)$ for which this subsequence is nonempty (if $i_a^- > i_a^+$, then $\phi(a)$ has no generic pieces). This will conclude the proof as the new piece will be the only generic one.

Let

$$\mathcal{G} = \Gamma(I) \cap \bigcup_{a \in A(H)} \bigcup_{i \notin [i_a^-, i_a^+]} A(P_i^a),$$

so in particular $|\mathcal{G}| \leqslant \|H\| \cdot (2c+2)$ as each piece P_i^a with $i \notin [i_a^-, i_a^+]$ contains at most one generic arc in $\Gamma(I)$. Since I is $2\|H\|(c+f+1)$-long, we may find in I a subset J of vertices not incident with $F \cup \mathcal{G} \cup \partial(I)$ such that $|J| \geqslant \|H\|$. Indeed,

$$|I| \geqslant 4\|H\|(c+f+1) \geqslant 2f + 2(c+1)\|H\| + 2c + \|H\| \geqslant 2|F| + |\mathcal{G}| + |\partial(I)| + \|H\|.$$

Now assign to each $a \in A(H)$ a different vertex $v(a) \in J$. We can use these vertices are "pivots" for the newly constructed generic pieces. Indeed, for every $a \in A(H)$ such that $i_a^- \leqslant i_a^+$ consider the path

$$P_\pm^a = (\mathsf{tail}(\mathsf{first}(P_{i_a^-}^a)), v(a))(v(a), \mathsf{head}(\mathsf{last}(P_{i_a^+}^a))).$$

This path is well-defined as the two arcs it consists of are generic (because $v(a)$ was chosen not to be incident with $\partial(I)$). Moreover, it is edge-disjoint with $\mathcal{G} \cup F$ and with all other $P_\pm^{a'}$ ($a' \neq a$), as $v(a) \neq v(a')$. Inserting paths $P_\pm^a$ in place of the subsequences described above gives a desired shortening and finishes the proof. □

Remark 4.1. In all the results throughout this chapter $\|H\|$ can actually be replaced with $\|H\| - |H|$ (i.e. the number of arcs of H). The weakening of these results does not substantially affect the main theorem and is meant to make the presentation more clear.

We call an (r,c)-type τ *short* if for every $a \in A(H)$ the length of $\tau(\phi(a))$ is at most $4c+6$.

Corollary 4.6. *For $r \geqslant 1$ we have the following:*

1. *If an interval I is $4\|H\|(r+c)$-long and c-flat, and a partial immersion ϕ in I is minimal, then $\tau^{(r,c)}(\phi)$ is short.*
2. *For every pair of integers $r, c \in \mathbb{N}$ there is a constant $t(r,c)$ such that there are exactly $t(r,c)$ short (r,c)-types. In particular, for any interval I as above, there are at most $t(r,c)$ different I-admissible minimal (r,c)-types.*

Proof. It is enough to take $F = \emptyset$ in Lemma 4.5 to see that the lengths of sequences $\tau^{(r,c)}(\phi(a))$ for $a \in A$ are uniformly bounded by $2(2c+3)$. To prove that the number of types is bounded in terms of r and c it is now enough to observe that the number of different equivalence classes R_+, R_- is so. The last claim is true as these are the equivalence classes introduced on a set of size bounded in terms of r and c. □

Let I_1, I_2 be two consecutive c-flat, $2(r+c)$-long σ-intervals and $I = I_1 \cup I_2$. We say that two (r,c)-types τ_1 of I_1 and τ_2 of I_2 are *compatible* if there exists a partial immersion ϕ in I such that $\phi|_{I_1}$ and $\phi|_{I_2}$ are compatible, $\phi|_{I_1}$ has type τ_1, and $\phi|_{I_2}$ has type τ_2.

The *gluing* of the types $\tau_1 \oplus \tau_2$ is defined as the set of all (r,c)-types of all such ϕ. The following lemma proves that this definition is correct, i.e. that types of two immersions store enough information to ensure the possibility of gluing them.

Lemma 4.7. *Let I_1, I_2 be two consecutive c-flat, $4\|H\|(r+c)$-long σ-intervals and $I = I_1 \cup I_2$. If two (r,c)-types τ_1 of I_1 and τ_2 of I_2 are* compatible*, then for every partial immersion ϕ_1 in I_1 of type τ_1 and for every partial immersion ϕ_2 in I_2 of type τ_2, the immersions ϕ_1 and ϕ_2 are compatible.*

Proof. By the definition of compatible types, we have at least one pair of compatible partial immersions η_1 in I_1 and η_2 in I_2, of types τ_1 and τ_2 respectively. Consider any $\eta \in \eta_1 \oplus \eta_2$. We will use η as a "model" for understanding how to successfully glue ϕ_1 and ϕ_2.

Fix $a \in A(H)$ and mappings π_i sending the pieces of path $\eta_i(a)$ to the corresponding pieces of path $\phi_i(a)$, in accordance with the equality of types, where $i = 1, 2$. Moreover, color in $\eta(a)$ all arcs contained in $A(I_1, I_1) \cup \partial^+(I_1) \cup \Gamma^-(I_1)$ blue, all arcs contained in $A(I_2, I_2) \cup \partial^-(I_2) \cup \Gamma^+(I_2)$ green, and all other arcs (i.e. the "gluing" arcs in $A(I_1, I_2)$) red. Thus, every piece in $\eta(a)$ consists of alternating (and possibly empty) green and blue paths with a red arc between each two consecutive paths of different colors.

Suppose that some piece P of $\eta(a)$ contains a red arc e. Assume that the maximal monochromatic nonred subpath of P directly preceding this arc is blue and the maximal monochromatic nonred subpath of P directly following this arc is green (the reasoning in the other case is analogous). Then the blue path is a subpath of some piece P_1 of η_1 and the green path is a subpath of some piece P_2 of η_2, where the only difference is the arc e and possibly the other extreme arcs (if they are colored red as well). We will show that we may directly glue pieces $\pi_1(P_1)$ of ϕ_1 and $\pi_2(P_2)$ of ϕ_2 using the arc

$$e' = (\mathsf{tail}(\mathsf{last}(\pi_1(P_1))), \mathsf{head}(\mathsf{first}(\pi_2(P_2))))$$

and moreover — that such gluing can be performed simultaneously in all places corresponding to red arcs of η. This yields a construction of a gluing $\phi \in \phi_1 \oplus \phi_2$ therefore certifying the desired compatibility. Moreover, $\tau^{(r,c)}(\phi) = \tau^{(r,c)}(\eta)$ and η was chosen arbitrarily, which implies that the types of immersions being glued are enough to determine all possible types of gluings.

First of all note that if the arc e is singular of color x, then this fact is stored in the equivalence relations R_+ of $\tau^{(r,c)}(\eta_1)$ and R_- of $\tau^{(r,c)}(\eta_2)$. Therefore, by the equality of types of η_i and ϕ_i, the arc e' is singular of color x as well, hence $e' = \mathsf{last}(\pi_1(P_1)) = \mathsf{first}(\pi_2(P_2))$ and the gluing is possible.

If e is generic then, again by the respective equivalence relations, no other generic gluing arc joins vertices $\mathsf{tail}(\mathsf{last}(P_1))$ and $\mathsf{head}(\mathsf{last}(P_1))$. This means that e' is generic and does not correspond to any other red arc of η than e. Therefore again the gluing along this arc can be performed. $\square$

Boundaried intervals and signatures. In the process of replacing protrusions we will need to consider intervals not as subsets of an ordered vertex set of a larger tournament, but as standalone structures which can be used to replace one another. We introduce the notion of an (r, c)-boundaried interval to enable such considerations.

Definition 4.5. *An (r, c)-boundaried interval is a digraph D on vertex set $V(D) = S^+ \cup I \cup S^-$ equipped with an ordering σ_I of I. Furthermore, we require the following:*

- $|I| \geqslant 4\|H\|(r + c)$;
- $D[I]$ *is a tournament;*
- σ_I *has width at most c;*
- $S^- \subseteq [3c] \times \{-\}$ *and* $S^+ \subseteq [3c] \times \{+\}$;
- *each vertex of S^+ has only one incident arc and this arc belongs to the set $\overrightarrow{A}(I, S^+)$;*
- *each vertex of S^- has only one incident arc and this arc belongs to the set $\overrightarrow{A}(S^-, I)$.*

The r-boundary of D is the pair of sets I^- and I^+ consisting of the first and the last r vertices of I in σ_I, respectively. For D defined above, we will shortly write $D = I \cup S^\pm$.

Note that this notion emulates a $4\|H\|(r + c)$-long c-flat interval I in the following sense. Arcs whose heads are contained in S^+ correspond to $\partial^+(I)$, and arcs whose tails are contained in the set S^- correspond to $\partial^-(I)$. The names of the auxiliary vertices in $S^\pm$ correspond to the ξ-colors of the respective backward arcs. The r-boundary consists of precisely those vertices whose generic entry or exit is remembered in the (r, c)-type of a partial immersion.

Formally, every $4\|H\|(r + c)$-long c-flat σ-interval I in T can be uniquely encoded with an (r, c)-boundaried interval $D^{(r,c)}(I)$, whose structure resembles the structure of $T[I]$ and $\partial(I)$ as follows:

- the vertices and arcs of I are kept along with their ordering in T;
- every singular arc $a \in \partial^\pm(I)$ is mapped to an arc joining $(\xi(a), \pm) \in S^\pm$ with the endpoint of a contained in I;

- the projection of $S^{\pm}$ onto the first coordinate is precisely $\{\xi(a) \mid a \in \partial^{\pm}(I)\}$.

The notion of a partial immersion can be naturally adjusted to the setting of boundaried intervals. The only difference is the lack of the "generic interface" i.e. there are no auxiliary edges in boundaries intervals used to emulate $\Gamma(I)$. These can be, however, emulated by storing the information from the type (marker X or number in $[r]$ if the generic arc is incident with the r-boundary, and the equivalence relations $R_{\pm}$) instead of the identity of particular generic arcs. Formally, a piece of a scattered path in (r, c)-boundaried interval can begin or end with an element in $\{\mathsf{X}\} \cup [r]$ instead of a generic arc. In particular, this slightly modified variant of partial immersions can be equipped with precisely the same definition of admissible (r, c)-type as in the former case.

Definition 4.6. *An* (r,c)-signature *is a subset of the set of all short* (r,c)*-types. The* (r,c)-signature $\Sigma^{(r,c)}(I)$ *of a* $4\|H\|(r+c)$*-long* c*-flat* σ*-interval* I *is the set of all* I*-admissible minimal* (r,c)*-types. The* (r,c)-signature $\Sigma^{(r,c)}(D)$ *of an* (r,c)*-boundaried interval* $D = I \cup S^{\pm}$ *is the set of all* I*-admissible minimal* (r,c)*-types.*

The intuition behind this definition is that if I is appropriately long and flat, then the signature of I stores the information about all possible interactions of I with minimal partial immersions. Note that $\Sigma^{(r,c)}(D^{(r,c)}(I)) = \Sigma^{(r,c)}(I)$.

We say that two (r, c)-boundaried intervals $I \cup S^{\pm}$ and $I' \cup S'^{\pm}$ are *exchangeable* if they have equal (r, c)-signatures, $S^{\pm} = S'^{\pm}$ (both intervals use precisely the same colors on the boundary), and the incidence structure of r-boundaries of those intervals with backward arcs is the same, i.e. for every $i \in [r]$ the set of colors of singular arcs incident with both $S^{\mp}$ and the i-th vertex of $I^{\pm}$ is the same as analogously defined set of colors for i-th vertex of $I'^{\pm}$. Intuitively this means that in T we may replace the interval I with I'.

Corollary 4.8. *For every pair of integers* $r, c \in \mathbb{N}$ *there is a constant* $s(r, c)$ *such that there are exactly* $s(r, c)$ *different* (r, c)*-signatures.*

Proof. We may set $s(r, c) = 2^{t(r,c)}$, where $t(r, c)$ is the constant provided by Corollary 4.6. □

For future discussion of algorithmic aspects, we will need the following observation.

Lemma 4.9. *Consider* r *and* c *fixed and let* T*,* σ*, and* I *be as in Definition 4.6. Then given* T*,* σ*, and* I*, the signature* $\Sigma^{(r,c)}(I)$ *can be computed in polynomial time.*

Proof. It suffices to decide, for every short (r, c)-type τ, whether τ is I-admissible. This can be done in polynomial time using the algorithm for rooted immersion containment of Fomin and Pilipczuk (Theorem 9.1 in [FP19]). □

Let $\mathcal{S}^{(r,c)}$ be the set of all (r, c)-signatures; we have $|\mathcal{S}^{(r,c)}| = s(r, c)$, where $s(r, c)$ is the constant given by Corollary 4.8. Let $\mathcal{S}^{(r,c)}_{\sigma}$ be the set of those (r, c)-signatures which are equal to $\Sigma^{(r,c)}(I)$ for some σ-interval I, i.e. contain only I-admissible minimal (r, c)-types. Then $\mathcal{S}^{(r,c)}_{\sigma} \subseteq S^{(r,c)}$, so $|\mathcal{S}^{(r,c)}_{\sigma}| \leqslant s(r, c)$. We argue that $\mathcal{S}^{r,c}_{\sigma}$ has a structure of a semigroup in the following sense.

Lemma 4.10. *Let* I_1*,* I_2 *be two* $4\|H\|(r + c)$*-long* c*-flat* σ*-intervals such that* $I = I_1 \cup I_2$ *is a* c*-flat* σ*-interval. Then* $\Sigma^{(r,c)}(I)$ *is uniquely determined by* $\Sigma^{(r,c)}(I_1)$ *and* $\Sigma^{(r,c)}(I_2)$*.*

Proof. It is enough to prove that:
(i) for every $\tau \in \Sigma^{(r,c)}(I)$ there exist $\tau_1 \in \Sigma^{(r,c)}(I_1)$ and $\tau_2 \in \Sigma^{(r,c)}(I_2)$ which are compatible and $\tau \in \tau_1 \oplus \tau_2$; and
(ii) for every two compatible $\tau_1 \in \Sigma^{(r,c)}(I_1)$ and $\tau_2 \in \Sigma^{(r,c)}(I_2)$ and every $\tau \in \tau_1 \oplus \tau_2$ that is minimal in I, we have $\tau \in \Sigma^{(r,c)}(I)$.

For the proof of (i) take a minimal partial immersion ϕ in I with $\tau^{(r,c)} = \tau$. We know that $\tau \in \tau^{(r,c)}(\phi|_{I_1}) \oplus \tau^{(r,c)}(\phi|_{I_2})$, so it is enough to prove that $\phi|_{I_1}$ and $\phi|_{I_2}$ are minimal. But this follows directly from Observation 4.2.

For the proof of (ii) note that $\tau_1 \oplus \tau_2$ consists only of I-admissible types. □

Lemma 4.10 implies that the set $\mathcal{S}^{(r,c)}_{\sigma} \cup \{0\}$ can be endowed with an associative binary product operation such that for every two consecutive intervals I_1, I_2, the product of their signatures is the signature of their union $I_1 \cup I_2$. Formally, we set the product for all pairs of consecutive intervals as above; Lemma 4.10 ensures that this is well-defined. Next, for all pairs of elements $\tau_1, \tau_2 \in \mathcal{S}_\sigma^{r,c}$ for which their product is not yet defined, we set $\tau_1 \cdot \tau_2 = 0$. Also, we set $0 = 0 \cdot 0 = 0 \cdot \tau = \tau \cdot 0$ for all $\tau \in \mathcal{S}_\sigma^{(r,c)}$. In this way, $\mathcal{S}_\sigma^{(r,c)} \cup \{0\}$ becomes a monoid; the empty signature is the neutral element of multiplication.

By Lemma 2.8 we obtain the following.

Corollary 4.11. *Suppose I is a c-flat $4\|H\|(r+c)\ell^{3s(r,c)}$-long σ-interval. Then there exists a sequence of consecutive $4\|H\|(r+c)$-long c-flat σ-intervals $(I_i)_{i=1}^{\ell}$ whose (r,c)-signatures are equal and equal to the signature of their union. Moreover, given r, c, T, σ, and I, such a sequence can be found in polynomial time.*

Proof. Let $(X_i)_{i=1}^{\ell^{3s(r,c)}}$ be an arbitrary sequence of consecutive $4\|H\|(r+c)$-long c-flat σ-intervals included in I. Note that $\Sigma_\sigma^{(r,c)} \colon A^* \to \mathcal{S}^{(r,c)}$ is a properly defined morphism, sending a single interval to its signature and a concatenation of two consecutive intervals to the unique signature of their union (which is not 0). Applying Lemma 2.8 finishes the proof; the algorithmic statement follows from the remark after Lemma 2.8. □

4.2 Finding protrusions

In order to find an appropriately large subgraph of T which does not "affect" the behavior of T with respect to H-hitting Immersions in Tournaments, we roughly proceed as follows. First, we find a suitable ordering σ of $V(T)$ and an appropriately long interval X in σ such that X has a constant-size singular interface towards the remainder of T. Then, inside X, we find (again, an appropriately long) subinterval I of a very specific structure: I can be divided into $2k+3$ subintervals with the same signatures as itself. This is where we use Simon Factorization through Corollary 4.11. In the next part we use this extra structure to prove that one of these subintervals can be replaced with a strictly smaller replacement in such a way that after the substitution, we obtain an equivalent instance of the problem.

We proceed to a formal implementation of this plan. The first lemma gives the ordering σ and the interval X.

Lemma 4.12. *Let T be a tournament with $\mathbf{ctw}(T) \leqslant c$ and $|T| \geqslant (2c+1)(x+1)(k+1)$. If T contains at most k arc-disjoint immersion copies of H, then there exists an ordering σ of $V(T)$ and an H-free σ-interval X such that $|X| \geqslant x$ and X is c_H-flat with respect to σ.*

Moreover, given T, k, c, x as above, one can in polynomial time either conclude that T contains more than k arc-disjoint immersion copies of H, or find an ordering σ and an interval X satisfying the above properties.

Proof. We first argue the existential statement, and then address the algorithmic aspects at the end.

Let σ_0 be an ordering of $V(T)$ of width $\mathbf{ctw}(T) \leqslant c$. Consider any $k+1$ vertex-disjoint σ_0-intervals, each containing at least $(2c+1)(x+1)$ vertices. By assumption, at least one of them, say I, induces an H-free graph. Let $B = \partial(I)$. Then $|B| \leqslant 2c$.

Note that we have $\mathbf{ctw}(T[I]) \leqslant c_H$ by Corollary 2.3. Let σ_I be an ordering of vertices of I of width $\mathbf{ctw}(T[I])$ and let σ be the ordering obtained from σ_0 by reordering I according to σ_I, that is,

- If $u, v \in I$ then $\sigma(u) < \sigma(v) \iff \sigma_I(u) < \sigma_I(v)$.
- If $u \notin I$ or $v \notin I$, then $\sigma(u) < \sigma(v) \iff \sigma_0(u) < \sigma_0(v)$.

Consider the set

$$I_{-B} \coloneqq I - \left(\{\mathsf{tail}(a) \mid a \in \partial^+(I)\} \cup \{\mathsf{head}(a) \mid a \in \partial^-(I)\}\right),$$

i.e. the vertices in I which are not incident with B. Since $|B| \leqslant 2c$, we have that $|I_{-B}| > (2c+1)x$ and I_{-B} is the union of at most $2c+1$ σ-intervals. Hence, at least one of these intervals has size at least x. Call it X.

Let

$$I^- \coloneqq \{v \in I \mid \sigma(v) < \sigma(w) \text{ for every } w \in X\} \quad \text{and} \quad I^+ = I - I^- - X.$$

Note that as no arc from B is adjacent to X, $\overrightarrow{A}(X, I^-)$ contains precisely all σ-backward arcs in T with tail in X, and $\overrightarrow{A}(I^+, X)$ contains all those with head in X. On the other hand, both $\overrightarrow{A}(X, I^-)$ and $\overrightarrow{A}(I^+, X)$ have size bounded by the width of σ_I, which is at most c_H. It follows that X is c_H-flat with respect to σ.

As for the algorithmic aspect, an ordering σ_0 with optimum width can be computed using Lemma 2.1. Then we can divide σ_0 into $k+1$ disjoint σ_0-intervals of size at least $(2c+1)(x+1)$ each, and for each of them check in polynomial time whether it induces an H-free subtournament using the algorithm from [FP19]. If this check fails for each of the intervals, we may terminate the algorithm and conclude that T contains $k+1$ arc-disjoint copies of H. Otherwise, we have one interval I such that $T[I]$ is H-free, and it is straightforward to turn the remainder of the reasoning into a polynomial-time algorithm computing σ and X. $\square$

In the remainder of this section let $x \geqslant 4\|H\|$, r, c be fixed positive integers. Moreover, let T be a tournament for which there exists an optimal solution of size at most k and let σ be an ordering of T. Finally let X be an H-free c-flat $x(3c+k+1)(2k+3)^{3s(r,c)}$-long σ-interval. We assume that there is a coloring ξ mapping all σ-backward arcs incident to X to colors in $[3c]$ such that not two such arcs of the same color participate in the same σ-cut.

We may apply Corollary 4.11 to X to find a collection of $2k+3$ consecutive σ-intervals I_i with $|I_i| = x(3c+k+1)$ for all $i \in [2k+3]$ such that all I_i have equal (r,c)-signatures, and this common signature, call it Σ, is equal to the (r,c)-signature of their union I. Then I is an H-free c-flat $x(3c+k+1)(2k+3)$-long σ-interval. That such I can be found in polynomial time (given r, c, x, T, X, σ, and ξ) follows from Corollary 4.11.

Now comes a key step in the proof: we argue that from the equality of types of $I, I_1, \ldots, I_{2k+3}$ it follows that every optimum solution will contain a bounded number of arcs incident with I.

Lemma 4.13. *For every optimal solution $F \subseteq A(T)$, we have $|F \cap A(I, V(T))| \leqslant 2c$.*

Proof. Fix an optimal solution F. By assumption, $|F| \leqslant k$.

Define

$$F_\star = (F - A(I, V(T))) \cup \partial(I).$$

That is, $F_\star$ is obtained from F by removing all arcs incident with I and replacing them with all I-singular arcs. We will prove that $F_\star$ is a solution in T. Observe that since F is optimal, this will imply that

$$\begin{aligned} |F \cap A(I, V(T))| &= |F| - |F - A(I, V(T))| \\ &\leqslant |F_\star| - |F - A(I, V(T))| = |F_\star \cap A(I, V(T))| = |\partial(I)| \leqslant 2c, \end{aligned}$$

and in consequence we may set $X_\star \coloneqq I$.

Suppose for the sake of contradiction that $F_\star$ is not a solution in T, i.e. there exists an immersion model ϕ of H in $T - F_\star$. We may choose ϕ so that $\phi|_I$ is minimal. As

$$|I| \geqslant x(3c+k+1) \geqslant 4\|H\|(c+(2c+k)+1) \geqslant 4\|H\|(c+|F_\star|+1),$$

by Lemma 4.5 we have that $\tau^{(r,c)}(\phi|_I)$ is short. Note that ϕ is incident with I for otherwise it would have not been hit by F, contradicting F being a solution. Moreover, $A(\phi) \cap A(I, V(T)) \subseteq \Gamma(I)$ as all singular arcs incident with I are in $F_\star$ which is disjoint with ϕ.

Observe that since each of the at most k arcs in F is incident with at most 2 intervals I_j, there exists $i \in [2k+3]$ such that $i \neq 1$, $i \neq 2k+3$, and I_i is not incident with F. Pick vertices $v_\alpha \in V(I_1)$, $v_\omega \in V(I_{2k+3})$ which are not incident with $\partial(I) \cup \partial(I_i) \cup F$. This is possible since

$$|I_1|, |I_{2k+3}| \geqslant 3c+k+1 > c+c+k \geqslant |\partial(I)| + |\partial(I_i)| + |F| \geqslant |\partial(I) \cup \partial(I_i) \cup F|.$$

By equality of signatures of I and I_i and the shortness of $\tau^{(r,c)}(\phi|_I)$, we conclude that there exists a partial immersion ϕ' in I_i of the same type as $\phi|_I$, which moreover has a purely generic boundary (i.e. does not contain any I-singular arcs).

Now it is enough to observe that we can connect ϕ' with $\phi|_{T-I}$ to get an immersion of H in T disjoint with F. We shall use v_α and v_ω as pivots for two-arc generic paths, similarly as in the proof of Lemma 4.5. In other words, to enrich ϕ' and $\phi|_{T-I}$ to a full immersion model of H, we insert a collection of two-arc generic paths with middle vertices in v_α or v_ω, joining vertices on ϕ' with vertices on $\phi|_{T-I}$.

Formally, for every $a \in A(H)$ for which $\phi(a)$ is incident with both I and $T-I$ we do the following. Consider all arcs used for gluing $\phi|_I(a)$ with $\phi|_{T-I}(a)$. As argued above, all these arcs are generic. For each such arc $g \in \Gamma^+(I)$ consider the piece P of the scattered path $\phi(a)$ with $g = \mathsf{last}(P)$ and let P' be the corresponding piece of the corresponding (via equality of the (r,c)-types) scattered path $\phi'(a)$. Now replace the arc $\mathsf{last}(P')$ with a two-arc path

$$(\mathsf{tail}(\mathsf{last}(P')), v_\omega)(v_\omega, \mathsf{head}(g))$$

joining P' with the piece Q of $\phi|_{T-I}(a)$ which is glued in ϕ along g to P. Finally remove g from Q. We perform a similar replacement for each $g \in \Gamma^-(I)$ (changing the roles in each pair first/last, head/tail, v_α/v_ω). The choice of v_α and v_ω ensures that the new connections are arc-disjoint and disjoint with F.□

We define a digraph T° based on T and $I = I_1 \cup \ldots \cup I_{2k+3}$ as follows: start with T, and

- remove all vertices of I_2;
- for every arc $a \in \partial^+(I_2)$, replace a with an arc with the same head as a and tail in a fresh vertex $s^+_{\xi(a)}$;
- for every arc $a \in \partial^-(I_2)$, replace a with an arc with the same tail as a and head in a fresh vertex $s^-_{\xi(a)}$.

We call the constructed graph a *c-boundaried co-interval*. The intuition of this construction is as follows. We pinch off one of the $2k+3$ intervals and keep the singular arc interface in a fashion similar as in (r,c)-boundaried intervals. The only difference is that we do not keep track of the r-boundary vertices.

Now we can define the *gluing* of T° with an (r,c)-boundaried interval $B = I \cup S^\pm$ with signature Σ, simply by identifying the singular arcs of the same color (note that different vertices from $S^\pm$ can be therefore mapped to the same vertex) and completing the obtained structure to a tournament by making all missing arcs generic. Note that in order for this to be well-defined, we need to require that the sets of colors of the singular arcs in T° and in B are identical — if it is so, we will say that T° and B are *compatible*. Also, note that we require that the signature of B is Σ: that is, the possible types of partial immersions present in B are exactly the same as in the substituted interval I_2.

Denote by $T^\circ \oplus B$ the tournament obtained from gluing T° and B. Note that in this tournament we have naturally defined ordering of vertices: in T° it is inherited from the ordering σ of T and within B it is inherited from the ordering σ_I of the boundaried interval. Finally all the vertices of the substituted interval appear in the ordering between the two interval parts (prefix and suffix) of the co-interval. The following observation is obvious.

Observation 4.14. *If two exchangable (r,c)-boundaried intervals $B = I \cup S^\pm$ and $B' = I' \cup S^\pm$ are compatible with T°, then the (r,c)-signatures of I in $T^\circ \oplus B$ and I' in $T^\circ \oplus B'$ are equal.*

We now observe, by inspecting the proof of Lemma 4.13, that if in I we replace one of its subintervals with an exchangeable interval, then the conclusion of Lemma 4.13 — that the modified I' will still have constant incidence with every optimal solution in T — will still hold. Let us summarize this section with putting together these observations and recalling all needed assumptions.

Corollary 4.15. *Suppose that $x \geqslant 4\|H\|$, r, c are fixed positive integers, T is a tournament of for which there exists an optimal solution of size at most k and σ is an ordering of T. Suppose further that X is an H-free c-flat $x(3c+k+1)(2k+3)^{3s(r,c)}$-long σ-interval with all incident σ-backward arcs colored according to ξ with colors $[3c]$ so that no two arcs of the same color participate in the same σ-cut.*

Then there exists an (r,c)-signature Σ, a c-boundaried co-interval T° of this signature, and an (r,c)-boundaried interval $B = I \cup S^\pm$ such that $T = T^\circ \oplus B$ and moreover for every $B' = I' \cup S'^\pm$ exchangeable with B and for every optimal solution $F \subseteq A(T')$ where $T' = T^\circ \oplus B'$ we have: $|F \cap A(I', T')| \leqslant 2c$. Moreover, given x, r, c, T, σ, and X, such Σ, T°, and B can be computed in polynomial time.

Proof. It is enough to see that the argument from the proof of Lemma 4.13 can be applied to both T and T' to obtain constant incidence of any optimum solution with I (and I'). Constant incidence with an interval means in particular constant incidence with each of its subintervals. As for the algorithmic statement, it follows directly from Corollary 4.11 and Lemma 4.9. □

4.3 Replacing protrusions

In the entire section we fix $c := c_H$, where c_H is the constant from Corollary 2.3, and $r := 6\|H\|c$. Moreover we assume that $x \geqslant 4\|H\|$; the precise value of x will be determined later.

Definition 4.7. Protrusion *is any (r,c)-boundaried interval $X = I \cup S^\pm$ which is H-free and $x(3c+k+1)$-long. For brevity we will refer to X as to I. The set $\Sigma^{(r,c)}(X)$ is the* signature *of the protrusion.*

Let T be a tournament equipped with a vertex ordering σ and let I be an H-free interval and such that $T = T^\circ \oplus X$, where X is a protrusion of signature Σ. Let $I \subseteq V(T)$ be the interval defined by the protrusion. Fix the coloring ξ of σ-backward arcs incident with I.

Recall that from Corollary 4.15 follows that for every (r,c)-boundaried co-interval T° compatible with X if $T^\circ \oplus X$ admits an optimal solution F of size not greater than k, then there are at most $2c$ arcs in F incident with X.

For every protrusion $X = I \cup S^\pm$ define a function $f_X\colon 2^{\mathcal{S}^{(r,c)}} \to \{0,1,2,\ldots,2c\} \cup \{\infty\}$ as follows: $f_X(S)$ is the minimum number of arcs in $A(I,I)$ needed to hit all partial immersions in I whose signatures belong to S, or ∞ if this number is greater than $2c$.

We introduce an equivalence relation $\sim$ on the set of all protrusions. Let $X = I \cup S^\pm$, $X' = I' \cup S'^\pm$. We say that $X \sim X'$ if:

- $S^\pm = S'^\pm$;
- $\Sigma^{(r,c)}(X) = \Sigma^{(r,c)}(X')$; and
- $f_X(S) = f_{X'}(S)$ for every $S \subseteq \mathcal{S}^{(r,c)}$.

Note that the number of equivalence classes of $\sim$ is finite and bounded by a constant depending only on H. This means that if in each class we pick a representative with the minimal number of vertices (call each such element a *small* protrusion), then all small protrusions will have size bounded from above uniformly by a constant s_H depending on the digraph H only.

Same arguments as in the proof of Lemma 4.9 give the following.

Lemma 4.16. *Given protrusions X and X' it can be decided in polynomial time whether $X \sim X'$.*

We now argue that equivalent protrusions are replaceable.

Lemma 4.17. *Suppose that $T = T^\circ \oplus X$ is a tournament satisfying the conclusion of Corollary 4.15, where X is a boundaried (r,c)-interval of signature Σ and T° is a c-boundaried co-interval which is compatible with X.*

Then for every X' such that $X \sim X'$, the optimal solution in $T^\circ \oplus X$ is of size not greater than k if and only if the optimal solution in $T^\circ \oplus X'$ is of size not greater than k.

Proof. We will prove the forward implication only; the proof of the other one is completely analogous due to the symmetry of the roles of X and X'.

Let F be an optimal solution in T and suppose that $|F| \leqslant k$. Let

$$F_0 = F - A(X, V(T)), \qquad F_1 = F \cap A(V(T) \setminus X, X) \quad \text{and} \quad F_X = F \cap A(X,X).$$

By the assumption that $T = T^\circ \oplus X$ satisfies the conclusion of Corollary 4.15, we have $|F_1 \cup F_X| \leqslant 2c$.

Claim 4.1. *Every X-generic arc of F_1 is incident with the r-boundary of X.*

Proof (of claim). Suppose for the sake of contradiction that in F_1 there exists an X-generic arc g not incident with the r-boundary of X. This means that this arc is incident with an *internal* vertex of X, i.e. the vertex being σ-smaller than all vertices in X_r^+ and σ-larger than all vertices in X_r^-.

By the optimality of F we conclude that in $T - (F - g)$ there exist immersion models of H, and moreover each such model uses the arc g. Let us choose such model ϕ with the property that the partial immersion $\phi|_X$ cannot be non-trivially shortened. We will show a way to "reroute" the generic connection containing g to avoid F, all X-singular arcs and all other generic arcs used in ϕ. This way we expose an immersion model of H in $T - F$, which is a contradiction.

This is the part of the proof where the r-boundary of X (and keeping track of the r-boundaries in the definition of types) comes into play. In the beginning of this section we picked r large enough to ensure the existence of a "pivot" vertex, which can be used for rerouting.

Suppose that $g \in \Gamma^+(X)$; the case $g \in \Gamma^-(X)$ is analogous. Denote $v = \mathsf{tail}(g)$. Note that at most $4c$ vertices of X_r^+ are incident with $F_1 \cup F_X$, as $|F_1 \cup F_X| \leqslant 2c$. Moreover $|\overrightarrow{A}(X_r^+, \{v\})| \leqslant c$, because X is c-flat. The number of vertices of X_r^+ incident with $\partial^-(X)$ also does not exceed c.

For every $a \in A(H)$ the number of arcs on $\phi|_X(a)$ with tail in v does not exceed 1 as otherwise $\phi|_X(a)$ could be non-trivially shortened contradicting minimality of $\phi|_X$. So the number of arcs in $\Gamma^+(X) \cap (F_1 \cup F_X \cup A(\phi|_X))$ incident with X_r^+ is not greater than $\|H\|$. Together this means that if

$$r > 6c + \|H\|,$$

which is true, then there exists a vertex $v_\omega \in X_r^+$ which is not incident with any of the special arcs mentioned above. The arc (v, v_ω) exists and is not contained in $F \cup A(\phi)$. Neither is the arc joining v_ω with the respective vertex in $T - X$. The arc g may therefore be replaced by a path composed of the two arcs described above. And so we have obtained an immersion model of H in $T - F$, a contradiction. □

Let S be the set of all short types τ such that every X-admissible partial immersion of type τ intersects with F_X. Because $f_X(S) = f_{X'}(S)$, we know that there exists a set $F_{X'}$ of arcs of $A(X')$ in $T' := T^\circ \oplus X'$ such that $|F_{X'}| \leqslant |F_X|$ and each partial immersion whose type is in S intersects with $F_{X'}$. Let $F_1' \subseteq A(T' - X', X')$ be the set directly corresponding to F_1 by a natural bijection that maps each singular arc of to the singular arc of the same color, and each generic arc to a generic arc whose the endpoint in T° is the same and the other endpoint has the same index in the respective r-boundary (by Claim 4.1, this endpoint must belong to the respective boundary). We will prove that the set $F' := F_0 \cup F_1' \cup F_{X'}$ is a solution in T', thereby completing the proof due to $|F'| \leqslant |F|$.

Suppose for the sake of contradiction that there exists an immersion model ϕ of H in $T' - F'$. Choose one such that $\phi|_{X'}$ is minimal. Therefore, by Lemma 4.5, $\phi|_{X'}$ is short. Moreover $\tau^{(r,c)}(\phi|_{X'}) \notin S$, so in T there exists a partial immersion ϕ_X in X which is disjoint with F and such that $\tau^{(r,c)}(\phi_X) = \tau^{(r,c)}(\phi|_{X'})$. We will show a way to reconstruct an immersion model of H in T from ϕ_X and $\phi|_{T^\circ}$ which will be disjoint with F and hence will give the desired contradiction. We may follow closely the strategy used in the proof of Lemma 4.7, taking an additional care about ensuring disjointness with F.

We know that ϕ_X is disjoint with F, so for sure the direct gluing along singular arcs is possible, as it was so with $\phi|_{X'}$ and ϕ_{T°. Moreover, gluing along generic arcs with one endpoint u on the r-boundary of X and the other one v in T° is possible since the arc uv is generic, not incident with F and will not have to be introduced more than once (this is “guarded” by the equivalence relations $R^\pm$). □

We are now set up with all tools needed to prove our main theorem.

Proof (of Theorem C). We prove that provided $|T| > C \cdot k^C$, where the constant C will be defined later, one can compute an instance (T', k) equivalent to (T, k) and satisfying $|T'| < |T|$. The conclusion will follow from applying such reduction (at most) $|T|$ times.

First of all note that from Theorem 2.2 for a digraph being a disjoint union of k exemplars of H, we conclude that if T does not contain k arc-disjoint immersion copies of H, then $\mathbf{ctw}(T) \leqslant c_0 k^2$ for some constant c_0 (depending on H). In particular, it follows that if $\mathbf{ctw}(T) > k^2 c_0$ (which can be established in polynomial time using Lemma 2.1), then T is a **no**-instance. So from now on we may assume that $\mathbf{ctw}(T) \leqslant c_0 k^2$.

Let $c = c_H$, $r = 6\|H\|c$, $x' = \max\{4\|H\|, s_H + 1\}$ and $x = x'(3c + k + 1)(2k + 3)^{3s(r,c)}$, where $s(r, c)$ is defined in Corollary 4.8. Let C be a constant satisfying $Ck^C \geqslant (2k^2 c_0 + 1)(x + 1)(k + 1)$, e.g. $C = \max\{3s(r, c) + 4, 5^{3s(r,c)} \cdot 3c_0 \cdot 4x' \cdot (6c + 2)\}$.

Suppose that T is a tournament satisfying $\mathbf{ctw}(T) \leqslant k^2 c_0$ and $|T| > Ck^C$. Applying Lemma 4.12, we either conclude that T admits more than k arc-disjoint copies of H (so (T, k) is a **no**-instance), or find an ordering σ of $V(T)$ and an H-free c-flat x-long σ-interval J. Both conclusions can be effectively gained in polynomial time.

In the latter case, we may use Corollary 4.11 in a manner described in Section 4.2 to find in J an H-free c-flat $x'(3c + k + 1)(2k + 3)$-long σ-interval I of (r, c)-signature Σ, which can be decomposed to $2k + 3$ consecutive $x'(3c + k + 1)$-long σ-intervals I_i, each of (r, c)-signature Σ. Both I and Σ are found in polynomial time.

Let $T = T^\circ \oplus X$ be the decomposition where T° is a c-boundaried co-interval and $X = (I_2 \cup S^\pm, \Sigma)$ is a protrusion corresponding to I_2. Clearly, X is compatible with T°. Using Lemma 4.16 we may check in polynomial time all small protrusions to find one X' such that $X' \sim X$. Let us define $T' = T^\circ \oplus X'$.

By Lemma 4.17 we conclude that (T', k) is an instance of H-hitting Immersions in Tournaments equivalent to (T, k). Moreover as $|X'| \leqslant s_H < |X|$, we have that $|T'| < |T|$. □

Part II

Immersions and Duality

Chapter 5

Objects Dual to Tree-cut Decompositions

The goal of this chapter is to introduce a variant of the definition of tree-cut width that is functionally equivalent to the existing ones, but which can be related to naturally defined dual objects: brambles and tangles.

In Section 5.1 we enrich the notation for the purpose of this chapter and discuss the definitions of tree-cut width of Wollan [Wol15] and of Giannopoulou et al. [GPR+21]. In Section 5.2 we introduce our notion of width of a tree-cut decomposition and define the dual objects: brambles and tangles. Section 5.3 is devoted to the proofs of Theorems D and E. In Section 5.4 we derive Theorem F: the game characterization of tree-cut width

Theorem D (restated). *For a graph G and a positive integer k, the following conditions are equivalent:*
- *G has no tree-cut decomposition of width $< k$;*
- *G has a bramble of order $\geqslant k$; and*
- *G has a tangle of order $\geqslant k$.*

Theorem E (restated). *For a graph G and positive integers a and b, the following conditions are equivalent:*
- *G has no tree-cut decomposition of adhesion-width $< a$ and bag-width $< b$;*
- *G has a bramble of adhesion-order $\geqslant a$ and bag-order $\geqslant b$; and*
- *G has a tangle of adhesion-order $\geqslant a$ and bag-order $\geqslant b$.*

Theorem F (restated). *A graph has a tree-cut decomposition with adhesion-width $\leqslant a$ and bag-width $\leqslant b$ if and only if there is a strategy to catch the robber using a cops and b dogs.*

Throughout this chapter we allow the existence of *parallel edges* (multiple edges with same pair of endpoints) but we do not allow loops (edges with both endpoints at the same vertex). Note in particular that this slightly modifies the meaning of the word *graph* in the statements of presented results.

5.1 Preliminaries

Edge separations. Let G be a graph. An *(edge) separation*[2] in G is a near-partition $\{A, B\}$ of $V(G)$ that has two elements, called the *sides*. The *order* of the separation $\{A, B\}$ is $|\delta(A)| = |\delta(B)|$. Separations of order at most 2 are called *thin*, whereas those of order at least 3 are called *bold*.

3-edge-connectedness. Two vertices u and v in a graph G are 3-*edge-connected* if there exist three edge-disjoint paths connecting u and v. By Menger's Theorem, the relation of 3-edge-connectedness is an equivalence relation on $V(G)$. We will denote this relation by $\sim_{3\mathrm{CC}}$; the graph G will always be clear from the context. The equivalence classes of $\sim_{3\mathrm{CC}}$ are called the 3-*edge-connected components* of G.

[2]We remark that most of literature in structural graph theory use term *separation* for a vertex separation, that is a pair $\{A, B\}$ of subsets of vertices with $A \cup B = V(G)$ and no edge between $A \setminus B$ and $B \setminus A$. Throughout this work we will solely work with edge separations as defined above, hence for brevity we call them simply *separations*.

A set of vertices $X \subseteq V(G)$ is 3-*edge-connected* in G if the vertices of X are pairwise 3-edge-connected in G. Equivalently, X is 3-edge-connected if it is entirely contained in a single 3-edge-connected component of G.

Suppose $\sim$ is an equivalence relation on the vertex set of G. We define the *quotient graph* ${}^{G}\!/\!\sim$ as follows. The vertices of ${}^{G}\!/\!\sim$ are the equivalence classes of $\sim$, and each edge uv of G with $u \not\sim v$ gives rise to one edge AB in ${}^{G}\!/\!\sim$, where A and B are the equivalence classes of $\sim$ to which u and v belong, respectively. Note that thus, the number of parallel edges in ${}^{G}\!/\!\sim$ connecting a pair A, B of equivalence classes of $\sim$ is equal to the number of edges in G whose one endpoint is in A and the other is in B. Also, edges of G whose endpoints belong to the same equivalence class of $\sim$ do not contribute to the edge set of ${}^{G}\!/\!\sim$.

In Subsection 5.3.2 we will study the structure of the quotient graph ${}^{G}\!/\!\sim_{3\mathrm{CC}}$, for any graph G.

Immersions. The notion of an undirected immersion is analogous to the directed one considered in part I. We say that a graph G admits a graph H as an *(unoriented) immersion* if there exists an *immersion model* of H in G: an embedding ϕ defined on the vertices and edges of H as follows:

- ϕ maps vertices of H to pairwise different vertices of G;
- ϕ maps each edge e of H with endpoints u and v to a path in G with endpoints $\phi(u)$ and $\phi(v)$; and
- paths in $\{\phi(e) \colon e \in E(H)\}$ are pairwise edge-disjoint.

Walls are central to the notions presented in this chapter. A $k \times k$ *wall* (cf. Figure 5.1) is a graph on $2k^2$ vertices, constructed from k (horizontal) paths $P_1, \ldots, P_k$. Each P_i has vertex set $\{v_1^i, \ldots, v_{2k}^i\}$ where v_j^i and v_{j+1}^i are adjacent for all $1 \leqslant j \leqslant 2k-1$, and there are the following additional edges between the paths:

- $v_j^i v_j^{i+1}$ if i, j are odd, $1 \leqslant i < k$, $1 \leqslant j \leqslant 2k$;
- $v_j^i v_j^{i+1}$ if i, j are even, $1 \leqslant i < k$, $1 \leqslant j \leqslant 2k$.

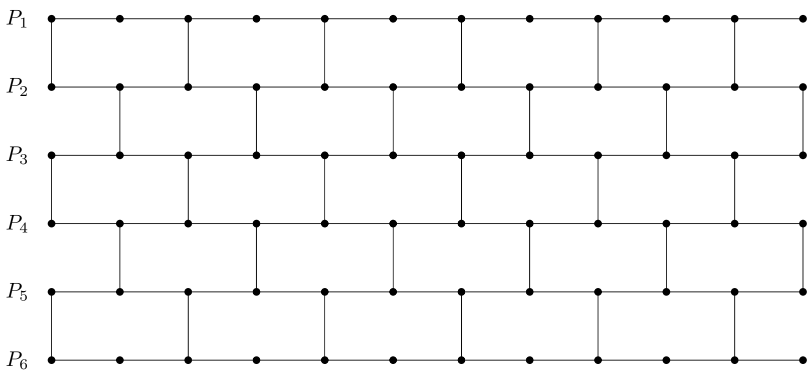

Figure 5.1: The 6×6 wall.

Tree-cut width. We now recall the concept of tree-cut width, as defined by Wollan [Wol15], and then give an equivalent definition proposed by Giannopoulou et al. [GPR+21]. First, we need a notion of a decomposition.

Definition 5.1 (Tree-cut decomposition). *A* tree-cut decomposition *of a graph G is a pair $\mathcal{T} = (T, \mathcal{X})$ such that T is a tree and $\mathcal{X} = \{X_t \subseteq V(G) \mid t \in V(T)\}$ is a near-partition of the vertex set of G, indexed by the nodes of T: for every node $t \in V(T)$, we call* bag *of t its associated set X_t in $\mathcal{X}$.*

We now introduce some useful definitions, which will eventually lead to a concept of the width of a tree-cut decomposition. Let us fix a tree-cut decomposition $\mathcal{T} = (T, \mathcal{X})$ of a graph G. For a pair of vertices u, v of G (not necessarily distinct), the *trace* of $\{u, v\}$ in $\mathcal{T}$ is the (unique) path in T connecting the node s satisfying $u \in X_s$, and the node t satisfying $v \in X_t$. Note that if $u = v$, then the trace of $\{u, v\}$ consists of only one node, the one whose bag contains u. The trace of an edge of G is the trace of its endpoints.

For an edge st of T, the *adhesion* of st, denoted $\mathrm{adh}(st)$, is the set of all edges of G whose traces contain st. Equivalently, an edge $uv \in E(G)$ belongs to $\mathrm{adh}(st)$ if and only if the node s' satisfying $u \in X_{s'}$ and the node t' satisfying $v \in X_{t'}$ lie in different connected components of $T - st$. An edge st of T is called *thin* if $|\,\mathrm{adh}(st)| \leqslant 2$, and *bold* otherwise. In the notation $\mathrm{adh}(\cdot)$, the decomposition $\mathcal{T}$ can be specified in the subscript if it is not clear from the context.

Next, for each node t of T we define the *torso* G_t. Intuitively, G_t is obtained from G by identifying, for every connected component C of $T - t$, all the vertices residing in the bags of C into a single vertex. Formally, G_t is defined as the quotient graph ${}^{G}/{\sim_t}$, where $\sim_t$ is an equivalence relation on $V(G)$ defined as follows: $u \sim_t v$ if either $u = v$ or $u \neq v$ and the trace of $\{u, v\}$ in $\mathcal{T}$ does not contain t.

Finally, the 3-*center* $\widetilde{G}_t$ of the torso G_t is obtained from G_t by iteratively suppressing vertices of degree at most 2 in G_t, but only those that do not belong to X_t. Here, to *suppress* a vertex of degree at most 2 means to either delete it, provided it had at most one neighbor, or delete it and add an edge connecting its two former neighbors, provided it had exactly two neighbors. It is not hard to see that the order of performing the suppressions does not matter. However, note that the suppression of a vertex of degree 2 can reduce the degree of another vertex outside of X_t, leading in turn to its suppression.

With all these notions in place, we can recall the definition of tree-cut width originally proposed by Wollan in [Wol15].

Definition 5.2 (Wollan's tree-cut width). *The* Wollan's width *of a tree-cut decomposition* $\mathcal{T} = (T, \mathcal{X})$ *of a graph* G *is defined as*

$$\max\left(\{|\operatorname{adh}(e)|\colon e \in E(T)\} \cup \{|V(\widetilde{G}_t)|\colon t \in V(T)\}\right).$$

The Wollan's tree-cut width *of* G *is the minimum width of a tree-cut decomposition of* G.

In order to simplify arguments in their study of algorithmic problems related to immersions, Giannopoulou et al. introduced in [GPR+21] an alternative definition of tree-cut width. Let G be a graph and $\mathcal{T} = (T, \mathcal{X})$ be a tree-cut decomposition of G. They define, for every node t of T, the following quantity:

$$w_t \coloneqq |X_t| + |\{t' \in N_T(t)\colon |\operatorname{adh}(tt')| \geqslant 3\}|.$$

They then propose the following adjustment of the definition of tree-cut width, which they show to be equivalent to Wollan's tree-cut width.

Definition 5.3 (GPRTW's tree-cut width). *The* GPRTW's width *of a tree-cut decomposition* $\mathcal{T} = (T, \mathcal{X})$ *of a graph* G *is defined as*

$$\max\left(\{|\operatorname{adh}(e)|\colon e \in E(T)\} \cup \{w_t\colon t \in V(T)\}\right).$$

The GPRTW's tree-cut width *of* G *is the minimum width of a tree-cut decomposition of* G.

Theorem 5.1 ([GPR+21]). *For every graph G, the GPRTW's tree-cut width of G and the Wollan's tree-cut width of G are equal.*

We point out that the equality described in Theorem 5.1 does not apply for every single tree-cut decomposition separately: there are tree-cut decompositions whose Wollan's width and GPRTW's width differ. An example can be obtained by taking a complete binary tree of depth $\geqslant 3$ as the graph G, and constructing a tree-cut decomposition $\mathcal{T}$ of G whose underlying tree is a star. The bag of the center of $\mathcal{T}$ consists of the unique vertex of G of degree 2, while every other vertex of G is placed in a different leaf bag of $\mathcal{T}$. It is easy to verify that $\mathcal{T}$ has Wollan's width 3, while its GPRTW's width is $\frac{|V(G)|-1}{2}$. There is no contradiction between this example and the statement of Theorem 5.1: the proof provided in [GPR+21] applies a modification of the given decomposition in order to get the desired bound.

Finally, let us recall the main result of [Wol15]: a grid theorem for tree-cut width and immersions. By Theorem 5.1, we may state it equivalently in terms of GPRTW's tree cut-width and in terms of Wollan's tree-cut width.

Theorem 5.2 ([Wol15]). *There exists a function $f\colon \mathbb{N} \to \mathbb{N}$ such that every graph that does not contain a $k \times k$ wall as an immersion has Wollan's (or GPRTW's) tree-cut width at most $f(k)$.*

5.2 Objects

In this section we introduce our definition of tree-cut width and relate it to Wollan's definition. Next, we define the dual objects — brambles and tangles — and state our main result.

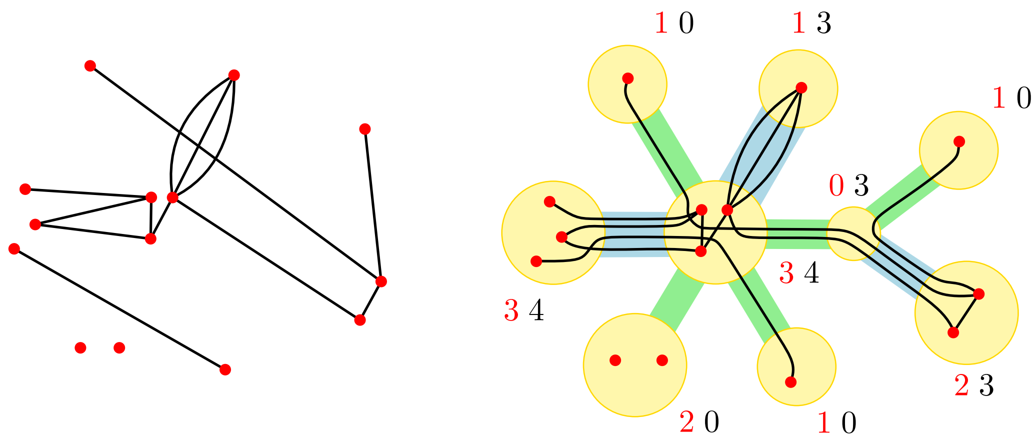

Figure 5.2: A tree-cut decomposition (right) of a graph (left). Bold adhesions are depicted in blue, while thin adhesions are in green. Red numbers stand for sizes of bags, and black number for sizes of adhesions.

5.2.1 Tree-cut decompositions

The idea is to use the same notion of tree-cut decompositions, as introduced in Definition 5.1, but to redefine the width. We do it as follows.

Let $\mathcal{T} = (T, \mathcal{X})$ be a tree-cut decomposition of a graph G. The *adhesion* of a node $t \in V(T)$ is defined as the union of adhesions of bold edges incident to t, namely,

$$\mathrm{adh}(t) := \bigcup \{\mathrm{adh}(st) \colon s \text{ is such that } st \in E(T) \text{ and } |\,\mathrm{adh}(st)| \geqslant 3\}\,.$$

We point that even if an edge e of G participates in the adhesions of two bold edges of T incident to t, e is counted only once when computing the cardinality of the adhesion of t. Also note that since only bold edges contribute to the adhesion of a node, for each $t \in V(T)$ we have either $|\,\mathrm{adh}(t)| = 0$ or $|\,\mathrm{adh}(t)| \geqslant 3$. This is illustrated in Figure 5.2.

We now present our concept of width. The key aspect is that we will work with two separate width measures, respectively corresponding to the adhesions and to the bags.

Definition 5.4 (Adhesion-width and bag-width). *Let $\mathcal{T} = (T, \mathcal{X})$ be a tree-cut decomposition of a graph G. The* adhesion-width *of $\mathcal{T}$ and the* bag-width *of $\mathcal{T}$ are respectively defined as*

$$\mathbf{aw}(\mathcal{T}) = \max_{t \in V(T)} |\,\mathrm{adh}(t)| \qquad \textit{and} \qquad \mathbf{bw}(\mathcal{T}) = \max_{t \in V(T)} |X_t|.$$

Obviously, we can combine the two width measures into one by taking the maximum, and thus we arrive at our proposition for the notion of tree-cut width.

Definition 5.5 (ab-tree-cut width). *The* ab-tree-cut width *of a graph G is the least number k such that G has a tree-cut decomposition with adhesion-width $\leqslant k$ and bag-width $\leqslant k$.*

Thus, we have now three notions of tree-cut width: Wollan's tree-cut width (Definition 5.2) and GPRTW's tree-cut width (5.3) that are equivalent, and the ab-tree-cut width (Definition 5.5). Let us note that, similarly to Wollan's definition, the new notions of width are closed under taking immersions.

Proposition 5.3. *Let G be a graph that contains a graph H as an immersion. Suppose G admits a tree-cut decomposition with adhesion-width $\leqslant a$ and bag-width $\leqslant b$, for some positive integers a, b. Then so does H.*

Proof. Let $\mathcal{T} = (T, \mathcal{X})$ be a tree-cut decomposition of G with $\mathbf{aw}(\mathcal{T}) \leqslant a$ and $\mathbf{aw}(\mathcal{T}) \leqslant b$. Fix an immersion model ϕ of H in G and for each $t \in V(T)$ define $Y_t := \phi^{-1}(X_t) \subseteq V(H)$ to be the set of vertices of H whose images are contained in the bag X_t. We prove that $\mathcal{T}' := (T, \{Y_t\}_{t \in V(T)})$ is a tree-cut decomposition of H with the desired properties.

Note that $\{Y_t\}_{t \in V(T)}$ is a near-partition of $V(H)$. As ϕ maps vertices of H injectively, we have $|Y_t| \leqslant |X_t|$ for every $t \in V(T)$, which implies $\mathbf{bw}(\mathcal{T}') \leqslant \mathbf{bw}(\mathcal{T})$. Observe that the trace in $\mathcal{T}'$ of an edge $e \in E(H)$ with endpoints u and v is contained in the union of traces in $\mathcal{T}$ of edges of the path $\phi(e)$.

Moreover, for every edge $st \in E(T)$, if e is an edge of H in $\mathrm{adh}_{\mathcal{T}'}(st)$ with endpoint u and v, then there is at least one edge f on the path $\phi(e)$ with $f \in \mathrm{adh}_{\mathcal{T}}(st)$.

Consider a partial function $\eta\colon E(G) \rightharpoonup E(H)$ defined as follows: if $f \in E(G)$ belongs to $\phi(e)$ for some $e \in E(H)$, then we set $\eta(f) \coloneqq e$, and otherwise f is not in the domain of η. Note that the validity of this definition is asserted by the fact that the paths in $\{\phi(e)\colon e \in E(G)\}$ are pairwise edge-disjoint. With this notion in place, the observations from the previous paragraph show that

$$\mathrm{adh}_{\mathcal{T}'}(st) \subseteq \eta\left(\mathrm{adh}_{\mathcal{T}}(st)\right) \qquad \text{for every } st \in E(T).$$

This implies that

$$\mathrm{adh}_{\mathcal{T}'}(t) \subseteq \eta\left(\mathrm{adh}_{\mathcal{T}}(t)\right) \qquad \text{for every } t \in V(T).$$

It follows that $\mathbf{aw}(\mathcal{T}') \leqslant \mathbf{aw}(\mathcal{T})$, as required. □

5.2.2 Comparison with original tree-cut width

We now verify that our new definition is functionally equivalent to the one of Wollan. Formally, we prove the following statement which involves the adjusted definition of tree-cut width proposed by Giannopoulou et al.

Theorem 5.4. *Let G be a graph and $\mathcal{T}$ be a tree-cut decomposition of G. Then:*

- *If $\mathcal{T}$ has GPRTW's width k, then $\mathcal{T}$ has adhesion-width $\leqslant k^2$ and bag-width $\leqslant k$.*
- *If $\mathcal{T}$ has adhesion-width a and bag-width b, then $\mathcal{T}$ has GPRTW's width $\leqslant a + b + 2$.*

Proof. We show the first implication. Let us assume that $\mathcal{T}$ has GPRTW's width k. In particular, we have $\max\{|\mathrm{adh}(e)|\colon e \in E(T)\} \leqslant k$ and $\max\{w_t\colon t \in V(T)\} \leqslant k$. Since by definition $w_t \geqslant |X_t|$ for any $t \in V(T)$, we immediately obtain that $\mathcal{T}$ has bag-width $\leqslant k$. Concerning the adhesion-width, note that $|\{t' \in N_T(t)\colon |\mathrm{adh}(tt')| \geqslant 3\}| \leqslant k$ for every $t \in V(T)$. Since in addition $|\mathrm{adh}(tt')| \leqslant k$ for every $tt' \in E(T)$, we get that $\mathcal{T}$ has adhesion-width $\leqslant k^2$.

We now show the other implication. Suppose that $\mathcal{T}$ has adhesion-width a and bag-width b. If $a = 0$ then $\mathcal{T}$ only contains thin adhesions, and we deduce that GPRTW's width is bounded by the size of a largest bag in $\mathcal{T}$, plus the size of a largest thin adhesion in $\mathcal{T}$. Hence $\mathcal{T}$ has GPRTW's width $\leqslant b+2$ in that case. Otherwise by definition, $a \geqslant 3$. Clearly, $|\mathrm{adh}(e)| \leqslant a$ for every $e \in E(T)$ in that case. Since for every $t \in V(T)$, w_t is bounded by the size of X_t plus the size of the adhesion of t, we conclude that GPRTW's width is bounded by the maximum of the following two quantities: the largest size of an adhesion of an edge in $\mathcal{T}$, and the largest cumulated size of X_t and the adhesion of t, for a node t in $\mathcal{T}$. This is bounded by $\max(a, a + b)$, hence the conclusion. □

By combining Theorem 5.4 and Theorem 5.1, we immediately get the following.

Corollary 5.5. *Let G be a graph and let k and ℓ be the Wollan's tree-cut width and the ab-tree-cut width of G, respectively. Then*

$$\frac{k}{2} - 1 \leqslant \ell \leqslant k^2.$$

From Corollary 5.5 we infer that Theorem 5.2 is also true for ab-tree-cut width.

5.2.3 Brambles

We now move to the first definition of a dual object: a bramble. First, we introduce *slabs*, which are elements from which the brambles are composed.

Definition 5.6 (Slabs). *A* slab *in a graph G is a pair (H, K) where H is a connected subgraph of G and K, called the* core *of the slab, is a non-empty subset of vertices of H that is 3-edge-connected in G. Two slabs (H, K) and (H', K')* touch *if they intersect on their cores, that is, $K \cap K' \neq \emptyset$.*

With slabs in place, brambles are defined as follows.

Definition 5.7 (Brambles). *A* bramble *in a graph is a family of pairwise touching slabs.*

Note that since 3-edge-connectedness is transitive, all cores of all the slabs in a bramble must be contained in a single 3-connected component of the graph G.

Next, we need to define the order of a bramble. Mirroring the situation in Definition 5.4, there will be two notions of an order: one corresponding to adhesions and one corresponding to the bags. For the first one, we need an appropriate notion of hitting a slab.

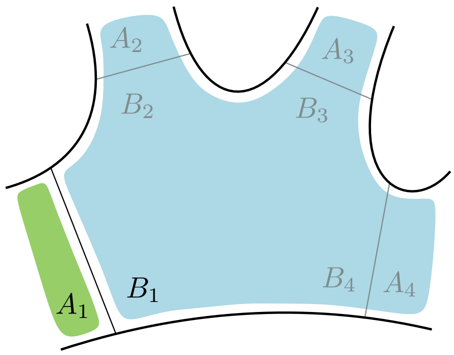

Figure 5.3: The separations $(A_1, B_1), \ldots, (A_4, B_4)$ form a star; A_1 is in green and B_1 in blue.

Definition 5.8 (Disconnecting sets). *Let G be a graph and (H, K) be a slab in G. A set of edges $F \subseteq E(G)$* disconnects *(H, K) if there are two vertices $u, v \in K$ that are disconnected by $F \cap E(H)$ in H. Further, F is a* disconnecting set *for a bramble $\mathcal{B}$ if F disconnects every slab in $\mathcal{B}$.*

Note that if the core K of a slab (H, K) has size 1, then there is no edge subset that disconnects (H, K). We now proceed with defining the order(s) of a bramble.

Definition 5.9 (Orders of a bramble). *Let $\mathcal{B}$ be a bramble in a graph G. The* adhesion-order *of $\mathcal{B}$ is the minimum size of a disconnecting set for $\mathcal{B}$, or $+\infty$ if no such disconnecting set exists. The* bag-order *of $\mathcal{B}$ is defined as $\min_{(H,K)\in\mathcal{B}} |K|$, or $+\infty$ if $\mathcal{B}$ is empty. The* order *of $\mathcal{B}$ is the minimum of the adhesion-order and the bag-order of $\mathcal{B}$.*

5.2.4 Tangles

We now move to the second definition of a dual object: a tangle. For this, we need to take a closer look at (edge) separations. The following definitions are essentially taken from the presentation in the book of Diestel [Die16], which in turns cites the work of Diestel and Oum [DO21] as the source of inspiration.

Let us fix a graph G. Consider a separation $\{A, B\}$ of G. With such a separation we can associate two *oriented separations* (A, B) and (B, A). We say that the oriented separation (A, B) *points toward* B. For a set $\mathcal{S}$ of separations, define $\vec{\mathcal{S}} \coloneqq \{(A, B), (B, A) \colon \{A, B\} \in \mathcal{S}\}$ as the set of all oriented separations associated with the elements of S. An *orientation* of a set $\mathcal{S}$ of separations is a subset $\mathcal{L} \subseteq \vec{\mathcal{S}}$ which contains precisely one oriented separation associated with every element of $\mathcal{S}$. We say that $\mathcal{L}$ *avoids* a collection Σ of sets of oriented separations if no subset of $\mathcal{L}$ belongs to Σ.

A set σ of oriented separations is *consistent* if it does not simultaneously contain separations (A, B) and (C, D) such that $B \cap D = \emptyset$. Intuitively, this means that there are no two separations in σ that "point away" from each other. A non-empty consistent set of oriented separations σ is a *star* if $A \cap C = \emptyset$ for all distinct $(A, B), (C, D) \in \sigma$. A star is illustrated in Figure 5.3.

For a positive integer a, we let $\mathcal{S}_a$ be the set of all separations of order $< a$ in G. Further, for a positive integer b, we let $\Sigma_{a,b}$ be the set of all stars $\sigma \subseteq \vec{\mathcal{S}}_a$ satisfying the following conditions:

$$\left|\bigcup\{\delta(A) \colon (A, B) \in \sigma \text{ and } |\delta(A)| \geqslant 3\}\right| < a,$$
$$\left|\bigcap\{B \colon (A, B) \in \sigma\}\right| < b.$$

We can now present the definition of a tangle.

Definition 5.10 (Tangles). *For a pair of positive integers a and b, an (a, b)-*tangle *is a consistent orientation of $\mathcal{S}_a$ that avoids $\Sigma_{a,b}$. We will say that an (a, b)-tangle has* adhesion-order a, bag-order b, *and* order $\min(a, b)$.

5.2.5 Main result

All the pieces are now set and we can state our main result.

Theorem E (restated). *For any graph G and positive integers a and b, the following are equivalent:*

(A1) G has a bramble of adhesion-order $\geqslant a$ and bag-order $\geqslant b$;

(A2) G has a tangle of adhesion-order $\geqslant a$ and bag-order $\geqslant b$;

(A3) G has no tree-cut decomposition of adhesion-width $< a$ and bag-width $< b$.

When cast to the variants of definitions with a single parameter, Theorem E takes the following form.

Theorem D (restated). *For any graph G and a positive integer k, the following are equivalent:*

(B1) G has a bramble of order $\geqslant k$;

(B2) G has a tangle of order $\geqslant k$;

(B3) G has ab-tree-cut width $\geqslant k$.

5.3 Proof of Theorem E

In this section we prove Theorem E. The outline is as follows. First, we verify that all the three statements always hold for $b = 1$, which proves the theorem in this case. Hence, for the remainder of the proof we may assume that $b \geqslant 2$. Then, in successive subsections we prove implications (A1)⇒(A2), (A2)⇒(A3), and (A3)⇒(A1), in this order. The first two implications are rather straightforward, while the main weight of the proof lies in the last implication. In particular, the subsection containing its proof is preceded by an analysis of the structure of 3-edge-connected components in an arbitrary graph.

Proof (of Theorem E for $b = 1$). We show that when $b = 1$, all three conditions (A1), (A2), and (A3) hold for every positive integer a.

For (A1), fix some vertex v and consider the bramble consisting of one slab (H, K) where $K = \{v\}$ and $H = G[\{v\}]$. This bramble has bag-order 1 and adhesion-order $+\infty$. Hence (A1) holds for $b = 1$ and any positive integer a.

For (A2), we argue similarly. Let a be any positive integer. Fix any vertex v and consider the orientation $\mathcal{L}$ of $\mathcal{S}_a$ defined as follows: for $\{A, B\} \in \mathcal{S}_a$, we include (A, B) in $\mathcal{L}$ provided $v \in B$, and otherwise we include (B, A). Clearly $\mathcal{L}$ is consistent and avoids $\Sigma_{a,1}$, because every separation in $\mathcal{L}$ points towards the side that contains v. So $\mathcal{L}$ is an $(a, 1)$-tangle, as required.

Finally, (A3) trivially holds for $b = 1$ and any positive integer a, because G is assumed to be non-empty. □

Therefore, for the remainder of the proof we assume that $b \geqslant 2$.

5.3.1 Easy implications

Proof (of Theorem E, (A1)⇒(A2)). Let $\mathcal{B}$ be a bramble of adhesion-order $\geqslant a$ and bag-order $\geqslant b$, where $b \geqslant 2$.

We claim that for each separation $\{A, B\}$ in $\mathcal{S}_a$, exactly one of the sides A and B contains the core K of at least one slab $(H, K) \in \mathcal{B}$ as a subset. Indeed, if none of A and B contained such a core, then $\delta(A)$ would be a disconnecting set for $\mathcal{B}$. Since $|\delta(A)| < a$, this is a contradiction. On the other hand, if both A and B contained cores of some slabs from $\mathcal{B}$, then these cores would not intersect, again a contradiction.

We construct a tangle $\mathcal{L}$ in G by orienting every separation $\{A, B\} \in \mathcal{S}_a$ so that it points to the part that contains the core of at least one slab from $\mathcal{B}$. Note that $\mathcal{L}$ is consistent, because otherwise $\mathcal{B}$ would contain two slabs with disjoint cores. It remains to prove that $\mathcal{L}$ avoids $\Sigma_{a,b}$.

For contradiction, suppose that there exists a star $\sigma \subseteq \mathcal{L}$ such that

$$\left|\bigcup\{\delta(A) \colon (A, B) \in \sigma \text{ and } |\delta(A)| \geqslant 3\}\right| < a, \tag{5.1}$$

$$\left|\bigcap\{B \colon (A, B) \in \sigma\}\right| < b. \tag{5.2}$$

Let

$$F := \bigcup\{\delta(A) \colon (A, B) \in \sigma \text{ and } |\delta(A)| \geqslant 3\}.$$

Consider any slab $(H, K) \in \mathcal{B}$. Since the bag-order of $\mathcal{B}$ is at least b, we have $|K| \geqslant b$. By (5.2), there exists a separation $(A, B) \in \sigma$ such that A and K intersect. Note that it cannot happen that $K \subseteq A$, because $(A, B) \in \mathcal{L}$, so A cannot fully contain any core of a slab from $\mathcal{B}$. We infer that $\delta(A)$ disconnects the slab (H, K). Since K is 3-edge-connected in G, this in particular implies that $|\delta(A)| \geqslant 3$. So $\delta(A) \subseteq F$ and F disconnects (H, K). Since the slab (H, K) was chosen arbitrarily from $\mathcal{B}$, we conclude that F is a disconnecting set for $\mathcal{B}$. However, now (5.1) stands in contradiction with the assumption that the adhesion-order of $\mathcal{B}$ is at least a. □

Proof (of Theorem E, (A2)⇒(A3)). Suppose toward a contradiction that G admits a tree-cut decomposition $\mathcal{T} = (T, \mathcal{X})$ with $\mathbf{aw}(\mathcal{T}) < a$ and $\mathbf{bw}(\mathcal{T}) < b$, as well as a tangle $\mathcal{L}$ with adhesion-order $\geqslant a$ and bag-order $\geqslant b$.

First, note that for every edge $ab \in E(T)$ we can define a separation $\{A, B\}$ of G, where A consists of all vertices u such that the node of T whose bag contains u lies in the same connected component of $T - ab$ as a, and $B = V(G) \setminus A$. Since $\mathbf{aw}(\mathcal{T}) < a$, it follows that the order of $\{A, B\}$ is smaller than a, that is, $\{A, B\} \in \mathcal{S}_a$. Hence, exactly one of (A, B) and (B, A) belongs to $\mathcal{L}$. Define an orientation $\vec{T}$ of the tree T as follows: orient ab towards b if $(A, B) \in \mathcal{L}$, and towards a otherwise. Since T has less edges than nodes, it follows that in $\vec{T}$ there exists a node r whose outdegree in $\vec{T}$ is zero. In other words, every edge incident to r in T points toward r in $\vec{T}$. Thus, the set σ of oriented separations corresponding to the edges incident to r forms a star in $\mathcal{S}_a$. Since T has adhesion-width $< a$ and bag-width $< b$, we immediately get that $\sigma \in \Sigma_{a,b}$. Hence $\mathcal{L}$ does not avoid $\Sigma_{a,b}$, a contradiction. □

5.3.2 Structure of 3-edge-connected components

Before we proceed to the proof of the last implication, we need to prove some auxiliary results about tree-cut decompositions of a graph and its 3-edge-connected components.

Consider a graph G and let $\sim_{3\mathrm{CC}}$ be the equivalence relation on $V(G)$ defined as being in the same 3-edge-connected component. We define

$$G_{3\mathrm{CC}} \coloneqq {}^{G}/_{\sim_{3\mathrm{CC}}}.$$

Our goal now is to understand the structure of $G_{3\mathrm{CC}}$. First, we observe that $G_{3\mathrm{CC}}$ has no non-trivial 3-edge-connected components.

Lemma 5.6. *In $G_{3\mathrm{CC}}$, no two different vertices are 3-edge-connected.*

Proof. Consider any two different vertices A and B of $G_{3\mathrm{CC}}$. Recall that A and B are two different 3-edge-connected components of G, hence let us pick arbitrary vertices $a \in A$ and $b \in B$. Since a and b belong to different 3-edge-connected components of G, there exists a separation $\{A', B'\}$ of G of order < 3 such that $a \in A'$ and $b \in B'$. Note that since the order of $\{A', B'\}$ is < 3, every 3-edge-connected component C of G has to be entirely contained either in A' or in B'. In particular, $A \subseteq A'$ and $B \subseteq B'$. Thus, $\{A', B'\}$ naturally induces a separation $\{\widehat{A}, \widehat{B}\}$ of $G_{3\mathrm{CC}}$ defined by placing each $C \in V(G_{3\mathrm{CC}})$ in $\widehat{A}$ provided $C \subseteq A'$, and in $\widehat{B}$ provided $C \subseteq B'$. Clearly, $A \in \widehat{A}$, $B \in \widehat{B}$, and the order of $\{\widehat{A}, \widehat{B}\}$ is the same as of $\{A', B'\}$. This means that $\{\widehat{A}, \widehat{B}\}$ witnesses that A and B are not 3-edge-connected in $G_{3\mathrm{CC}}$. □

It appears that the structure of graph satisfying the condition stated in Lemma 5.6 can be nicely described: they are *cacti*. More precisely, a *cactus* is a graph where every 2-(vertex)-connected component is either a single edge or a cycle (here, we allow cycles of length 2, that is, pairs of parallel edges). We have the following observation, which is essentially (up to technical details in definitions) known in the literature [Din76, MNS17].

Lemma 5.7. *A graph G is a cactus if and only if no two different vertices of G are 3-edge-connected.*

Proof. Assume first that G is a cactus. Take any two different vertices u, v of G and for contradiction suppose that there are three edge-disjoint paths P_1, P_2, P_3 connecting u and v. Let u_1, u_2, u_3 be the vertices directly succeeding u on P_1, P_2, P_3, respectively. Within the closed walk $P_1 \cup P_2$ one can find a cycle that contains u, u_1, and u_2, which witnesses that all these three vertices belong to the same 2-connected component of G. The same can be argued about the triple u, u_2, and u_3, so all the four vertices u, u_1, u_2, and u_3 belong to the same 2-connected component of G. However, the first edges of P_1, P_2, P_3 are pairwise different, which means that u has degree at least 3 within this 2-connected component. This contradicts the assumption that G is a cactus.

Assume now that G is not a cactus, which means that there exists a 2-connected component C of G that is neither a single edge nor a cycle. This means that C has a vertex of degree 3, say u. Let e_1, e_2, e_3 be an arbitrary triple of distinct edges of C incident to u, and let u_1, u_2, u_3 be the other endpoints of

e_1, e_2, e_3, respectively. Since C is 2-connected, $C - u$ is connected, hence in $C - u$ there exists a tree T whose set of leaves is $\{u_1, u_2, u_3\}$. Now T with vertex u and edges e_1, e_2, e_3 added is a graph consisting of u, another vertex v of degree 3 (a leaf of T if two of u_1, u_2, u_3 are equal, or an internal node of T otherwise), and 3 internally vertex-disjoint paths connecting u and v. Since this graph is a subgraph of G, it follows that u and v are 3-edge-connected in G. □

By combining Lemma 5.6 and Lemma 5.7 we get the following.

Corollary 5.8. *For any graph G, the graph $G_{3\mathrm{CC}}$ is a cactus.*

We now relate the ab-tree-cut width of a graph G with the ab-tree-cut width of its 3-edge-connected components. For this, we need to associate with each 3-edge-connected component A a suitable *torso* of A, which is a graph that reflects connections between vertices of A that are realized either by edges within A or paths that are internally disjoint with A. Formally, the torso of A is the graph $\mathsf{torso}(A)$ obtained from G as follows. Observe that for every connected component Z of $G - A$, there are at most two edges in G having one endpoint in Z and the other in A. Then, for every such component Z,

- remove Z completely, provided Z has at most one neighbor in A; or
- remove Z and replace it with a new edge f_Z connecting the neighbors of Z in A, provided Z has exactly two neighbors in A.

Note that this second operation does not create loops. The edge f_Z is called the *replacement edge* of the component Z.

We have the following simple observations about the torso operation.

Lemma 5.9. *For every graph G and a 3-edge-connected component A of G, the graph G contains $\mathsf{torso}(A)$ as an immersion.*

Proof. It suffices to map every vertex of A to itself, every edge of $G[A]$ to itself, and every replacement edge f_Z to any path in G that connects the endpoints of f_Z and has all the internal vertices in Z. □

Lemma 5.10. *For every graph G and a 3-edge-connected component A of G, the graph $\mathsf{torso}(A)$ is 3-edge-connected.*

Proof. Consider any pair of vertices $u, v \in A$. Since A is 3-edge-connected in G, there are three edge-disjoint paths P_1, P_2, P_3 in G that connect u and v. These can be naturally projected to paths Q_1, Q_2, Q_3 in $\mathsf{torso}(A)$ as follows: for every maximal subpath of P_i, $i \in \{1, 2, 3\}$, whose all internal vertices do not belong to A, say they belong to some connected component Z of $G - A$, replace this subpath with the replacement edge f_Z. It can be easily seen that each connected component Z of $G - A$ will participate in at most one such replacement, because there are at most two edges connecting Z with A. Hence, Q_1, Q_2, Q_3 remain edge-disjoint and u and v are 3-edge-connected in $\mathsf{torso}(A)$. □

The next theorem is the main outcome of this section. Intuitively, it will allow us to focus on a single 3-edge-connected component of G when constructing a bramble of high order. A statement in the work of Wollan that mirrors this step is [Wol15, Lemma 5].

Theorem 5.11. *Let G be a graph and a, b be two positive integers. Then G admits a tree-cut decomposition of adhesion-width $\leqslant a$ and bag-width $\leqslant b$ if and only if for every 3-edge-connected component A of G, the same can be said about the graph $\mathsf{torso}(A)$.*

Proof. The forward implication follows by combining Lemma 5.9 with Proposition 5.3. We are left with proving the backward implication.

Assume then that for every 3-edge-connected component A of G, there is a tree-cut decomposition $\mathcal{T}^A = (T^A, \mathcal{X}^A)$ of $\mathsf{torso}(A)$ such that $\mathbf{aw}(\mathcal{T}^A) \leqslant a$ and $\mathbf{bw}(\mathcal{T}^A) \leqslant b$. The goal is to to "glue" the decompositions $\mathcal{T}^A$ into a single tree-cut decomposition $\mathcal{T}$ of G so that the guarantees about the width measures are preserved. The gluing will be done along the graph $G_{3\mathrm{CC}}$, which by Corollary 5.8 is a cactus.

We execute the gluing as follows. Fix any spanning forest of $G_{3\mathrm{CC}}$ and let S be its edge set. Further, let $R = E(G_{3\mathrm{CC}}) \setminus S$ be the remaining edges of $G_{3\mathrm{CC}}$. Note that R contains exactly one edge from every 2-connected component C of $G_{3\mathrm{CC}}$ that is a cycle (that is, is not a single edge). Call this edge r_C.

Recall that in the construction of $G_{3\mathrm{CC}} = {}^{G}/_{\sim_{3\mathrm{CC}}}$, every edge of $G_{3\mathrm{CC}}$ originates in some edge of G that connects two different 3-edge-connected components. Let $\alpha\colon E(G_{3\mathrm{CC}}) \to E(G)$ be this origin mapping: $\alpha(e)$ is the edge of G from which e originates. Note that α is injective.

We construct a forest T from trees T^A as follows. First, take the disjoint union of trees T^A. Then, consider every edge $e \in S$, say with endpoints $A, B \in V(G_{3\mathrm{CC}})$. Let u and v be the endpoints of $\alpha(e)$, where $u \in A$ and $v \in B$. Clearly, there are nodes $p \in V(T^A)$ and $q \in V(T^B)$ such that $u \in X^A_p$ and $v \in X^B_q$. Then add the edge pq to the forest T, and call this edge $\gamma(e)$. This concludes the construction of T. It is easy to see that since S is a forest and each T^A is a tree, T is also a forest.

We now associate the nodes of T with bags $\mathcal{X} = \{X_s : s \in V(T)\}$ inherited from decomposition $\mathcal{T}^A$ in the natural manner: if a node s of T originates from the tree T^A, then we set $X_s \coloneqq X^A_s$. Since 3-edge-connected components of G form a partition $V(G)$, it follows that $\mathcal{X}$ is a near-partition of $V(G)$. Thus, $\mathcal{T} \coloneqq (T, \mathcal{X})$ is[3] a tree-cut decomposition of G.

It remains to analyze the width measures of $\mathcal{T}$. Since the bags are directly taken from decompositions $\mathcal{T}^A$, which have bag-width $\leqslant b$, we immediately see that $\mathbf{bw}(\mathcal{T}) \leqslant b$. The argument for the adhesion-width is a bit more complicated, because the adhesions may actually change during gluing.

We first show the following claim about connectedness in $G_{3\mathrm{CC}}$ and in G.

Claim 5.1. *Let A be a 3-edge-connected component of G and let C be a connected component of $G_{3\mathrm{CC}} - A$. Then the vertices of $\bigcup_{D \in V(C)} D$ lie in the same connected component of $G - A$.*

Proof. Let $\widehat{C} \coloneqq \bigcup_{D \in V(C)} D$. Suppose for contradiction that there is a partition $\{L, R\}$ of $\widehat{C}$ such that in G there is no edge with one endpoint in L and second in R and L, R are non-empty. Since C is connected in $G_{3\mathrm{CC}}$, there must exist a 3-edge-connected component $B \in V(C)$ such that both $L \cap B$ and $R \cap B$ are non-empty. Pick any $u \in L \cap B$ and $v \in R \cap B$. Since B is 3-edge-connected in G, there are three edge-disjoint paths P_1, P_2, P_3 in G connecting u and v. Note that in G there are at most 2 edges with one endpoint in $\widehat{C}$ and the other outside of $\widehat{C}$: these are the images of the at most two edges between A and C under α. At most two of the paths P_1, P_2, P_3 can contain any of these at most two edges, hence one of them, say P_1, must have all the vertices contained in $\widehat{C}$. But then P_1 contains an edge with one endpoint in L and second in R, a contradiction. □

Claim 5.1 provides us with an understanding of the replacement edges in torsos of the 3-edge-connected components of G. Precisely, let A be a 3-edge-connected component of G and let C be any 2-connected component of $G_{3\mathrm{CC}}$ that is a cycle and contains A. By Claim 5.1, all the vertices of $\bigcup_{D \in V(C) \setminus \{A\}} V(D)$ lie in the same connected component of $G - A$. Call this component $Z^A(C)$. Note that since C ranges over all 2-connected components of $G_{3\mathrm{CC}}$ that are cycles containing A, the components $Z^A(C)$ are pairwise different. Finally, if e^1, e^2 are the two edges of C that are incident to A, then the replacement edge of $Z^A(C)$ in $\mathsf{torso}(A)$ connects the endpoints of $\alpha(e^1)$ and $\alpha(e^2)$ that lie in A, or is non-existent if these endpoints coincide.

Using all these observations we can understand the traces of edges of G in the decomposition $\mathcal{T}$. For an edge e of G, by $\mathsf{trace}_{\mathcal{T}}(e)$ we mean the edge set of the trace of e in $\mathcal{T}$. Similarly, if e is an edge of $\mathsf{torso}(A)$, for some 3-edge-connected component A, then $\mathsf{trace}_{\mathcal{T}^A}(e)$ is the edge set of the trace of e in $\mathcal{T}^A$. The following claim explains how the traces of the edges of G behave in $\mathcal{T}$. The proof is a straightforward check using the observations presented above, hence we omit it. It can be followed on Figure 5.4.

Claim 5.2. *Let e be an edge of G with endpoints u and v, and let A and B be the 3-edge-connected components of G such that $u \in A$ and $v \in B$. Then:*

- *If $A = B$, then*
$$\mathsf{trace}_{\mathcal{T}}(e) = \mathsf{trace}_{\mathcal{T}^A}(e).$$
- *If $A \neq B$ and $\alpha^{-1}(e) \in S$, then*
$$\mathsf{trace}_{\mathcal{T}}(e) = \{\gamma(\alpha^{-1}(e))\}.$$
- *If $A \neq B$ and $\alpha^{-1}(e) \in R$, say $e = r_C$ for some 2-connected component C of $G_{3\mathrm{CC}}$ that is a cycle, then*
$$\mathsf{trace}_{\mathcal{T}}(e) = \bigcup_{D \in V(C)} \mathsf{trace}_{\mathcal{T}^D}\left(f_{Z^D(C)}\right) \cup \bigcup_{g \in E(C) \setminus \{e\}} \{\gamma(g)\}.$$

Here, if the replacement edge $f_{Z^D(C)}$ does not exist, we take $\emptyset$ for the corresponding trace.

For a 3-edge-connected component A we define a mapping $\eta^A : E(\mathsf{torso}(A)) \to E(G)$ as follows:

- If e is not a replacement edge, then set $\eta^A(e) = e$.

[3] Formally, we required tree-cut decompositions to be trees and not just forests, but we can always add arbitrary edges to make T connected without increasing any of the width measures.

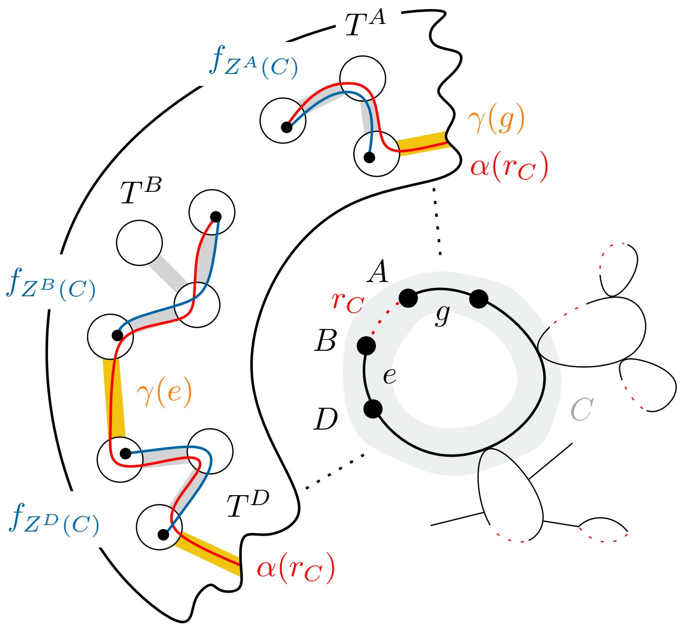

Figure 5.4: A close look at the trace of $\alpha(r_C)$.

- If e is a replacement edge, say $e = f_Z$ for some connected component Z of $G - A$, then observe that there is a unique 2-connected component C of G_{3CC} that is a cycle containing A and for which $Z = Z^A(C)$. Then set $\eta^A(e) = \alpha(r_C)$.

With this notation in place, Claim 5.2 immediately gives the following characterization of adhesions in $\mathcal{T}$.

Claim 5.3. *Let e be an edge of T. Then:*

- *If $e = \gamma(g)$ for some edge $g \in S$, say belonging to a 2-connected component C of G_{3CC} that is a cycle, then*
$$\mathrm{adh}_{\mathcal{T}}(e) = \{\alpha(g), \alpha(r_C)\}.$$
- *If $e \in E(T^A)$ for some 3-edge-connected component A, then*
$$\mathrm{adh}_{\mathcal{T}}(e) = \eta^A\left(\mathrm{adh}_{\mathcal{T}^A}(e)\right).$$

From the first point of Claim 5.3 it follows that all edges of T that originate from the spanning forest S are thin in $\mathcal{T}$, hence they do not contribute to the adhesions of the nodes of $\mathcal{T}$. Then, from the second point of Claim 5.3 we observe that $|\mathrm{adh}_{\mathcal{T}}(s)| \leqslant |\mathrm{adh}_{\mathcal{T}^A}(s)|$ for every node s of T that originates from T^A. Since $\mathbf{aw}(\mathcal{T}^A) \leqslant a$ for every 3-edge-connected component A, it follows that $\mathbf{aw}(\mathcal{T}) \leqslant a$. □

5.3.3 From the non-existence of a decomposition to a bramble

We are now ready to prove the last implication of Theorem E. The proof closely follows the line of argumentation for the treewidth case presented by Diestel in [Die16], which in turn is based on a proof by Mazoit [Maz13].

Proof (of Theorem E, (A3)⇒(A1)). Assume that G has no tree-cut decomposition of adhesion-width $< a$ and bag-width $< b$. We deduce by Theorem 5.11 that there exists a 3-edge-connected component A of G such that every tree-cut decomposition $\mathcal{T}$ of $\mathsf{torso}(A)$ satisfies at least one of the conditions: $\mathbf{aw}(\mathcal{T}) \geqslant a$ or $\mathbf{bw}(\mathcal{T}) \geqslant b$. We will construct a suitable bramble using the component A. Denote $G_A \coloneqq \mathsf{torso}(A)$ for brevity.

A tree-cut decomposition $\mathcal{T} = (T, \mathcal{X})$ of G_A shall be called *good* if for every node t, we have

$$\mathrm{adh}_{\mathcal{T}}(t) < a, \text{and}$$
$$\text{if } |X_t| \geqslant b, \text{ then } t \text{ is a leaf of } T.$$

In other words, a good tree decomposition has adhesion-width $< a$, and the only nodes whose bags are allowed to be of size $\geqslant b$ are the leaves. If $\mathcal{T} = (T, \mathcal{X})$ is a good tree-cut decomposition and a leaf t of T

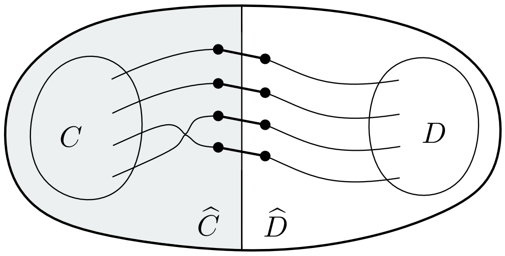

Figure 5.5: A separation $\{\widehat{C}, \widehat{D}\}$ of G_A such that $|\delta(\widehat{C})| = |\delta(\widehat{D})| = |\mathcal{P}|$. Edges in $\delta(\widehat{C})$ are depicted in bold.

satisfies $|X_t| \geqslant b$, then X_t shall be called a *petal* of $\mathcal{T}$. Note that the assumption that G_A has no tree-cut decomposition of adhesion-width $< a$ and bag-width $< b$ implies that every good tree-cut decomposition of G_A has a petal.

Our first goal is to construct a bramble $\mathcal{B}_A$ of adhesion-order $\geqslant a$ and bag-order $\geqslant b$ in G_A. For this, let $\mathcal{M} \subseteq 2^A$ be the family of all petals of all good tree-cut decompositions of G_A. Further, let $\mathcal{F} \subseteq \mathcal{M}$ be an inclusion-wise minimal subfamily of $\mathcal{M}$ satisfying the following two conditions:

(i) For each good tree-cut decomposition $\mathcal{T}$ of G_A, $\mathcal{F}$ contains at least one petal of $\mathcal{T}$.

(ii) $\mathcal{F}$ is *upward-closed*: if $C, D \in \mathcal{M}$ are such that $C \subseteq D$ and $C \in \mathcal{F}$, then also $D \in \mathcal{F}$.

We observe the following.

Claim 5.4. *Suppose C is an inclusion-wise minimal element of $\mathcal{F}$. Then there exists a good tree-cut decomposition $\mathcal{T}^C$ of G_A such that C is a petal of $\mathcal{T}^C$, and moreover C is the only petal of $\mathcal{T}^C$ that belongs to $\mathcal{F}$.*

Proof. Observe that since C is an inclusion-wise minimal element of $\mathcal{F}$, the family $\mathcal{F} \setminus \{C\}$ is upward-closed, that is, satisfies (ii). As $\mathcal{F}$ is inclusion-wise minimal subject to satisfying both (ii) and (i), it follows that $\mathcal{F} \setminus \{C\}$ does not satisfy (i), which implies the claim. □

The next claim is the key observation.

Claim 5.5. *The elements of $\mathcal{F}$ pairwise intersect.*

Proof. Suppose otherwise: there exist sets $C, D \in \mathcal{F}$ that are disjoint. By possibly replacing each of C and D by its subset, we may assume that C and D are inclusion-wise minimal elements of $\mathcal{F}$. By Claim 5.4, there exist good tree-cut decompositions $\mathcal{T}^C = (T^C, \mathcal{X}^C)$ and $\mathcal{T}^D = (T^D, \mathcal{X}^D)$ of G_A such that C is the only petal of $\mathcal{T}^C$ that belongs to $\mathcal{F}$ and D is the only petal of $\mathcal{T}^D$ that belongs to $\mathcal{F}$. Let t^C be the leaf of T^C whose bag is C, and define t^D analogously.

Let $\mathcal{P}$ be a maximum-size family of edge-disjoint paths in G_A connecting C with D. By Menger's theorem, there exists a separation $\{\widehat{C}, \widehat{D}\}$ of G_A such that $|\delta(\widehat{C})| = |\delta(\widehat{D})| = |\mathcal{P}|$. In particular, every path $P \in \mathcal{P}$ contains exactly one edge in $\delta(\widehat{C})$, all the vertices on P before this edge belong to $\widehat{C}$, and all the vertices on P after this edge belong to $\widehat{D}$. This is illustrated in Figure 5.5.

We now construct a decomposition $\mathcal{T} = (T, \mathcal{X})$ of G_A as follows. Construct T by taking the disjoint union of T^C and T^D, removing the nodes t^C (from T^C) and t^D (from T^D), and adding the edge $s^C s^D$, where s^C is the unique neighbor of t^C in T^C and s^D is the unique neighbor of t^D in T^D. (Note here that s^C and s^D exist, as otherwise either $C = A$ or $D = A$, implying that the other one is empty, but every element of $\mathcal{F}$ has size at least $b \geqslant 2$.) Next, define the bags $\mathcal{X} = \{X_t : t \in V(T)\}$ as follows: for every node t originating from T^C we set $X_t \coloneqq X_t^C \cap \widehat{D}$, and for every node t originating from T^D we set $X_t \coloneqq X_t^D \cap \widehat{C}$. Thus, $\{X_t : t \in V(T^C) \setminus \{t^C\}\}$ is a near-partition of $\widehat{D}$ and $\{X_t : t \in V(T^D) \setminus \{t^D\}\}$ is a near-partition of $\widehat{C}$, implying that $\mathcal{X}$ is a near-partition of A. So $\mathcal{T}$ is a tree-cut decomposition of G_A.

We now bound the adhesion-width of $\mathcal{T}$. Similarly as before, for an edge e, the edge set of the trace of e in $\mathcal{T}$ is denoted by $\mathsf{trace}_{\mathcal{T}}(e)$; similarly for decompositions $\mathcal{T}^C$ and $\mathcal{T}^D$. Observe that for every edge e of G, say with endpoints u and v, the trace of e in $\mathcal{T}$ can be characterized as follows.

- If $u, v \in \widehat{C}$, then $\mathsf{trace}_{\mathcal{T}}(e) = \mathsf{trace}_{\mathcal{T}^D}(e)$.
- If $u, v \in \widehat{D}$, then $\mathsf{trace}_{\mathcal{T}}(e) = \mathsf{trace}_{\mathcal{T}^C}(e)$.
- Suppose $u \in \widehat{C}$ and $v \in \widehat{D}$. Then there exists a path $P \in \mathcal{P}$ such that e lies on P. Let P^C be the prefix of P consisting of edges with at least one endpoint in $\widehat{C}$, and P^D be the suffix of P consisting of edges with at least one endpoint in $\widehat{D}$ (thus, $E(P^C) \cap E(P^D) = \{e\}$). Then

$$\mathsf{trace}_{\mathcal{T}}(e) \subseteq \left(\bigcup_{f \in E(P^C)} \mathsf{trace}_{\mathcal{T}^C}(f) \setminus \{t^C s^C\} \right) \cup \left(\bigcup_{f \in E(P^D)} \mathsf{trace}_{\mathcal{T}^D}(f) \setminus \{t^D s^D\} \right) \cup \{s^C s^D\}.$$

Let $\eta^C \colon E(G_A) \rightharpoonup E(G_A)$ be a partial function defined as follows:

- each edge e with both endpoints in $\widehat{C}$ is mapped to itself;
- each other edge e is not in the domain of η^C, unless it belongs to some path $P \in \mathcal{P}$, in which case it is mapped to the unique edge of P with one endpoint in $\widehat{C}$ and the other in $\widehat{D}$.

Define a partial function $\eta^D \colon E(G_A) \rightharpoonup E(G_A)$ symmetrically using $\widehat{D}$ instead of $\widehat{C}$. Then from the above characterization of traces it follows that:

- For each node $t \in V(T^C)$, we have
$$\mathrm{adh}_{\mathcal{T}}(t) \subseteq \eta^D\left(\mathrm{adh}_{\mathcal{T}^C}(t)\right).$$
- For each node $t \in V(T^D)$, we have
$$\mathrm{adh}_{\mathcal{T}}(t) \subseteq \eta^C\left(\mathrm{adh}_{\mathcal{T}^D}(t)\right).$$

Since both $\mathcal{T}^C$ and $\mathcal{T}^D$ have adhesion-width $< a$, it follows that $\mathbf{aw}(\mathcal{T}) < a$.

Finally, observe that every bag in $\mathcal{T}$ is a subset of a bag originating either from $\mathcal{T}^C$ or from $\mathcal{T}^D$. Since non-leaf nodes of $\mathcal{T}$ originate from non-leaf nodes of $\mathcal{T}^C$ and $\mathcal{T}^D$, it follows that only the leaves of $\mathcal{T}$ may have bags of size $\geqslant b$, hence $\mathcal{T}$ is good. Further, for every leaf t of $\mathcal{T}$, the bag at t in $\mathcal{T}$ is either a subset of a leaf bag in $\mathcal{T}^C$ other than C, or a subset of a leaf bag in $\mathcal{T}^D$ other than D. Since C and D are the only petals of $\mathcal{T}^C$ and $\mathcal{T}^D$, respectively, that belong to $\mathcal{F}$, and $\mathcal{F}$ is upward-closed, it follows that $\mathcal{T}$ has no petals that belong to $\mathcal{F}$. This is a contradiction with property (i) of $\mathcal{F}$. ☐

We can now define the bramble $\mathcal{B}_A$ as follows: for each connected subgraph H of G_A such that $V(H) \in \mathcal{F}$, include the slab $(H, V(H))$ in $\mathcal{B}_A$. Note that since G_A is 3-edge-connected by Lemma 5.10, these are indeed slabs in G_A. Further, Claim 5.5 shows that $\mathcal{B}_A$ is a bramble. We now verify the orders of $\mathcal{B}_A$.

Claim 5.6. *$\mathcal{B}_A$ has adhesion-order $\geqslant a$.*

Proof. Consider any set of edges $F \subseteq E(G_A)$ satisfying $|F| < a$. Construct a tree-cut decomposition $\mathcal{T}^F$ of G_A as follows. For each connected component D of $G_A - F$, construct a node t^D with $V(D)$ as its bag. Finally, construct a root node r with an empty bag and make it adjacent to all the nodes t^D. (Thus, the tree underlying $\mathcal{T}^F$ is a star with r being the center.) Since the adhesions of all the nodes are contained in F, it follows that $\mathbf{aw}(\mathcal{T}^F) < a$. Further, every node other than r is a leaf, and r's bag is empty, so we conclude that $\mathcal{T}^F$ is a good tree-cut decomposition of G_A. By property (i) of $\mathcal{F}$, there is a connected component D of $G_A - F$ such that $V(D) \in \mathcal{F}$. Observe that the slab $(D, V(D))$ has been included in $\mathcal{B}_A$ and $E(D) \cap F = \emptyset$, so in particular this slab is not disconnected by F. Since F was chosen arbitrarily, we conclude that $\mathcal{B}_A$ has no disconnecting set of size $< a$. ☐

Claim 5.7. *$\mathcal{B}_A$ has bag-order $\geqslant b$.*

Proof. It suffices to note that every element of $\mathcal{F}$ is a petal of some good tree-cut decomposition of G_A, and hence has size $\geqslant b$. ☐

Now that the bramble $\mathcal{B}_A$ in G_A is constructed, we can modify it to obtain a bramble $\mathcal{B}$ in G. Consider any slab $(H, K) \in \mathcal{B}_A$ and recalling that H is a subgraph of G_A, construct a subgraph H' of G from H as follows: for every replacement edge present in H, say edge f_Z for some connected component Z of $G - A$, replace f_Z by an arbitrary path connecting the endpoints of f_Z that has all internal vertices in Z. Note that H' remains connected and K, as a subset of A, is 3-edge-connected in G. Hence (H', K) is a slab in G. We define $\mathcal{B}$ as the set of all slabs (H', K) obtained from slabs $(H, K) \in \mathcal{B}_A$ as described above.

Since the cores of slabs did not change, $\mathcal{B}$ is a bramble in G and its bag-order is the same as that of $\mathcal{B}_A$, which is $\geqslant b$. To see that the adhesion-order of $\mathcal{B}$ is not smaller than that of $\mathcal{B}_A$, observe the following: if $F' \subseteq E(G)$ is a disconnecting set for $\mathcal{B}$, then replacing every edge of F' with an endpoint outside of A, say in a component Z of $G - A$, with the replacement edge f_Z, turns F' into a disconnecting set F for $\mathcal{B}$ such that $|F| \leqslant |F'|$. Therefore, we conclude that $\mathcal{B}$ is a bramble in G of adhesion-order $\geqslant a$ and bag-order $\geqslant b$. ☐

5.4 Cops, dogs, and robber game

In this section we use the equivalence provided by Theorem E to give a characterization of ab-tree-cut width expressed in terms of an analogue of the cops and robber game, which we call the *cops, dogs, and robber game.*

The game is played on a graph G by two players, one controlling cops and dogs, and the other controlling the robber. There are two parameters of the game:

- the number of cops a; and
- the number of dogs b.

The game starts with the robber player placing the robber at some vertex r_0. At all times, the cops occupy a set F consisting of at most a edges of G. This set is initially empty, that is, $F_0 = \emptyset$. Then the players proceed in rounds. Each round $i \in \{1, 2, 3, \ldots\}$, consists of the following steps:

- The cop player announces a set of edges F_i of size $\leqslant a$ to which the cops will move in this round.
- The robber player moves the robber from vertex r_{i-1} to any vertex r_i which is reachable from r_{i-1} by a path in G which does not pass through any edge occupied by a cop that does not move, that is, a path that does not intersect $F_{i-1} \cap F_i$. However, we also require that r_i and r_{i-1} are 3-edge-connected in G.
- The cops execute the announced move.

Note that so far the dogs do not get to play, but they are important in the winning condition. Namely, after every round the players verify to how many vertices the robber could potentially move, that is, how many vertices of G are 3-edge-connected with r_i and can be reached from r_i by a path that avoids the edges of F_i. If this number is $\leqslant b$, then the cop player can unleash the b dogs on those vertices and immediately catch the robber, thus winning the game. The robber player wins the game if she can avoid getting caught indefinitely.

Note that the restriction about 3-edge-connectedness of the moves essentially means that the robber is confined to the 3-edge-connected component of G to which r_0 belongs.

We say that a graph G is *searchable* by a cops and b dogs, if there exists a strategy for the cop player to win the game using a cops and b dogs. The main result of this section is the following equivalence.

Theorem F (restated). *Let G be a graph and a, b be positive integers. Then G is searchable by $< a$ cops and $< b$ dogs if and only if G has a tree-cut decomposition of adhesion-width $< a$ and bag-width $< b$.*

Proof. First, observe again that if $b = 1$, then both statements trivially do not hold for every positive integer a. Therefore, we shall assume that $b \geqslant 2$.

We show the left-to-right implication by proving its contrapositive. Assume that G has no tree-cut decomposition with adhesion-width $< a$ and bag-width $< b$. We describe a winning strategy for the robber player to win against $< a$ cops and $< b$ dogs.

By Theorem E, in G there exists a bramble $\mathcal{B}$ of adhesion-order $\geqslant a$ and bag-order $\geqslant b$. The robber player will maintain the following invariant: if at the end of round i the robber is placed at vertex r_i, then there is a slab $(H_i, K_i) \in \mathcal{B}$ such that $r_i \in K_i$ and (H_i, K_i) is not disconnected by F_i (the set of edges occupied by the cops at the end of round i). Since $F_0 = \emptyset$, to have this invariant satisfied at the start of the game, it suffices that the robber player chooses r_0 to be any vertex from the core of any slab $(H_0, K_0) \in \mathcal{B}$.

We now explain how the invariant is maintained in round i of the game. When the cop player announces the new set F_i to which the cops will move, the robber chooses any slab (H_i, K_i) that is not disconnected by F_i. Such a slab exists by the assumption that the adhesion-order of $\mathcal{B}$ is at least a and $b \geqslant 2$. Since $\mathcal{B}$ is a bramble, the cores K_{i-1} and K_i intersect. The robber player can therefore choose any vertex $r_i \in K_{i-1} \cap K_i$. As the invariant was satisfied in round $i-1$, the set $F_{i-1} \cap F_i$ (in fact, even F_{i-1}) does not disconnect r_{i-1} from r_i within H_{i-1}. Since $r_{i-1}, r_i \in K_{i-1}$ and K_{i-1} is 3-edge-connected in G, we conclude that it is allowed for the robber to move from r_{i-1} to r_i, and this move is duly executed by the robber player.

To see that in this way the robber player can evade being caught indefinitely, observe that provided the invariant is maintained, after round i the robber is allowed to move from r_i to any vertex of K_i. As the bag-order of $\mathcal{B}$ is $\geqslant b$, we have $|K_i| \geqslant b$, hence $< b$ dogs are never sufficient to catch the robber if she follows the described strategy.

We now proceed to the right-to-left implication. This amounts to describing a strategy for $< a$ cops and $< b$ dogs to search G, assuming that G has a tree-cut decomposition $\mathcal{T} = (T, \mathcal{X})$ satisfying $\mathbf{aw}(\mathcal{T}) < a$ and $\mathbf{bw}(\mathcal{T}) < b$.

Let us root the tree T in an arbitrary node, which naturally imposes an ancestor-descendant relation in T. After round 0, when we have $F_0 = \emptyset$, in rounds $1, 2, 3, \ldots$ the cop player will select nodes $t_1, t_2, t_3, \ldots$ so that each t_{i+1} is a child of t_i, and play

$$F_i \coloneqq \mathrm{adh}_{\mathcal{T}}(t_i).$$

While doing this, the cop player will maintain the following invariant: after round i, either she has already won, or the robber must be placed at a vertex that belongs to a bag of a (strict) descendant of t_i.

The strategy for maintaining the invariant is simple. In round 1 the cop player chooses t_1 to be the root of T, while in round $i \geqslant 2$ she chooses t_i to be the child of t_{i-1} such that r_{i-1} belongs to the bag of either t_{i-1} or any of its descendants. We now verify that the invariant is maintained after round i; we do this only for $i \geqslant 2$, as for $i = 1$ the check is almost the same. For brevity, let T_i be the subtree of T rooted at t_i.

First, observe that either the edge $t_i t_{i-1}$ is thin in $\mathcal{T}$, or

$$\mathrm{adh}(t_i t_{i-1}) \subseteq \mathrm{adh}(t_i) \cap \mathrm{adh}(t_{i-1}) = F_i \cap F_{i-1}.$$

In either case, the robber cannot move from r_{i-1} to any vertex w outside of $\bigcup_{s \in V(T_i)} X_s$, because either w and r_{i-1} are not 3-edge-connected in G, or every path connecting r_{i-1} with w intersects $F_i \cap F_{i-1}$. Hence, the robber player must choose $r_i \in \bigcup_{s \in V(T_i)} X_s$. However, if she chooses $r_i \in X_{t_i}$, then she immediately loses after this round: the set of vertices to which the robber can move once the cops are on F_i would be confined to a subset of X_{t_i}, which is of size $< b$. Hence, to avoid being captured the robber player must choose r_i to be a vertex contained in a bag of a strict descendant of t_i, and the invariant is maintained.

To see that following the strategy results in catching the robber, observe that eventually the cop player will chose t_i to be a leaf of T. Then she wins, as there is no vertex at which the robber can be placed after this round. $\square$

A graph is *k-searchable* if it is searchable by k cops and k dogs. Theorem F then implies the following.

Corollary 5.12. *A graph has ab-tree-cut width $\leqslant k$ if and only if it is k-searchable.*

Chapter 6

Digraphs without Onion Star Immersions

The motivation behind the results presented in this chapter was to define an appropriate directed counterpart of tree-cut width and tree-cut decompositions — in a way allowing to build a sound structural theory based on natural duality statements. This goal was not reached, since the directed setting turns out to deviate significantly from the undirected one.

However, we are able to identify an obstruction (cf. Figure 6.1) for which we can state and prove two duality results corresponding to the opening steps in Wollan's proof of the (undirected) Wall Immersion Theorem [Wol15]. These results are Theorems G and H.

Theorem G (restated). *There exists a function $f\colon \mathbb{N} \to \mathbb{N}$ such that the following holds. Suppose a digraph D contains a set X consisting of $2t+1$ vertices such that for all distinct $x, y \in X$, we have $\mu(x, y) > f(t)$. Then D contains the t-onion star as an immersion.*

Theorem H (restated). *There exists a function $g\colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that the following holds. Suppose D is a digraph and x, y are two distinct vertices in D such that $\mu(x, y) > g(t, k)$ and $\mu(y, x) > g(t, k)$, for some $t, k \in \mathbb{N}$. Then at least one of the following holds:*

- *D contains the t-onion star as an immersion;*
- *in D there is a family of $2k$ arc-disjoint paths consisting of k paths from x to y and k paths from y to x.*

Throughout the chapter we allow the existence of parallel arcs (multiple arcs with the same pair of endpoints), but we do not allow loops (arcs with both endpoints at the same vertex). Note that this in particular affects the way the word *digraph* should be understood.

6.1 Preliminaries

Paths. For the purpose of this chapter, a *path* in a digraph D is a sequence $(a_1, a_2, \ldots, a_k)$ of pairwise distinct arcs of D with the property that $\mathsf{head}(a_i) = \mathsf{tail}(a_{i+1})$ for every $i \in [k-1]$. Given a path $P = (a_1, a_2, \ldots, a_k)$ we put $\mathrm{first}(P) \coloneqq a_1$, $\mathrm{last}(P) \coloneqq a_k$, and $A(P) \coloneqq \{a_1, a_2, \ldots, a_k\}$. Path P is *simple* if $\mathsf{head}(a_i) = \mathsf{tail}(a_j)$ if and only if $j = i+1$, i.e., if the tail of the first arc and the heads of all arcs form $k+1$ distinct vertices. Let $<_P$ be the natural linear order of $A(P)$ defined by $a_i <_P a_j$ if and only if $i < j$. We define *trimmed paths*:

$$\begin{aligned}
P(\to a_i) &\coloneqq (a_1, a_2, \ldots, a_{i-1}),\\
P(\to a_i] &\coloneqq (a_1, a_2, \ldots, a_i),\\
P(a_i \to) &\coloneqq (a_{i+1}, a_{i+2}, \ldots, a_k),\\
P[a_i \to) &\coloneqq (a_i, a_{i+1}, \ldots, a_k),\\
P(a_i, a_j) &\coloneqq (a_{i+1}, a_{i+2}, \ldots, a_{j-1}).
\end{aligned}$$

If $s = \mathsf{tail}(a_1)$ and $t = \mathsf{head}(a_k)$, then P is called an *(s,t)-path*. If $\mathcal{P}$ is a family of paths and Q is a path, then we put $A(\mathcal{P}) \coloneqq \bigcup_{P \in \mathcal{P}} A(P)$ and $\mathcal{P}(Q) \coloneqq \{P \in \mathcal{P} \mid A(P) \cap A(Q) \neq \emptyset\}$. For two arc-disjoint paths

$P = (p_1, \ldots, p_n)$ and $Q = (q_1, \ldots, q_m)$ such that $\mathsf{head}(p_n) = \mathsf{tail}(q_1)$, we denote by PQ the *concatenation* of P and Q, i.e., the path $(p_1, \ldots, p_n, q_1, \ldots, q_m)$.

If $\mathcal{P}$ and $\mathcal{Q}$ are two families of paths in a digraph D, then by *intersection graph* of the pair $(\mathcal{P}, \mathcal{Q})$ we mean the undirected bipartite graph with bipartition classes $\mathcal{P}$, $\mathcal{Q}$, and an edge between $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$ if P and Q have at least one common arc.

Onions. Let $t \geqslant 1$. An *onion*[4] is the digraph $\overrightarrow{O}$ which consists of two vertices x and y, two arcs yx and one arc xy. The vertices x and y are called *roots* of the onion, with x being also called the *sink* of the onion, and y being the *source* of the onion. If ϕ is an immersion model (cf. Chapter 2) of $\overrightarrow{O}$ in D, then the vertices $\phi(x)$, $\phi(y)$ of D will be called roots (accordingly sink and source) of this immersion. A t*-onion-star* is the digraph $\overrightarrow{S_t}$ with the set of vertices $\{x, y_1, \ldots, y_t, z_1, \ldots, z_t\}$ and set of arcs $\bigcup_{i=1}^{t} A_i$, where each A_i consists of single arcs $y_i x$ and $x z_i$, and double arcs $x y_i$ and $z_i x$ (see Figure 6.1).

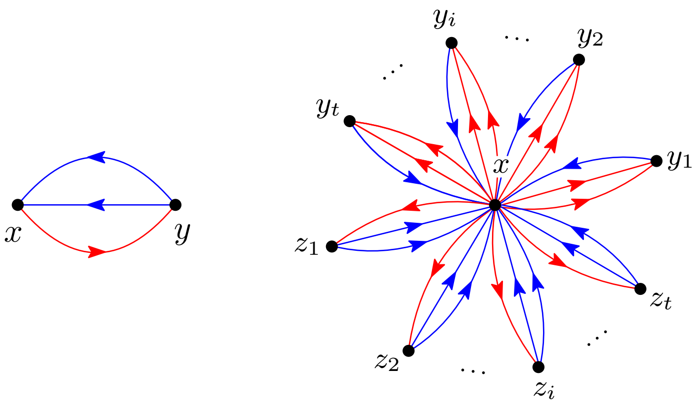

Figure 6.1: An onion (left) with sink x and source y and a t-onion-star (right). Blue and red arcs are incoming and outgoing arcs of x, respectively.

Observe that the graph $\overrightarrow{S_t}[\{y_1, \ldots, y_t, z_1, \ldots, z_t\}]$ is arcless and for every $i = 1, 2, \ldots, t$, the graphs $\overrightarrow{S_t}[\{x, y_i\}]$ and $\overrightarrow{S_t}[\{x, z_i\}]$ are onions. Moreover, x is the source of $\overrightarrow{S_t}[\{x, y_i\}]$ and the sink of $\overrightarrow{S_t}[\{x, z_i\}]$.

Ramsey-type results for bipartite graphs. Let n be a positive integer. By $b(n)$ we denote the minimum integer such that in every balanced graph on $2b(n)$ vertices, with bipartition classes X and Y, either there exists an induced subgraph isomorphic to $K_{n,n}$, or there exist two sets $A \subseteq X, B \subseteq Y$, each of size n, that are anti-complete to each other. The existence of $b(n)$ for each n follows from [Tho82].

Theorem 6.1 (Thomason [Tho82]). *For every $n \geqslant 1$ we have $b(n) \leqslant 2^n(n-1)+1$.*

In other words, if a bipartite graph G has sufficiently many vertices then we can always find in it two sets of size n, each contained in a different bipartition class, that are either complete or anti-complete to each other. In the same flavor, the following classical result of Kővári, Sós and Turán [KST54] gives a lower bound on the number of edges of a (not necessarily bipartite) graph G so that it contains a large complete balanced subgraph.

Theorem 6.2 (Kővári, Sós, Turán [KST54]). *There exists a function $c \colon \mathbb{N} \to (0, \infty)$ such that if an n-vertex graph G has at least $c(k)n^{2-1/k}$ edges for some $k \in \mathbb{N}$, then it contains $K_{k,k}$ as a subgraph.*

6.2 Onion Harvesting Lemma

This section is devoted to the proof of Lemma 6.3, the main conceptual piece of this chapter.

Let $\mathcal{P}$ be a family of paths in a digraph D and let $x \in V(D)$. We say that $\mathcal{P}$ *starts at* x if for every $P \in \mathcal{P}$ it holds that $\mathsf{tail}(\mathrm{first}(P)) = x$, and that it *ends at* x if for every $P \in \mathcal{P}$ it holds that $\mathsf{head}(\mathrm{last}(P)) = x$. Consider a pair $(\mathcal{P}, \mathcal{Q})$ where $\mathcal{P}$ and $\mathcal{Q}$ are families of pairwise arc-disjoint simple paths in D such that there exists $x \in V(D)$ with $\mathcal{P}$ starting at x and $\mathcal{Q}$ ending at x. We say that $(\mathcal{P}, \mathcal{Q})$ is a *well-crossing pair*

[4]We note that onions are directed variants of pumpkins, which were studied e.g. in [JPS+14].

rooted at x (or simply *well-crossing pair*) if for every $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$ we have $A(P) \cap A(Q) \neq \emptyset$, i.e., if the intersection graph of $(\mathcal{P}, \mathcal{Q})$ is complete.

We shall prove that, for every $t \geqslant 1$, if a digraph D contains a well-crossing pair $(\mathcal{P}, \mathcal{Q})$ with both families $\mathcal{P}$ and $\mathcal{Q}$ sufficiently large (in terms of t), then D admits a t-onion-star as an immersion. This is formalized by the following lemma.

Lemma 6.3 (Onion Harvesting Lemma). *Let $t \geqslant 1$ be an integer and D be a digraph. There exists a function $F \colon \mathbb{N} \to \mathbb{N}$ with the following property: if in D there exists a well-crossing pair $(\mathcal{P}, \mathcal{Q})$ with $|\mathcal{P}| \geqslant F(t)$ and $|\mathcal{Q}| \geqslant F(t)$, then in D there exists an immersion model of a t-onion-star.*

Let $(\mathcal{P}, \mathcal{Q})$ be a well-crossing pair. In the following, we call *crossing* of $(\mathcal{P}, \mathcal{Q})$ an arc e that belongs to $A(P) \cap A(Q)$ for some $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$, and specify P*-crossing* (resp. Q*-crossing*, (P, Q)*-crossing*) for a crossing of $(\mathcal{P}, \mathcal{Q})$ contained in P (resp. Q, $A(P) \cap A(Q)$). A crossing e is (P, Q)*-safe* (or simply *safe*) if there exist at least $|\mathcal{Q}|/3$ paths in $\mathcal{Q} \setminus \{Q\}$ whose $<_P$-minimal crossing with P precedes e in $<_P$. Otherwise, it is said to be (P, Q)*-dangerous* (or simply *dangerous*). We say that a path P is Q*-dangerous* if it contains a dangerous (P, Q)-crossing. If a path $Q \in \mathcal{Q}$ is trimmed to $\overline{Q}$, then we say that P is $\overline{Q}$*-dangerous* if there exists a dangerous (P, Q)-crossing belonging to $\overline{Q}$.

In order to prove Lemma 6.3, we will repeatedly "harvest" immersions of single onions rooted at the root of $(\mathcal{P}, \mathcal{Q})$, keeping appropriately large well-crossing pairs disjoint with previously found onions to enable gaining new ones.

Lemma 6.4. *Let $n \geqslant 1$ be an integer. There exists a function $f \colon \mathbb{N} \to \mathbb{N}$ with the following property: If in a digraph D there exists a well-crossing pair $(\mathcal{P}, \mathcal{Q})$ rooted at x with $|\mathcal{P}| = f(n)$ and $|\mathcal{Q}| = f(n)$, then in D there exist an immersion model ϕ of an onion with source x, and a well-crossing pair $(\mathcal{P}^*, \mathcal{Q}^*)$ rooted at x with $|\mathcal{P}^*| = n$, $|\mathcal{Q}^*| = n$ such that:*

- *$A(\mathcal{P}^*) \subseteq A(\mathcal{P})$, and $A(\mathcal{Q}^*) \subseteq A(\mathcal{Q})$;*
- *all paths in $\mathcal{P}^* \cup \mathcal{Q}^*$ are arc-disjoint with ϕ;*
- *no arc in $A(\mathcal{P}^*) \cap A(\mathcal{Q}^*)$ has the sink of ϕ for tail.*

Proof (of Lemma 6.4). Let c be the function from Theorem 6.2. Define functions g and f by:

$$g(n) = \max\left\{8, \left\lceil \frac{1}{2}(16c(2n))^{2n} \right\rceil\right\}, \qquad f(n) = \max\left\{6n, 3 \cdot \left\lceil \frac{1}{4}(36c(g(n)))^{g(n)} \right\rceil\right\}.$$

Note that since the values of f are divisible by 3, so are the sizes of $\mathcal{P}$ and $\mathcal{Q}$ for $(\mathcal{P}, \mathcal{Q})$ a well-crossing pair satisfying the assumptions of the lemma.

In the following, we consider one such pair $(\mathcal{P}, \mathcal{Q})$ rooted at a vertex x and fix $Q \in \mathcal{Q}$. Let e be the $<_Q$-smallest arc of Q such that $|\mathcal{P}(Q(e \rightarrow))| = |\mathcal{P}|/3$. Observe that e is a crossing satisfying $|\mathcal{P}(Q[e \rightarrow))| = |\mathcal{P}(Q(e \rightarrow))| + 1$. Let $\overline{Q} := Q[e \rightarrow)$. We distinguish two cases depending on the number of $\overline{Q}$-dangerous paths in $\mathcal{P}(\overline{Q})$. These cases are depicted in Figure 6.2 and Figure 6.3, respectively.

Case 1. Suppose that $\mathcal{P}(\overline{Q})$ contains at least two $\overline{Q}$-dangerous paths. Let d be the $<_Q$-greatest dangerous crossing in Q and let P_2 be the path such that d is a (P_2, Q)-crossing. Then d belongs to $\overline{Q}$. Let d' be the $<_{P_2}$-smallest (P_2, Q)-crossing belonging to $\overline{Q}(e' \rightarrow)$. Clearly since d is dangerous and belongs to $\overline{Q}$, d' must be dangerous (with possibly $d = d'$). Let e' be the $<_Q$-greatest dangerous crossing in $A(Q) \setminus A(P_2)$ and let $P_1 \in \mathcal{P}(\overline{Q}) \setminus \{P_2\}$ be the path such that e' is a (P_1, Q)-crossing (with possibly $e = e'$).

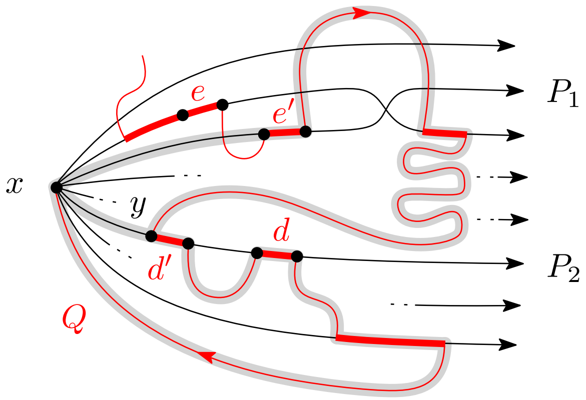

Figure 6.2: The situation of Case 1 when $e \neq e'$ and $d \neq d'$. Paths of $\mathcal{P}$ are depicted in black, Q is in red and Q-crossings are in thick red. An onion immersion model with source x and sink y is underlined in gray.

Let $y = \mathsf{tail}(d')$. We claim that the (x, y)-paths $P_1(\to e']Q(e', d')$, $P_2(\to d')$ and the (y, x)-path $Q[d' \to)$ form an immersion model ϕ of an onion with source x and sink y. Observe that $P_1(\to e']$ is arc-disjoint with $Q(e' \to)$ as e' is the $<_Q$-greatest dangerous (P_1, Q)-crossing, and having a safe (P_1, Q)-crossing in $A(P_1(\to e']) \cap A(Q(e' \to))$ would imply that e' is also safe, a contradiction. Similarly, the $<_{P_2}$-minimality of d' asserts that $P_2(\to d')$ is arc-disjoint with $Q[e' \to)$. Note that these arguments hold even if $e = e'$ and $d = d'$. Finally, $Q(e', d')$ and $Q[d' \to)$ are arc-disjoint as Q is simple. Hence, the path $P_1(\to e']Q(e', d')$ is a well-defined (although not necessarily simple) (x, y)-path, and paths $P_1(\to e']Q(e', d)$, $P_2(\to d')$, and $Q[d' \to)$ are pairwise arc-disjoint, yielding an immersion model ϕ of an onion with source x in D.

Let $\mathcal{Q}^\circ \subseteq \mathcal{Q}$ be the set of paths which are arc-disjoint both with $P_1(\to e']$ and $P_2(\to d']$ (so in particular $Q \notin \mathcal{Q}^\circ$). Since e' and d' are both dangerous, $|\mathcal{Q}^\circ| > |\mathcal{Q}| - |\mathcal{Q}|/3 - |\mathcal{Q}|/3 = |\mathcal{Q}|/3 \geqslant 2n$ by the definition of f. Let us pick an arbitrary collection $\mathcal{P}^* \subseteq \mathcal{P} \setminus \{P_1, P_2\}$ of paths not intersecting $Q[e' \to)$ with $|\mathcal{P}^*| = n$. Such a collection exists as $|\mathcal{P}(Q(e \to))| = |\mathcal{P}|/3$ yielding $\frac{2}{3}|\mathcal{P}| \geqslant 4n$ candidate paths for $\mathcal{P}^*$. Note that each path in such a collection has at most one arc whose tail is y (as the paths are simple), so there are at most $|\mathcal{P}^*| = n$ arcs in $A(\mathcal{Q}^\circ) \cap A(\mathcal{P}^*)$ whose tail is y. Let $\mathcal{Q}^* \subseteq \mathcal{Q}^\circ$ comprise of any n paths disjoint with these arcs. Clearly, since $\mathcal{P}^* \subseteq \mathcal{P}$ and $\mathcal{Q}^* \subseteq \mathcal{Q}$, $(\mathcal{P}^*, \mathcal{Q}^*)$ is a well-crossing pair. Hence ϕ, $\mathcal{P}^*$ and $\mathcal{Q}^*$ satisfy all the desired properties.

Case 2. Suppose that there is at most one $\overline{Q}$-dangerous path in $\mathcal{P}(\overline{Q})$. Let $\mathcal{S}$ be the set of those paths in $\mathcal{P}(\overline{Q})$ whose all crossings with $\overline{Q}$ are safe, and let

$$\overline{\mathcal{S}} \coloneqq \{P(\to \min{}_{<_P}\{A(P) \cap A(\overline{Q})\}) \mid P \in \mathcal{S}\}\,.$$

This is the set of paths from $\mathcal{S}$ trimmed so that they start at x, have no arcs in common with $\overline{Q}$ and are maximally long. Since $|\mathcal{P}(\overline{Q})| = |\mathcal{P}|/3 + 1$ and there is at most one dangerous path in $\mathcal{P}(\overline{Q})$, we derive $|\overline{\mathcal{S}}| = |\mathcal{S}| \geqslant |\mathcal{P}(\overline{Q})| - 1 = |\mathcal{P}|/3$. We shall first exhibit large well-crossing subfamilies of $\overline{\mathcal{S}}$ and $\mathcal{Q}$.

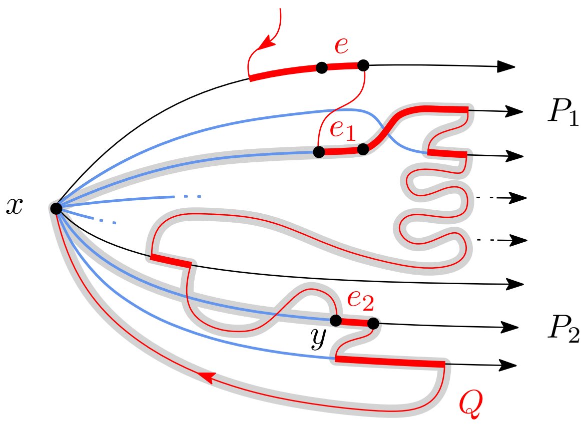

Figure 6.3: The situation of Case 2. Paths in $\overline{\mathcal{S}}$ are depicted in blue, Q is in red and Q-crossings are in thick red. Paths P_1 and P_2 are those of Claim 6.2. An onion immersion model with source x and sink y is underlined in gray.

Claim 6.1. *There exist subfamilies $\overline{\mathcal{S}}' \subseteq \overline{\mathcal{S}}$, $\mathcal{Q}' \subseteq \mathcal{Q} \setminus \{Q\}$ such that $|\overline{\mathcal{S}}'| = g(n)$, $|\mathcal{Q}'| = g(n)$ and $(\overline{\mathcal{S}}', \mathcal{Q}')$ is well-crossing.*

Proof (of claim). Let G be the intersection graph of $(\overline{\mathcal{S}}, \mathcal{Q})$. By definition, each trimmed path $\overline{P}$ in $\overline{\mathcal{S}}$ originates from a path P in $\mathcal{S}$ whose $\overline{(P,Q)}$-crossings e within $\overline{Q}$ are safe, i.e., there are at least $|\mathcal{Q}|/3$ distinct paths among $\mathcal{Q} \setminus \{Q\}$ crossing $\overline{P}$ before e with respect to $<_{\overline{P}}$. Since moreover $|\overline{\mathcal{S}}| \geqslant |\mathcal{P}|/3$, we get

$$|E(G)| = \sum_{\overline{P} \in \overline{\mathcal{S}}} \deg_G(\overline{P}) \geqslant \sum_{\overline{P} \in \overline{\mathcal{S}}} \frac{|\mathcal{Q}|}{3} \geqslant \frac{|\mathcal{P}|}{3} \cdot \frac{|\mathcal{Q}|}{3} = \left(\frac{f(n)}{3}\right)^2.$$

Note in addition that $\frac{4}{3}f(n) = \frac{|\mathcal{P}|}{3} + |\mathcal{Q}| \leqslant |V(G)| \leqslant 2f(n)$. By Theorem 6.2 there exists a constant $c_1 = c(g(n))$ such that if

$$|E(G)| \geqslant c_1 \cdot |V(G)|^{2-\frac{1}{g(n)}},$$

then G contains $K_{g(n),g(n)}$ as a subgraph, corresponding to the desired well-crossing pair $(\overline{\mathcal{S}}', \mathcal{Q}')$. For the above inequality to hold it is enough to guarantee that

$$\frac{|E(G)|}{|V(G)|^2} \geqslant \frac{c_1}{|V(G)|^{\frac{1}{g(n)}}}. \tag{6.1}$$

By $|V(G)| \leqslant 2f(n)$ and $|E(G)| \geqslant \frac{1}{9}f(n)^2$ we first get

$$\frac{|E(G)|}{|V(G)|^2} \geqslant \frac{\frac{1}{9}f(n)^2}{4f(n)^2} = \frac{1}{36}. \tag{6.2}$$

From the definition of f it follows that $f(n) \geqslant \frac{3}{4}(36c_1)^{g(n)}$, so equivalently

$$\frac{1}{36} \geqslant \frac{c_1}{\left(\frac{4}{3}f(n)\right)^{\frac{1}{g(n)}}}. \tag{6.3}$$

Moreover, $|V(G)| \geqslant \frac{4}{3}f(n)$ gives

$$\frac{c_1}{\left(\frac{4}{3}f(n)\right)^{\frac{1}{g(n)}}} \geqslant \frac{c_1}{|V(G)|^{\frac{1}{g(n)}}}. \tag{6.4}$$

Putting together inequalities eq. (6.2), eq. (6.3), eq. (6.4), we obtain eq. (6.1), which finishes the proof.□

Consider the pair $(\overline{\mathcal{S}}', \mathcal{Q}')$ given by Claim 6.1 and let $\mathcal{S}'$ be the set of paths in $\mathcal{S}$ from which $\overline{\mathcal{S}}'$ originates (i.e., the paths from $\overline{\mathcal{S}}'$ before trimming). For two distinct paths $P_1, P_2 \in \mathcal{S}'$, we denote

$$\overline{\mathcal{Q}}'(P_1, P_2) \coloneqq \{R(\max_{<_R}\{A(\{P_1,P_2\}) \cap A(R)\} \rightarrow) \mid R \in \mathcal{Q}'\}.$$

This is the set of paths from $\mathcal{Q}'$ trimmed so that they end at x, have no arcs in common with P_1 or P_2 and are longest possible.

Claim 6.2. *There exist distinct $P_1, P_2 \in \mathcal{S}'$ and subfamilies $\mathcal{P}^\circ \subseteq \overline{\mathcal{S}}'$, $\mathcal{Q}^\circ \subseteq \overline{\mathcal{Q}}'(P_1, P_2)$ such that $|\mathcal{P}^\circ| = 2n$, $|\mathcal{Q}^\circ| = 2n$ and $(\mathcal{P}^\circ, \mathcal{Q}^\circ)$ is well-crossing.*

Proof (of claim). Let $k \coloneqq g(n) = |\mathcal{Q}'| = |\overline{\mathcal{S}}'| = |\mathcal{S}'|$. Choose a pair $\{P_1, P_2\} \subseteq \mathcal{S}'$ uniformly at random (from the set of all $\binom{k}{2}$ two-element subsets). Let $\overline{\mathcal{Q}}' \coloneqq \overline{\mathcal{Q}}'(P_1, P_2)$ and let G be the intersection graph of $(\overline{\mathcal{S}}', \overline{\mathcal{Q}}')$. Clearly, $|\overline{\mathcal{Q}}'| = |\mathcal{Q}'| = k$. For $\overline{R} \in \overline{\mathcal{Q}}'$ let $N_{\overline{R}}$ be the random variable denoting the number of paths $P \in \overline{\mathcal{S}}'$ such that $A(\overline{R}) \cap A(P) \neq \emptyset$, i.e., the degree of $\overline{R}$ in G.

Fix $\overline{R} \in \overline{\mathcal{Q}}'$ and let $R \in \mathcal{Q}'$ be the path from which $\overline{R}$ originates (i.e., the untrimmed counterpart of $\overline{R}$). For every $P \in \overline{\mathcal{S}}'$ let $a(P) \coloneqq \max_{<_R}\{A(P) \cap A(R)\}$ be the $<_R$-greatest common arc of P and R. As $(\overline{\mathcal{S}}', \mathcal{Q}')$ is well-crossing, the k-element set $\{a(P) \mid P \in \overline{\mathcal{S}}'\} \subseteq A(R)$ is naturally linearly ordered by $<_R$. Label all paths from $\overline{\mathcal{S}}'$ with numbers from 1 to k accordingly with the order of appearance of arcs $a(P)$ on R, and let $\ell(P)$ be the label of P. Note that $N_{\overline{R}} \geqslant k - \max(\ell(P_1), \ell(P_2))$ and for every $i = 0, 1, \ldots, k-1$,

$$\mathbb{P}(\max(\ell(P_1), \ell(P_2)) = i + 1) = \frac{i}{\binom{k}{2}}.$$

By a straightforward computation we obtain

$$\mathbb{E}N_{\overline{R}} = \sum_{i=0}^{k-1} \mathbb{P}(N_{\overline{R}} = k-i-1)(k-i-1) = \frac{1}{\binom{k}{2}} \sum_{i=0}^{k-1} i(k-i-1) = \frac{k-2}{3}.$$

Note that the distribution of $N_{\overline{R}}$ depends only on labels of P_1 and P_2 and since these are chosen randomly, for different $\overline{R}, \overline{S} \in \overline{\mathcal{Q}}'$ the variables $N_{\overline{R}}$, $N_{\overline{S}}$ are identically distributed. By the above observations and since $k = g(n) \geqslant 8$, we have

$$\mathbb{E}|E(G)| \geqslant \mathbb{E} \sum_{\overline{S} \in \overline{\mathcal{Q}}'} N_{\overline{S}} \geqslant k \cdot \frac{k-2}{3} \geqslant \left(\frac{k}{2}\right)^2.$$

It follows that there exists a choice of P_1, P_2 such that

$$|E(G)| \geqslant \left(\frac{g(n)}{2}\right)^2. \tag{6.5}$$

Similarly as in the proof of Claim 6.1, we observe by Theorem 6.2 that there exists a constant $c_2 = c(2n)$ such that if

$$|E(G)| \geqslant c_2 \cdot |V(G)|^{2-\frac{1}{2n}},$$

then G contains $K_{2n,2n}$ as a subgraph, corresponding to the desired well-crossing pair $(\mathcal{P}^\circ, \mathcal{Q}^\circ)$. For the above inequality to hold it is enough to guarantee that

$$\frac{|E(G)|}{|V(G)|^2} \geqslant \frac{c_2}{|V(G)|^{\frac{1}{2n}}}.$$

This is obtained observing that

$$\frac{|E(G)|}{|V(G)|^2} = \frac{|E(G)|}{4(g(n))^2} \geqslant \frac{1}{16} \geqslant \frac{c_2}{(2g(n))^{\frac{1}{2n}}} = \frac{c_2}{|V(G)|^{\frac{1}{2n}}},$$

where the first inequality follows from eq. (6.5), and the second inequality follows from the definition of g ensuring that $g(n) \geqslant \frac{1}{2}(16c_2)^{2n}$. □

Consider P_1, P_2 and $(\mathcal{P}^\circ, \mathcal{Q}^\circ)$ as in the claim above and denote $e_i = \min_{<_{P_i}}\{A(\overline{Q}) \cap A(P_i)\}$ for $i = 1, 2$. Assume without loss of generality that $e_1 <_Q e_2$, i.e., that the $<_{P_1}$-smallest $(P_1, \overline{Q})$-crossing is further from x on $\overline{Q}$ than the $<_{P_2}$-smallest $(P_2, \overline{Q})$-crossing. Let $y = \mathsf{tail}(e_2)$. By the choice of e_1 and e_2, the (x, y)-paths $P_2(\to e_2)$, $P_1(\to e_1]Q(e_1, e_2)$ and the (y, x)-path $Q[e_2 \to)$ are arc-disjoint and hence form an immersion model ϕ of an onion with source x and sink y. Moreover, this model is disjoint with $A(\mathcal{P}^\circ \cup \mathcal{Q}^\circ)$ as $\mathcal{Q}^\circ \subseteq \overline{\mathcal{Q}}'$ with $\overline{\mathcal{Q}}'$ arc-disjoint with both P_1 and P_2.

The last step (to ensure that y will not be the tail of a crossing in the found well-crossing pair) is performed similarly as in the first case: we pick an arbitrary collection $\mathcal{P}^*$ of exactly n paths among the $2n$ paths in $\mathcal{P}^\circ$, and select a subfamily $\mathcal{Q}^*$ consisting of n paths which do not contain crossings of tail y among the $2n$ paths in $\mathcal{Q}^\circ$.

This concludes the proof of Lemma 6.4. □

Now, we observe that by iterating Lemma 6.4 t times, we can obtain an immersion model of a digraph consisting of t onions with a common source.

Corollary 6.5. *Let n, t be positive integers. Suppose that in a digraph D there exists a well-crossing pair $(\mathcal{P}, \mathcal{Q})$ rooted at x with $|\mathcal{P}| = |\mathcal{Q}| = f^{\langle t\rangle}(n)$, where f is the function from Lemma 6.4. Then in D there exists a family $\mathcal{H}$ of t arc-disjoint immersion models of onions with common source x, whose all sinks are pairwise different, and a well-crossing pair $(\mathcal{P}^*, \mathcal{Q}^*)$ rooted at x with $|\mathcal{P}^*| = n$, $|\mathcal{Q}^*| = n$ such that all paths in $\mathcal{P}^* \cup \mathcal{Q}^*$ are arc-disjoint with every element of $\mathcal{H}$.*

Proof. Put $\mathcal{P}_0 := \mathcal{P}$ and $\mathcal{Q}_0 := \mathcal{Q}$. For $i = 0, 1, \ldots, t-1$ let us apply Lemma 6.4 to the pair $(\mathcal{P}_i, \mathcal{Q}_i)$ to get an immersion model of an onion ϕ_{i+1} with sink y_{i+1}, and a well-crossing pair $(\mathcal{P}_{i+1}, \mathcal{Q}_{i+1})$.

Now it is enough to take $\mathcal{P}^* = \mathcal{P}_t$, $\mathcal{Q}^* = \mathcal{Q}_t$ and $\mathcal{H} = \{\phi_i\}_{i=1}^t$. Indeed, to see that the sinks $y_1, \ldots, y_t$ are distinct, recall that for every $i \in \{0, \ldots, t-1\}$ no arc in $A(\mathcal{P}_{i+1}) \cap A(\mathcal{Q}_{i+1})$ has y_{i+1} as a tail. Since $A(\mathcal{P}_i) \subseteq A(\mathcal{P}_j)$, and $A(\mathcal{Q}_i) \subseteq A(\mathcal{Q}_j)$ for every $0 \leqslant j \leqslant i$, vertices $y_1, \ldots, y_i$ do not appear as tails of arcs in $A(\mathcal{P}_i) \cap A(\mathcal{Q}_i)$, hence y_{i+1} must be distinct from $y_1, \ldots, y_i$. The remaining properties of the immersion model are straightforward to verify. □

An analogous lemma can be designed to enable finding onions whose sink (not source) is x. The proof basically follows from reversing all arcs in the arguments delivered previously.

Corollary 6.6. *Let n, t be positive integers. Suppose that in a digraph D there exists a well-crossing pair $(\mathcal{P}, \mathcal{Q})$ rooted at x with $|\mathcal{P}| = |\mathcal{Q}| = f^{\langle t\rangle}(n)$, where f is the function from Lemma 6.4. Then in D there exists a family $\mathcal{H}$ of t arc-disjoint immersion models of onions with common sink x, whose all sources are pairwise different, and a well-crossing pair $(\mathcal{P}^*, \mathcal{Q}^*)$ rooted at x with $|\mathcal{P}^*| = n$, $|\mathcal{Q}^*| = n$ such that all paths in $\mathcal{P}^* \cup \mathcal{Q}^*$ are arc-disjoint with every element of $\mathcal{H}$.*

Proof. It is enough to apply Corollary 6.5 to the digraph $\overleftarrow{D}$, in which each arc of D is replaced with a reversed arc, and to the (swapped) well-crossing pair $(\overleftarrow{\mathcal{Q}}, \overleftarrow{\mathcal{P}})$ consisting of families of reversed paths from $\mathcal{Q}$ and $\mathcal{P}$, respectively. $\square$

Finally, we use Corollary 6.5 and Corollary 6.6 to prove Lemma 6.3.

Proof (of Lemma 6.3). Take $F(t) = f^{\langle 4t\rangle}(1)$, where f is the function satisfying Lemma 6.4. Let x be the root of the pair $(\mathcal{P}, \mathcal{Q})$. Applying Corollary 6.5 to D, $(\mathcal{P}, \mathcal{Q})$, $n = f^{\langle 2t\rangle}(1)$ and $2t$ (in place of t) we find a family $\mathcal{H}^+ = \{\phi_i^+\}_{i=1}^{2t}$ of $2t$ arc-disjoint immersion models of onions with source x and mutually different sinks $\{y_i\}_{i=1}^{2t}$, along with a well-crossing pair $(\mathcal{P}^*, \mathcal{Q}^*)$ in which both families have size at least $f^{\langle 2t\rangle}(1)$ and are arc-disjoint with elements of $\mathcal{H}^+$. Now applying Corollary 6.6 to D, $(\mathcal{P}^*, \mathcal{Q}^*)$, $n = 1$ and $2t$ (in place of t), we find a family $\mathcal{H}^- = \{\phi_j^-\}_{j=1}^{2t}$ of $2t$ arc-disjoint immersion models of onions with sink x and mutually different sources $\{z_j\}_{j=1}^{2t}$. Moreover, $A(\mathcal{H}^+) \cap A(\mathcal{H}^-) = \emptyset$. It now remains to note that it is possible to find two sets $I, J \subseteq [2t]$ such that $|I| = |J| = t$ and $\{y_i : i \in I\} \cap \{z_j : j \in J\} = \emptyset$. The union of immersions ϕ_i^+ where $i \in I$ and ϕ_j^- where $j \in J$ forms an immersion of a t-onion-star. $\square$

6.3 Proofs of Theorems G and H

We are ready to prove our main theorems. We will first prove Theorem H and then Theorem G.

Proof (of Theorem H). We define

$$g(t, k) = 2^N (N - 1) + 1$$

where $N = \max\{k, F(t)\}$ and F is the function given by Lemma 6.3. Let $\mathcal{P}$ and $\mathcal{Q}$ be families of pairwise arc-disjoint simple paths such that elements of $\mathcal{P}$ start at x and end at y, elements of $\mathcal{Q}$ start at y and end at x, and $\min\{|\mathcal{P}|, |\mathcal{Q}|\} \geqslant g(t, k)$.

By Theorem 6.1 applied to the intersection graph of $(\mathcal{P}, \mathcal{Q})$, there exist families $\mathcal{P}' \subseteq \mathcal{P}$ and $\mathcal{Q}' \subseteq \mathcal{Q}$ of paths such that $|\mathcal{P}'|, |\mathcal{Q}'| \geqslant N$ and either $(\mathcal{P}', \mathcal{Q}')$ is a well-crossing pair, or $A(\mathcal{P}') \cap A(\mathcal{Q}') = \emptyset$. In the first case, we conclude by Lemma 6.3 the existence of an immersion model of a t-onion-star in D. As of the second outcome, it is precisely as desired. $\square$

Proof (of Theorem G). Define function f as

$$f(t) = 2t \cdot g_t^{\langle 4t^2\rangle}(2),$$

where $g_t(k) \coloneqq g(t, k)$ is the function from Theorem H.

We arbitrarily enumerate the vertices of X as $y, x_1, \ldots, x_{2t}$. Consider an auxiliary digraph D' obtained from D by adding a single vertex v, $\frac{f(t)}{2t}$ arcs vx_i and $\frac{f(t)}{2t}$ arcs x_iv for every $i \in [2t]$. First, observe that in D' there exists a family of $f(t)$ arc-disjoint (v, y)-paths. Indeed, let $\mathcal{P}'$ be the maximum family of arc-disjoint (v, y)-paths, and let k be the size of $\mathcal{P}'$. By Menger's Theorem, there exists a (v, y)-cut (A, B) of size k. Since $|\delta^+(v)| = f(t)$, we have $k \leqslant f(t)$. Observe that if $k < f(t)$, then $A \setminus \{v\} \neq \emptyset$, which implies that there exists $i \in [2t]$ such that $x_i \in A$. However, in such a case, $(A \setminus \{v\}, B)$ is an (x_i, y)-cut in D of size $k < f(t)$ contradicting the assumption of the theorem. Hence, $|\mathcal{P}'| = f(t)$. By a symmetric argument we obtain the existence of a family $\mathcal{Q}'$ of $f(t)$ arc-disjoint (y, v)-paths in D'. Clearly, we can assume that $\mathcal{P}'$ and $\mathcal{Q}'$ consist of simple paths.

Let $\mathcal{P}$ and $\mathcal{Q}$ respectively denote the families of paths $\mathcal{P}'$ and $\mathcal{Q}'$ once every path has been restricted to the arcs of D. Note that each $P \in \mathcal{P}$ (resp. $Q \in \mathcal{Q}$) is an (x, y)-path (resp. (y, x)-path) for some $x \in X \setminus \{y\}$. Moreover, for every $i \in [2t]$, the number of (x_i, y)-paths in $\mathcal{P}$ (resp. (y, x_i)-paths in $\mathcal{Q}$) is precisely $\frac{f(t)}{2t}$. Hence, if we denote by $\mathcal{P}_i$ the subset of $\mathcal{P}$ that consists of (x_i, y)-paths, for every $i \in [2t]$, we have that

$|\mathcal{P}_i| = \frac{f(t)}{2t} = g_t^{\langle 4t^2 \rangle}(2)$. Similarly, by $\mathcal{Q}_i$ we denote the subset of $\mathcal{Q}$ that consists of (y, x_i)-paths, and for every $i \in [2t]$ we have that $|\mathcal{Q}_i| = g_t^{\langle 4t^2 \rangle}(2)$.

Let $p \in [4t^2]$. We fix an arbitrary order on the elements of $[2t] \times [2t]$, and for the p-th element (i, j) we proceed as follows. We apply Theorem H to $\mathcal{P}_i$ and $\mathcal{Q}_j$ for $n = g_t^{\langle 4t^2 - p + 1 \rangle}(2)$. Either we conclude that there exists an immersion model of a t-onion-star in D, or replace $\mathcal{P}_i$ and $\mathcal{Q}_j$ by their subsets $\mathcal{P}_i'$ and $\mathcal{Q}_i'$ of size $g_t^{\langle 4t^2 - p \rangle}(2)$. Clearly the theorem holds if the first case occurs. In the other case, for every $i' \neq i, j' \neq j$, we reduce sets $\mathcal{P}_{i'}$ and $\mathcal{Q}_{j'}$ arbitrarily so that each of them contain exactly $g_t^{\langle 4t^2 - p \rangle}(2)$ paths, and proceed to the next iteration.

After $4t^2$ iterations, if no onion-star immersion model is found, for every i the families $\mathcal{P}_i$ and $\mathcal{Q}_i$ have precisely $g^{\langle 0 \rangle}(2) = 2$ elements each. Let $\mathcal{S}_i = \mathcal{P}_i \cup \mathcal{Q}_i$ and $\mathcal{S} = \bigcup_{i \in [2t]} \mathcal{S}_i$. We observe that paths from $\mathcal{S}$ are now pairwise arc-disjoint, and each $\mathcal{S}_i$ contains two paths $x_i \to y$ and two paths $y \to x_i$ for every $i \in [2t]$. Hence, there exists an t-onion-star immersion model in D, with center being y. This concludes the proof. $\square$

Part III

Bipartite Permutation Graphs

Chapter 7

Vertex Deletion into Bipartite Permutation Graphs

This chapter is devoted to proving Theorems I and J.

Theorem I (restated). *There is an $\mathcal{O}(9^k \cdot |G|^9)$-time algorithm for instances (G, k) of the vertex deletion into bipartite permutation graphs problem.*

Theorem J (restated). *There exists a polynomial-time 9-approximation algorithm for vertex deletion into bipartite permutation graphs problem.*

7.1 Posets and comparability graphs

A *partially ordered set* (shortly *partial order* or *poset*) is a pair $P = (X, \leqslant_P)$ that consists of a set X and a reflexive, transitive, and antisymmetric relation $\leqslant_P$ on X. For a poset $(X, \leqslant_P)$, let the *strict partial order* $<_P$ be a binary relation defined on X such that $x <_P y$ if and only if $x \leqslant_P y$ and $x \neq y$. Equivalently, $(X, <_P)$ is a strict partial order if $<_P$ is irreflexive and transitive. Two elements $x, y \in X$ are *comparable* in P if $x \leqslant_P y$ or $y \leqslant_P x$; otherwise, x, y are *incomparable* in P. A *linear order* $L = (X, \leqslant_L)$ is a partial order in which every two elements $x, y \in X$ are comparable. The *strict linear order* $(X, <_L)$ associated with L is a binary relation defined in a way that $x <_L y$ if and only if $x \leqslant_L y$ and $x \neq y$.

Let $P = (X, \leqslant_P)$ be a poset. A linear order $L = (X, \leqslant_L)$ is called a *linear extension* of P if $\leqslant_P \subseteq \leqslant_L$. Given a family of posets $\mathcal{P} = \{P_i = (X, \leqslant_{P_i}) : i \in I\}$, we say that P is the *intersection of* $\mathcal{P}$ if for every $x, y \in X$ we have $x \leqslant_P y$ if and only if $x \leqslant_{P_i} y$ for every $i \in I$. The *dimension* of a poset P is the minimal number of linear extensions of P that intersect to P. In particular, we say that P is *two-dimensional* if it is the intersection of two linear extensions of P.

For a graph $G = (V, E)$, a pair $(V, <)$ is a *transitive orientation* of G if and only if $<$ is a transitive and irreflexive relation on V that satisfies either $u < v$ or $v < u$ iff $uv \in E$ for every $u, v \in V$.

A *comparability graph* (*incomparability graph*) of a poset $P = (X, \leqslant_P)$ has X as the set of its vertices and the set including every two vertices comparable (incomparable, respectively) in P as the set of its edges. Note the following: if $(X, \leqslant_P)$ is a poset, then $(X, <_P)$ is a transitive orientation of the comparability graph of P. A graph $G = (V, E)$ is a *comparability graph* (*co-comparability graph*) if G is a comparability (incomparability, respectively) graph of some poset defined on V. So, G is a comparability graph if and only if G admits a transitive orientation. A graph G is a *permutation graph* if and only if G and the complement of G are comparability graphs [PLE71] (or equivalently, G and the complement of G admit transitive orientations). Baker, Fishburn, and Roberts [BFR72] proved that G is a permutation graph if and only if G is the incomparability graph of a two-dimensional poset.

7.2 The structure of (almost) bipartite permutation graphs

The characterization of bipartite permutation graphs presented below was proposed by Spinrad, Brandstädt, and Stewart [SBS87].

Suppose $G = (U, W, E)$ is a connected bipartite graph. A linear order $(W, <_W)$ satisfies *adjacency property* if for each vertex $u \in U$, the set $N(u)$ consists of vertices that are consecutive in $(W, <_W)$. A linear order $(W, <_W)$ satisfies *enclosure property* if for every pair of vertices $u, u' \in U$ such that $N(u)$ is a subset of $N(u')$, vertices in $N(u') - N(u)$ occur consecutively in $(W, <_W)$. A *strong ordering* of the vertices of $U \cup W$ consists of linear orders $(U, <_U)$ and $(W, <_W)$ such that for every $(u, w'), (u', w) \in E$, where $u, u' \in U$ and $w, w' \in W$, it holds that $u <_U u'$ and $w <_W w'$ imply $(u, w) \in E$ and $(u', w') \in E$. Note that (as G is connected) whenever $(U, <_U)$ and $(W, <_W)$ form a strong ordering of $U \cup W$, $(U, <_U)$ and $(W, <_W)$ satisfy the adjacency and enclosure properties.

Theorem 7.1 (Spinrad, Brandstädt, Stewart [SBS87]). *The following three statements are equivalent for a connected bipartite graph $G = (U, W, E)$:*

(a) (U, W, E) is a bipartite permutation graph.

(b) There exists a strong ordering of $U \cup W$.

(c) There exists a linear order $(W, <_W)$ of W satisfying adjacency and enclosure properties.

An example of a bipartite permutation graph $G = (U, W, E)$ with linear order $w_1 <_W w_2 <_W \ldots <_W w_8 <_W w_9$ of the vertices of W which satisfies the adjacency and the enclosure properties is shown in Figure 1.4 (reproduced below for convenience).

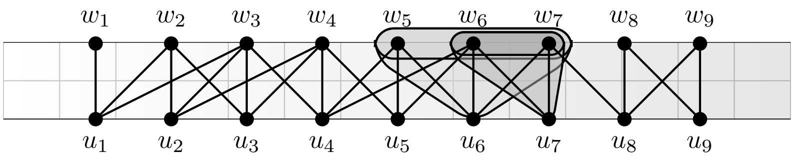

Figure 1.4 (reproduced): Embedding of a bipartite permutation graph (U, W, E) into a strip satysfying the adjacency and the enclosure properties.

Another characterization of bipartite permutation graphs can be obtained by listing all minimal forbidden induced subgraphs for this class of graphs. Such a list can be compiled by taking all odd cycles of length $\geqslant 3$ (forbidden structures for bipartite graphs) and all bipartite graphs from the list of forbidden structures for permutation graphs obtained by Gallai [Gal67]. The whole list is shown in Figure 7.1.

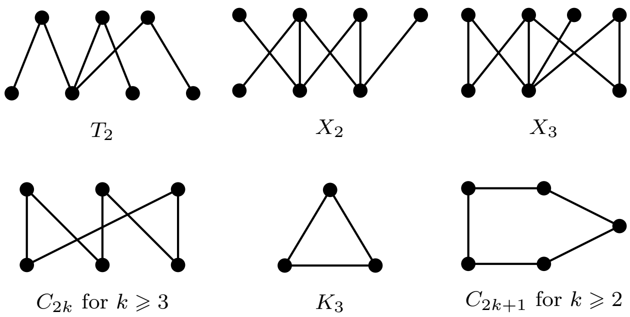

Figure 7.1: Forbidden structures for bipartite permutation graphs.

7.2.1 Almost bipartite permutation graphs

The goal of this section is to characterize graphs which do not contain small forbidden subgraphs for the class of bipartite permutation graphs. Following the terminology of van 't Hof and Villanger [vtHV13] we call such graphs almost bipartite permutation graphs.

Definition 7.1. *A graph $G = (V, E)$ is an* almost bipartite permutation graph *if G does not contain T_2, X_2, X_3, K_3, or C_k for $k \in [5, 9]$ as induced subgraphs.*

Suppose $G = (V, E)$ is a connected almost bipartite permutation graph. By *hole* we mean an induced cycle on at least five[5] vertices. We say that a hole is *even* (or *odd*) if it contains even (resp. odd) number of vertices. We say that $X \subseteq V$ is a *dominating set* if $X \cup N(X) = V(G)$.

[5] We deviate here from the generally accepted convention defining holes as chordless cycles of length at least 4. I.e., for the purpose of this chapter, C_4 is *not* a hole.

Proposition 7.2. *The vertex set of every hole in G is a dominating set.*

Proof. Let $C = \{c_0, c_1, \ldots, c_{m-1}\}$ be a hole in G. Hence, $m \geqslant 10$. Suppose, for a contradiction, that there exists a vertex in the set $V \setminus (C \cup N(C))$. As G is connected, there must exist $v \in V$ at distance two from C. Let $w \in N(v) \cap N(C)$ and let c_j be a neighbor of w in C. We now look at the neighborhood of w. As G contains no triangle, wc_{j-1} and wc_{j+1} are non-edges. Moreover, as G contains no copy of T_2, vertex w is adjacent to at least one of c_{j-2} and c_{j+2}, say c_{j-2}. Thus, w is nonadjacent to c_{j-3}. Therefore, the set $\{c_{j-3}, c_{j-2}, c_{j-1}, c_j, c_{j+1}, w, v\}$ induces a copy of X_2 in G, which leads to a contradiction. ☐

Let C be a shortest hole in G, m be the length of C, and $c_0, c_1, \ldots, c_{m-1}$ be the consecutive vertices of C, $m \geqslant 10$. In the remaining part of the paper we use the following notation with respect to C. For any integral number i by c_i we denote the unique vertex $c_{i \bmod m}$ from the cycle C. For any two different vertices c_i, c_j in C, by *the set of all vertices between c_i and c_j from C* we mean the set $\{c_i, c_{i+1}, \ldots, c_{i+k}\}$, where k is the smallest natural number such that $c_{i+k} = c_j$. Note that this notion is not symmetric, i.e., the set of all vertices between c_j and c_i from C contains c_i, c_j and all the vertices from C that are not between c_i and c_j.

Proposition 7.3. *For every vertex $v \in V$ either:*

(1) $N(v) \cap C = \{c_i\}$ for some $i \in [\![m-1]\!]$, or

(2) $N(v) \cap C = \{c_i, c_{i+2}\}$ for some $i \in [\![m-1]\!]$.

Proof. Since C is an induced cycle, (2) clearly holds for the vertices from C, so let v be a vertex in $V \setminus C$. As C is a dominating set, by Proposition 7.2, vertex v has at least one neighbor in C. If v has exactly one neighbor in C, then (1) holds and we are done. So assume that it has more than one neighbor. We now distinguish two cases. First, suppose that there exist two vertices $c_j, c_\ell \in N(v) \cap C$ at distance at least three in C such that v has no neighbor in the set of vertices between c_j and c_l, except c_j and c_l. Then, $\{c_j, c_{j+1}, \ldots, c_\ell, v\}$ induces a cycle C' on at least five vertices in G. As c_j and c_ℓ are at distance at least three in C, C' is shorter than C. In particular, C' contradicts either G containing no copy of C_ℓ, for $\ell \in \{5, \ldots, 9\}$, or C being a shortest hole in G. Therefore, this case never occurs.

Hence, v has either (i) exactly two neighbors in C and those are at distance two as there is no triangle in G, so (2) holds, or (ii) C has an even number of vertices and v is adjacent to every second vertex of C. It remains to show that the latter never occurs. Indeed, if it does, then without loss of generality $c_0 \in N(v)$. But observe that since C has at least ten vertices, the set $\{c_0, c_1, c_2, c_3, c_4, c_6, v\}$ induces a copy of X_3. This concludes the proof. ☐

Given Proposition 7.3, for every $i \in [\![m-1]\!]$ we can set $A_i = \big\{v \in V : N(v) \cap C = \{c_{i-1}, c_{i+1}\}\big\}$ and $B_i = \big\{v \in V : N(v) \cap C = \{c_i\}\big\}$. Note that sets $A_0, B_0, \ldots, A_{m-1}, B_{m-1}$ form a partition of V. Moreover, for every $i \in [\![m-1]\!]$ we have $c_i \in A_i$. Following our notation, for any integer i by A_i and B_i we denote the sets $A_{i \bmod m}$ and $B_{i \bmod m}$, respectively. Furthermore, for every $i \leqslant j$ we set:

$$A_G[i,j] = \begin{cases} A_i \cup B_{i+1} \cup A_{i+2} \cup B_{i+3} \cup \ldots \cup A_{j-1} \cup B_j & \text{if } j-i \text{ is odd,} \\ A_i \cup B_{i+1} \cup A_{i+2} \cup B_{i+3} \cup \ldots \cup B_{j-1} \cup A_j & \text{if } j-i \text{ is even,} \end{cases}$$

$$B_G[i,j] = \begin{cases} B_i \cup A_{i+1} \cup B_{i+2} \cup A_{i+3} \cup \ldots \cup B_{j-1} \cup A_j & \text{if } j-i \text{ is odd,} \\ B_i \cup A_{i+1} \cup B_{i+2} \cup A_{i+3} \cup \ldots \cup A_{j-1} \cup B_j & \text{if } j-i \text{ is even,} \end{cases}$$

$$\text{and} \quad V_G[i,j] = A_G[i,j] \cup B_G[i,j].$$

We write just $A[i,j]$, $B[i,j]$, and $V[i,j]$, respectively, instead of $A_G[i,j]$, $B_G[i,j]$, and $V_G[i,j]$, when there is no confusion.

We now characterize the neighborhoods of the vertices in sets A_i and B_i, see also Figure 7.2.

Proposition 7.4. *Let $i \in [\![m-1]\!]$. Then:*

(1) A_i and B_i are independent sets.

(2) For every $u \in A_i$ and every $w \in B_i$ we have $uw \in E$.

(3) For every $u \in A_i$ we have $B_i \subseteq N(u) \subseteq B[i-2, i+2]$.

(4) For every $w \in B_i$ we have $A_i \subseteq N(w) \subseteq A[i-2, i+2]$.

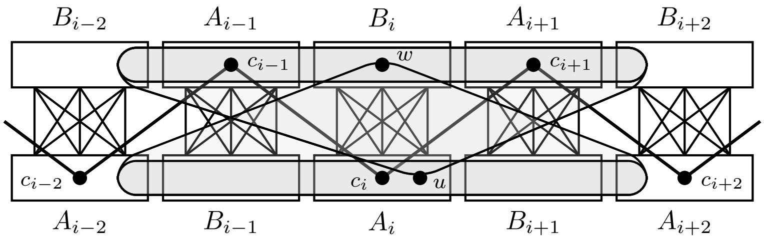

Figure 7.2: A possible neighborhood of u in A_i and w in B_i.

Proof. Statement (1) follows trivially from the fact that G contains no triangle. To show statement (2), assume for contradiction that $uw \notin E$ for some $u \in A_i$ and some $w \in B_i$. Since $u \in A_i$, we have $N(u) \cap C = \{c_{i-1}, c_{i+1}\}$, and since $w \in B_i$ we have $N(w) \cap C = \{c_i\}$. Hence, the set $\{c_{i-2}, u, c_i, c_{i+2}, c_{i-1}, w, c_{i+1}\}$ induces an X_2 in G, which cannot be the case.

To prove statements (3) and (4), consider a graph G induced by the set $U \cup W$, where

$$U = A[i-2, i+2] \quad \text{and} \quad W = B[i-2, i+2].$$

Since any edge with two endpoints in U (or two endpoints in W) could be extended by some vertices from $\{c_{i-2}, c_{i-1}, c_i, c_{i+1}, c_{i+2}\}$ to an odd cycle of length $\leqslant 7$ in G, the graph $G[U \cup W]$ is bipartite with bipartition classes U and W.

To see that (3) holds, first note that $B_i \subseteq N(u)$ by statement (2). Therefore, suppose that u has a neighbor v in the set $V \setminus (U \cup W)$.

Consider the case when $v \in A_j$ for some $j \notin [i-2, i+2]$. Since $N(u) \cap C = \{c_{i-1}, c_{i+1}\}$ and $N(v) \cap C = \{c_{j-1}, c_{j+1}\}$, u, v and the vertices between c_{j+1} and c_{i-1} in C as well as u, v and the vertices between c_{i+1} and c_{j-1} in C induce cycles in G of length $\leqslant m-2$. Since $|C| \geqslant 10$, at least one from these cycles has length $\geqslant 6$, and as such it cannot occur in G.

So suppose $v \in B_j$ for some $j \notin [i-2, i+2]$. Since $N(v) \cap C = \{c_j\}$, u, v and the vertices between c_{i+1} and c_j in C as well as u, v and the vertices between c_j and c_{i-1} in C induce holes in G of length $\leqslant m-2$, and as such they cannot occur in G. So, $N(u) \subseteq W$, which completes the proof of statement (3).

Statement (4) is proved by similar arguments. □

Proposition 7.4 asserts that all the neighbors of the vertices from A_i and from B_i are contained in the set $B[i-2, i+2]$ and $A[i-2, i+2]$, respectively. The next proposition describes the relations that hold between the neighborhoods of the vertices from $B[i-2, i+2]$ restricted to the set A_i and between the neighborhoods of the vertices from $A[i-2, i+2]$ restricted to the set B_i.

We say that two *sets* X and Y are *comparable* if X an Y are comparable with respect to $\subseteq$-relation (i.e. $X \subseteq Y$ or $Y \subseteq X$ holds).

Proposition 7.5. *Let $i \in [\![m-1]\!]$. For $(i \pm 2, i \pm 1) \in \{(i-2, i-1), (i+2, i+1)\}$, the following hold:*

(1) For every $w, w' \in B_{i\pm 2} \cup A_{i\pm 1}$ the sets $N(w) \cap A_i$ and $N(w') \cap A_i$ are comparable.
Moreover, if $w \in B_{i\pm 2}$ and $w' \in A_{i\pm 1}$, then $N(w) \cap A_i \subseteq N(w') \cap A_i$.

(2) For every $u, u' \in A_{i\pm 2} \cup B_{i\pm 1}$ the sets $N(u) \cap B_i$ and $N(u') \cap B_i$ are comparable.
Moreover, if $u \in A_{i\pm 2}$ and $u' \in B_{i\pm 1}$, then $N(u) \cap B_i \subseteq N(u') \cap B_i$.

Proof. To prove (1), we consider the case $(i\pm 2, i\pm 1) = (i-2, i-1)$, as the other one follows by symmetry. Suppose that $w, w' \in B_{i-2} \cup A_{i-1}$ are such that neither $N(w) \cap A_i \subseteq N(w') \cap A_i$ nor $N(w') \cap A_i \subseteq N(w) \cap A_i$ holds. It means that there are $u, u' \in A_i$ such that $wu \in E$, $w'u' \in E$, $wu' \notin E$, and $w'u \notin E$. Since $w, w' \in B_{i-2} \cup A_{i-1}$, we have $c_{i-2}w, c_{i-2}w' \in E$ and $c_{i-4}w, c_{i-3}w, c_{i-4}w', c_{i-3}w' \notin E$. Furthermore, $ww' \notin E$ and $uu' \notin E$ as G contains no triangle. Consequently, the set $\{c_{i-3}, w, w', c_{i-4}, c_{i-2}, u, u'\}$ induces a copy of T_2 in G, which cannot be the case. Moreover, if $w \in B_{i-2}$, $w' \in A_{i-1}$, then since $c_i \in (N(w') \cap A_i) \setminus (N(w) \cap A_i)$, the latter statement holds.

To show (2), we again only consider the case $(i\pm 2, i\pm 1) = (i-2, i-1)$. Suppose that $u, u' \in A_{i-2} \cup B_{i-1}$ are such that neither $N(u') \cap B_i \subseteq N(u) \cap B_i$ nor $N(u) \cap B_i \subseteq N(u') \cap B_i$ holds. It means that there are $w, w' \in B_i$ such that $uw, u'w' \in E$ and $u'w, uw' \notin E$. Since $u, u' \in A_{i-2} \cup B_{i-1}$, we have $uc_{i-1}, u'c_{i-1} \in E$ and $uc_{i+1}, u'c_{i+1} \notin E$. Furthermore, $uu' \notin E$ and $ww' \notin E$ as G contains no triangle. Hence, the set $\{c_{i-1}, w, w', c_{i+1}, u, c_i, u'\}$ induces a copy of X_3 in G, which cannot be the case.

To see the second part of the statement, assume that $N(u) \cap B_i \nsubseteq N(u') \cap B_i$ for some $u \in A_{i-2}$, $u' \in B_{i-1}$. That is, there is $w \in B_i$ such that $uw \in E$ and $u'w \notin E$. In particular, it means that $u \neq c_{i-2}$. Note that $uc_{i-1}, u'c_{i-1} \in E$. Consequently, the set $\{c_{i-3}, c_{i-2}, c_{i-1}, c_i, u, u', w\}$ induces a copy of X_3 in G, which is a contradiction. ☐

Proposition 7.5 allows us to order vertices of A_i based on two properties. We now define relation $<_{A_i}$ which combines them and we show that $<_{A_i}$ is a partial order (see Figure 7.3 for an illustration). We define for every $u, u' \in A_i$:

$$u <_{A_i} u' \text{ iff } \begin{array}{l} \text{there is } w \in B_{i-2} \cup A_{i-1} \text{ such that } u \in N(w) \text{ and } u' \notin N(w), \text{ or} \\ \text{there is } w \in A_{i+1} \cup B_{i+2} \text{ such that } u' \in N(w) \text{ and } u \notin N(w), \end{array}$$

Similarly, we define a relation $<_{B_i}$ for every $w, w' \in B_i$:

$$w <_{B_i} w' \text{ iff } \begin{array}{l} \text{there is } u \in A_{i-2} \cup B_{i-1} \text{ such that } w \in N(u) \text{ and } w' \notin N(u), \text{ or} \\ \text{there is } u \in B_{i+1} \cup A_{i+2} \text{ such that } w' \in N(u) \text{ and } w \notin N(u). \end{array}$$

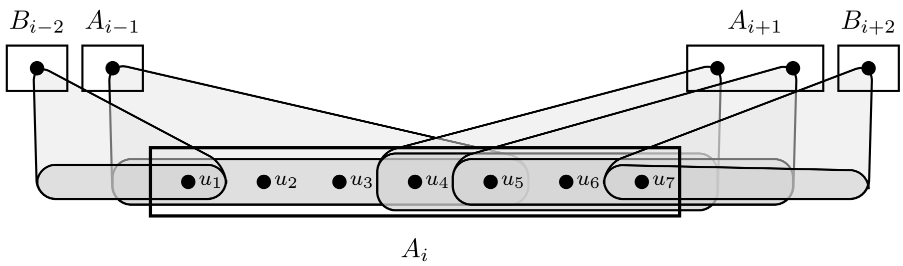

Figure 7.3: The neighborhoods of the vertices from $B_{i-2} \cup A_{i-1} \cup A_{i+1} \cup B_{i+2}$ restricted to A_i. We have $u_1 <_{A_i} \{u_2, u_3\} <_{A_i} u_4 <_{A_i} u_5 <_{A_i} u_6 <_{A_i} u_7$.

Proposition 7.6. *The following statements hold for every $i \in [\![m-1]\!]$:*

(1) $(A_i, <_{A_i})$ is a strict partial order. Moreover, $u, u' \in A_i$ are incomparable in $(A_i, <_{A_i})$ if and only if $N(u) = N(u')$.

(2) $(B_i, <_{B_i})$ is a strict partial order. Moreover, $w, w' \in B_i$ are incomparable in $(B_i, <_{B_i})$ if and only if $N(w) = N(w')$.

PROOF. Let $i \in [\![m-1]\!]$ be fixed. To prove that $(A_i, <_{A_i})$ is a strict partial order, we need to show that $<_{A_i}$ is irreflexive and transitive. The irreflexivity follows from the definition, in aim to show transitivity, we first prove that $<_{A_i}$ is antisymmetric. Suppose for a contradiction that there are $u, u' \in A_i$ such that $u <_{A_i} u'$, and $u' <_{A_i} u$. Suppose $u <_{A_i} u'$ is witnessed by a vertex $w \in B_{i-2} \cup A_{i-1}$ such that $u \in N(w)$ and $u' \notin N(w)$; the other case $w \in A_{i+1} \cup B_{i+2}$ is analogous. By Proposition 7.5.(1), there is no $w' \in B_{i-2} \cup A_{i-1}$ such that $u' \in N(w')$ and $u \notin N(w')$. Hence, since $u' <_{A_i} u$, there must be a vertex $w' \in A_{i+1} \cup B_{i+2}$ such that $u \in N(w')$ and $u' \notin N(w')$. We have $uc_{i+1}, u'c_{i+1} \in E$ and $uc_{i+2}, u'c_{i+2}, uc_{i+3}, u'c_{i+3} \notin E$ as $\{u, u'\} \subseteq A_i$. We have also $wc_{i+1}, wc_{i+2}, wc_{i+3}, w'c_{i+1}, w'c_{i+3} \notin E$ and $w'c_{i+2} \in E$ as $w \in B_{i-2} \cup A_{i-1}$ and $w' \in A_{i+1} \cup B_{i+2}$. Moreover, $uu', ww' \notin E$, by Proposition 7.4. Consequently, the set $\{w, w', c_{i+1}, c_{i+3}, u, u', c_{i+2}\}$ induces a copy of X_2 in G, which cannot be the case.

To show transitivity, suppose for a contradiction that there are vertices $u, u', u'' \in A_i$ such that $u <_{A_i} u'$ and $u' <_{A_i} u''$, but $u <_{A_i} u''$ does not hold. Suppose $u <_{A_i} u'$ is witnessed by a vertex $w \in B_{i-2} \cup A_{i-1}$ such that $u \in N(w)$ and $u' \notin N(w)$; the other case $w \in A_{i+1} \cup B_{i+2}$ is symmetric. We have $u'' \in N(w)$ as otherwise $u <_{A_i} u''$, by definition. Suppose $u' <_{A_i} u''$ is witnessed by a vertex $w' \in B_{i-2} \cup A_{i-1} \cup A_{i+1} \cup B_{i+2}$. Note that if $w' \in B_{i-2} \cup A_{i-1}$, then $u'' \notin N(w')$ and $u' \in N(w')$, which enforces also $u \in N(w')$ as $u <_{A_i} u'$ and we already proved that $<_{A_i}$ is antisymmetric. Thus, $u \in N(w')$ and $u'' \notin N(w')$, which shows $u <_{A_i} u''$. Hence, we must have $w' \in A_{i+1} \cup B_{i+2}$, and so $u' \notin N(w')$ and $u'' \in N(w')$. Moreover, $u \in N(w')$ as otherwise $u <_{A_i} u''$. As $\{u', u''\} \subseteq A_i$, $w \in B_{i-2} \cup A_{i-1}$, and $w' \in A_{i+1} \cup B_{i+2}$, we have $u'u'', ww' \notin E$, by Proposition 7.4. Consequently, $\{w, w', c_{i+1}, c_{i+3}, u', u'', c_{i+2}\}$ induces a copy of X_2 in G, which is not possible. We conclude that $(A_i, <_{A_i})$ is a strict partial order.

By definition, if $N(u) = N(u')$, then u and u' are incomparable in $(A_i, <_{A_i})$. Hence, for the second statement of (1), it is enough to show that $N(u) \neq N(u')$ implies that u and u' are comparable in

$(A_i, <_{A_i})$. Let w be a vertex such that $wu \in E$ and $wu' \notin E$. By Proposition 7.4.(2) and (3), $w \in B_{i-2} \cup A_{i-1} \cup A_{i+1} \cup B_{i+2}$. However, if $w \in B_{i-2} \cup A_{i-1}$ then $u <_{A_i} u'$ and if $w \in A_{i+1} \cup B_{i+2}$ then $u' <_{A_i} u$, by definition. Hence, u and u' are comparable in $<_{A_i}$.

The proof of (2) is similar. For antisymmetry, suppose that we have $w, w' \in B_i$ such that $w <_{B_i} w'$ and $w' <_{B_i} w$. Let $w <_{B_i} w'$ and $w' <_{B_i} w$ be witnessed by u and u' from $A_{i-2} \cup B_{i-1} \cup B_{i+1} \cup A_{i+2}$, respectively. Analogously to (1), by Proposition 7.5.(2), we can assume that $u \in A_{i-2} \cup B_{i-1}$ and $u' \in B_{i+1} \cup A_{i+2}$ and $uw, u'w \in E$, $uw', u'w' \notin E$. Observe that the set $\{c_{i-1}, w, w', c_{i+1}, u, c_i, u'\}$ induces a copy of X_3 in G, a contradiction.

For transitivity of $<_{B_i}$, suppose that for some w, w', and $w'' \in B_i$ we have $w <_{B_i} w'$ and $w' <_{B_i} w''$, but $w <_{B_i} w''$ does not hold. By symmetry of the proof of (1), we reach the case $u \in A_{i-2} \cup B_{i-1}$ and $u' \in B_{i+1} \cup A_{i+2}$, $uw, uw'', u'w, u'w'' \in E$ and $uw', u'w' \notin E$. Then, one can easily check that the set $\{w', w'', c_{i+1}, u, c_i, u', c_{i+2}\}$ induces a copy of X_2 in G, a contradiction.

Now, assume that $N(w) \neq N(w')$. Without loss of generality assume that there exists $u \in A_{i-2} \cup B_{i-1} \cup B_{i+1} \cup A_{i+2}$ such that $u \in N(w) \setminus N(w')$. Analogously as before, observe that if $u \in A_{i-2} \cup B_{i-1}$ then $w <_{B_i} w'$ and if $u \in B_{i+1} \cup A_{i+2}$ then $w' <_{B_i} w$. Therefore, w and w' are comparable, which finishes the proof. □

Finally, for every $i \in [\![m-1]\!]$ we order arbitrarily the elements inside every antichain of $(A_i, <_{A_i})$ and of $(B_i, <_{B_i})$, obtaining strict linear orders $(A_i, <_{A_i})$ and $(B_i, <_{B_i})$. We introduce a binary relation $\prec$ defined on the set V, such that $v \prec v'$ for $v, v' \in V$ if one of the following conditions holds for some $i \in [\![m-1]\!]$:

- $v, v' \in A_i$, $v <_{A_i} v'$, and v, v' are consecutive in $(A_i, <_{A_i})$,
- $v, v' \in B_i$, $v <_{B_i} v'$, and v, v' are consecutive in $(B_i, <_{B_i})$,
- v is the maximum of $(A_i, <_{A_i})$ and v' is the minimum of $(B_{i+1}, <_{B_{i+1}})$,
- v is the maximum of $(B_i, <_{B_i})$ and v' is the minimum of $(A_{i+1}, <_{A_{i+1}})$.

Informally, to get an embedding of G into a cylinder (when the shortest hole is even) or into a Möbius strip (when the shortest hole is odd) which locally satisfies the adjacency and the enclosure properties, we place the vertices v, v' satisfying $v \prec v'$ next to each other, v before v' assuming that the border of the cylinder or the Möbius strip are oriented as shown in Figure 1.5 (reproduced below for convenience). In what follows we extend $\prec$ relation as follows:

- For every $V' \subsetneq V$ by $<_{V'}$ we denote the transitive closure of $\prec$ restricted to V',
- For $v, v' \in V$ we set $v <_{cl} v'$ if $v, v' \in A[i-2, i+2]$ and $v <_{A[i-2,i+2]} v'$ for some $i \in [\![m-1]\!]$ or $v, v' \in B[i-2, i+2]$ and $v <_{B[i-2,i+2]} v'$ for some $i \in [\![m-1]\!]$.

Finally, the following lemma characterizes the global structure of an almost bipartite permutation graph.

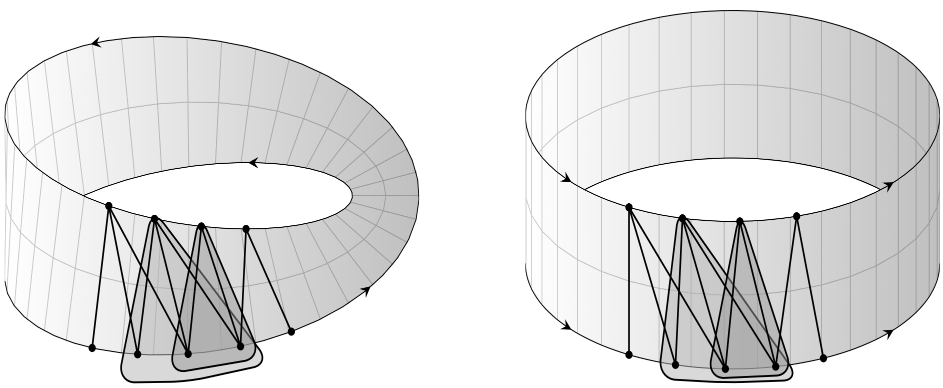

Figure 1.5 (reproduced): An embedding of a connected almost bipartite permutation graph in a cylinder or a Möbius strip that locally satisfies the adjacency and enclosure properties.

Lemma 7.7. *Let i, j be such that $i \leqslant j$, $|j-i| = m-3$. Let $U = A[i,j]$ and $W = B[i,j]$. Then $G[U \cup W]$ is a bipartite permutation graph with bipartition classes U and W.*

Moreover, $(U, <_U)$ and $(W, <_W)$ are strict linear orders that satisfy the adjacency and enclosure properties in $G[U \cup W]$.

PROOF. Proposition 7.4 asserts there is no edge between a vertex in $V[j-1,j]$ and a vertex in $V[i,i+1]$. In particular, $G[U \cup W]$ is a bipartite graph and $(U, <_U)$ and $(W, <_W)$ are strict linear orders. Given Theorem 7.1.(c), to prove the lemma we need to show that $(W, <_W)$ satisfies the adjacency and enclosure properties in $G[U \cup W]$.

To prove the adjacency property, consider $u \in A_k \subseteq U$ for some suitable k. Recall that by Proposition 7.4.(3), $B_k \subseteq N(u) \subseteq B[k-2, k+2]$. To show that $N(u)$ consists of consecutive vertices in W it suffices to note that:

- if $w \in A_{k+1}$, $w' \in B_{k+2}$ and $uw' \in E$ then $uw \in E$, by Proposition 7.5,
- if $w, w' \in A_{k+1}$ (resp. $w, w' \in B_{k+2}$) are such that $w <_{A_{k+1}} w'$ (resp. $w <_{B_{k+2}} w'$) and $uw' \in E$, then $uw \in E$, by Proposition 7.6,

and that analogous statements hold by symmetry for the part of $N(u)$ contained in $A_{k-1} \cup B_{k-2}$. If $u \in B_k \subseteq U$, the case analysis is similar (one needs to swap letters A and B in the reasoning above). Therefore, the adjacency property is proved.

To show that $(W, <_W)$ satisfies the enclosure property assume for a contradiction that there are $w, w', w'' \in W$ and $u, u' \in U$ such that $N(u') \subseteq N(u)$, $w <_W w' <_W w''$ and $uw, uw', uw'' \in E$, $u'w' \in E$, and $u'w, u'w'' \notin E$.

Claim 7.1. *There is $k \in [i,j]$ such that either $u, u' \in A_k$, or $u, u' \in B_k$.*

PROOF. If $u \in B_k$, then since $N(u') \cap C \subseteq N(u) \cap C = \{c_k\}$, we have $u' \in B_k$, so the claim holds. Therefore, assume that $u \in A_k$, and suppose that $u' \notin A_k$. Then $N(u') \cap C \subseteq N(u) \cap C = \{c_{k-1}, c_{k+1}\}$. Assuming $u <_U u'$ (the other case is symmetric), we have that $u' \in B_{k+1}$. Due to Proposition 7.4 and $w' <_W w''$ we have $w', w'' \in A[k-1, k+2]$. Moreover, as we already proved that $N(u')$ is consecutive in $(W, <_W)$ (adjacency property), and $c_{k+1} \in N(u')$, we have $c_{k+1} <_W w''$. Therefore $w'' \in A_{k+1} \cup B_{k+2}$. Note that:

- if $w'' \in B_{k+2}$, then, since $uw'' \in E$, we have that $w'' \in N(u) \cap B_{k+2}$. However, by Proposition 7.5.(2), we have $N(u) \cap B_{k+2} \subseteq N(u') \cap B_{k+2}$, so it implies that $u'w'' \in E$, a contradiction,
- if $w'' \in A_{k+1}$, then by Proposition 7.4.(2) we would have $u'w'' \in E$, a contradiction.

This concludes the proof of claim. ⌟

Suppose $u, u' \in A_k$. Since $u'w, u'w'' \notin E$ and $u'c_{k-1}, u'c_{k+1} \in E$, we have by adjacency property of $(W, <_W)$ that $w <_W c_{k-1} <_W c_{k+1} <_W w''$. Therefore, we must have $w \in B_{k-2} \cup A_{k-1}$ and $w'' \in A_{k+1} \cup B_{k+2}$. Observe that w'' witnesses that $u' <_{A_k} u$ by definition, however, w witnesses the opposite, that is $u <_{A_k} u'$. We have a contradiction by Proposition 7.6.

Suppose $u, u' \in B_k$. An analysis, which is analogous to the one in the previous case (again, it is enough to swap letters A and B in that reasoning above), gives us that we must have $w \in A_{k-2} \cup B_{k-1}$ and $w'' \in B_{k+1} \cup A_{k+2}$. Again, we obtain a contradiction by the definition of $<_{B_k}$ and Proposition 7.6. □

Lemma 7.7 provides an interesting view on classification of almost bipartite permutation graphs. Specifically, if m is even, then the graph may be drawn on a cylinder, whose boundary consists of two closed curves, one of which traverses the vertices of $A_0, B_1, \ldots A_{m-2}, B_{m-1}$, and the second one — the vertices of $B_0, A_1, \ldots, B_{m-2}, A_{m-1}$. If in turn m is odd, then the graph can be represented on a Möbius strip, whose boundary traverses the vertices of $A_0, B_1, \ldots, B_{m-2}, A_{m-1}$ and then $B_0, A_1, \ldots A_{m-2}, B_{m-1}$. In both cases we draw the vertices of A_i and of B_i on the opposite side of the strip according to the orders given by $<_{A_i}$ and $<_{B_i}$, for $i \in [\![m-1]\!]$ (recall Figure 1.5).

The following definitions are taken from [vtHV13]. A *hole cut* of G is a vertex set $X \subseteq V$ such that $G - X$ is a bipartite permutation graph. Lemma 7.7 asserts that for every $i \in [\![m-1]\!]$ the set $V[i, i+1]$ is a hole cut in G. A hole cut X of G is *minimum* if G does not have a hole cut whose size is strictly smaller than the size of X. A hole cut X of G is *minimal* if any proper subset of X is not a hole cut in G.

The next proposition describes the structure of every hole in G.

Proposition 7.8. *Suppose C' is a hole of length t in G for some $t \geqslant m$. Then, the consecutive vertices of C' can be labeled by $c'_0, c'_1, \ldots, c'_{t-1}$ so as the following conditions hold (the indices are taken modulo k):*

- $c'_i c'_{i+1} \in E$ *for every* $i \in [\![t-1]\!]$,
- $c'_i <_{cl} c'_{i+2}$ *for every* $i \in [\![t-1]\!]$,
- $\{c'' \in C' : c'_i <_{cl} c'' <_{cl} c'_{i+2}\} = \emptyset$ *for every* $i \in [\![t-1]\!]$.

Proof. Let c'_0, c'_1, c'_2, …, c'_{t-1} be consecutive vertices of C' denoted in such a way that $c'_0 <_{cl} c'_2$. We can assume it, since $c'_0, c'_2 \in N(c'_1)$, thus, by Proposition 7.4.(3) and (4), both c'_0, c'_2 belong to $A[\ell-2, \ell+2]$ or both belong to $B[\ell-2, \ell+2]$ for some $\ell \in [\![m-1]\!]$.

Now, we show that if there exists $j \in [\![t-1]\!]$ such that $c'_j <_{cl} c'_{j+2}$, then $c'_{j+1} <_{cl} c'_{j+3}$. Suppose, for contradiction that $c'_j <_{cl} c'_{j+2}$ and $c'_{j+1} \not<_{cl} c'_{j+3}$. Let $i \in [\![m-1]\!]$ be such that $c'_{j+2} \in A_i \cup B_i$. Similarly, as $c'_{j+1}, c'_{j+3} \in N(c'_{j+2})$, either $c'_{j+1}, c'_{j+3} \in B[i-2, i+2]$ if $c'_{j+2} \in A_i$ or $c'_{j+1}, c'_{j+3} \in A[i-2, i+2]$ if $c'_{j+2} \in B_i$. In both cases, Lemma 7.7 implies that $<_{cl}$ restricted to $V[i-4, i+2]$ is a strong ordering of $G[V[i-4, i+2]]$. Moreover, c'_{j+1}, c'_{j+3} are comparable in $<_{cl}$, by Proposition 7.4.(3) or (4), and the definition of $<_{cl}$. Since we assumed that $c'_{j+1} \not<_{cl} c'_{j+3}$, we must have $c'_{j+3} <_{cl} c'_{j+1}$, and from Theorem 7.1.(b) we get that $c'_j c'_{j+3} \in E$, so C' has a chord — contradiction. Therefore $c'_j <_{cl} c'_{j+2}$ implies $c'_{j+1} <_{cl} c'_{j+3}$ for every integer j. Applying the above observation repeatedly for $j = 0, 1, 2, \ldots$, we get that $c'_0 <_{cl} c'_2 <_{cl} c'_4 <_{cl} \ldots$ and $c'_1 <_{cl} c'_3 <_{cl} c'_5 <_{cl} \ldots$

For the last property, suppose for the sake of contradiction that there exists $j \notin \{i, i+2\}$ such that $c'_i <_{cl} c'_j <_{cl} c'_{i+2}$ (clearly, $j \neq i+1$, since c'_{i+1} belongs to the opposite side of the local bipartition). Then, by Lemma 7.7, $c'_j c'_{i+1} \in E$ due to the adjacency property. But then the edge $c'_j c'_{i+1}$ is a chord in C'. This completes the proof. □

The structure of holes described above implies that for every $i \in [\![m-1]\!]$ the sets $A[i-2, i+2]$ and $B[i-2, i+2]$ are hole cuts. We use this observation to prove the following statement about minimal hole cuts in G.

Proposition 7.9. *Every minimal hole cut X in G is fully contained in the set $V[i-2, i+2]$ for some $i \in [\![m-1]\!]$.*

Proof. Let X be a minimal hole cut. Since X is minimal, we can choose elements z_1, x_1, x_2, z_2 in V such that the following conditions hold:

- we have $z_1 \prec x_1 \leqslant_{cl} x_2 \prec z_2$, the set $X' = \{x : x_1 \leqslant_{cl} x \leqslant_{cl} x_2\}$ is non-empty and is contained in X, and the elements z_1, z_2 are not in X.

Note that either $\{z_1, z_2\} \cup X' \subseteq B[i-2, i+2]$ or $\{z_1, z_2\} \cup X' \subseteq A[i-2, i+2]$ for some $i \in [\![m-1]\!]$. Otherwise, we have $B[j, j+3] \subseteq X'$ or $A[j, j+3] \subseteq X'$ for some $j \in [\![m-2]\!]$. However, Proposition 7.8 and Proposition 7.4.(3) and (4) imply that the sets $A[j, j+3]$ and $B[j, j+3]$ are hole cuts for every $j \in [\![m-1]\!]$. Otherwise, the set $A[j, j+3]$ ($B[j, j+3]$) is contained in the neighbourhood of some vertex of a hole avoiding $A[j, j+3]$ ($B[j, j+3]$, respectively) in G and the neighborhood of this vertex intersects also B_{j-1} and A_{j+4} (A_{j-1} and B_{j+4}, respectively), which cannot be the case due to Proposition 7.4.(3) and (4). So, we have either $B[j, j+3] = X' = X$ or $A[j, j+3] = X' = X$ as X is a minimal hole cut. But then, we have $X \subseteq V[j, j+3]$, which completes the proof of our claim.

Therefore, for the rest of the proof we assume $\{z_1, z_2\} \cup X' \subseteq B[i-2, i+2]$; the other case is proved similarly. Moreover, we may assume that i is picked such that:

- $z_1 \in B[i-2, i]$, $z_2 \in B[i, i+2]$, and $X' \subseteq B[i-2, i+2]$.

See Figure 7.4 for an illustration.

Suppose Y' is the set consisting of all the neighbors of both z_1 and z_2; that is, $Y' = N(z_1) \cap N(z_2)$. Clearly, we have $Y' \subseteq A[i-2, i+2]$. To complete the proof of the proposition, we show that:

- every element of Y' is a member of X,
- $X' \cup Y'$ is a hole cut in G.

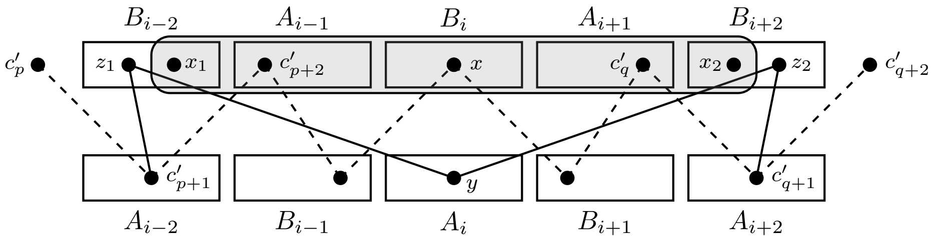

Figure 7.4: Illustration of the proof: the cycle C' is marked with a dashed line. The set X' is shaded.

Then, since $X' \subseteq X$, we have $X' \cup Y' = X$ by minimality of X and consequently $X \subseteq V[i-2, i+2]$. So, it remains to prove the claims about the set Y'.

Suppose we have $y \in Y'$ such that $y \notin X$. We fix some $x \in X'$. Clearly, since X is a minimal hole cut, there is a hole C' in $G - (X \setminus \{x\})$. Note that C' must contain x. Suppose $c'_0, \dots, c'_{\ell-1}$ for some $\ell \geqslant 9$ are consecutive vertices in C' chosen such that $c'_j <_{cl} c'_{j+2}$ for every $j \in [\![\ell - 1]\!]$ (indices are taken modulo ℓ). Now we pick $p, q \in [\![\ell - 1]\!]$ such that $c'_p <_{cl} z_1 \leqslant_{cl} c'_{p+2}$ and $c'_q \leqslant_{cl} z_2 <_{cl} c'_{q+2}$. Since $x \in C'$, we have $c'_{p+2} \leqslant_{cl} x$ and $x \leqslant_{cl} c'_q$. Note that c'_{p+1} is adjacent to z_1 and c'_{q+1} is adjacent to z_2 due to the adjacency property. Next we replace in C' all the vertices between c'_{p+2} and c'_q (inclusive; this set includes x) with the vertices z_1, y, z_2 and we obtain a cycle C'' containing no elements from X. Clearly, we can easily find a hole among the elements from C'' that avoids all the elements from X. E.g. we can consider a shortest "noncontractible" subcycle of C'' (i.e. not contained in any set of the form $V[t+1, t]$). It will be chordless and of length $\geqslant 5$ (since $m \geqslant 10$). This yields a contradiction as X is a hole cut.

To prove the second claim, suppose there is a hole C' in $G - (X' \cup Y')$. By Proposition 7.8 there are $c'_1, c'_2, c'_3 \in C'$ such that $c'_1 <_{cl} X' <_{cl} c'_3$ and $c'_1, c'_3 \in N(c'_2)$. However, since $c'_1 \leqslant_{cl} z_1 <_{cl} z_2 \leqslant_{cl} c'_3$ and $c'_1, c'_3 \in N(c'_2)$, we have $z_1 \in N(c'_2)$ and $z_2 \in N(c'_2)$. So, we have $c'_2 \in Y'$, which is a contradiction. ☐

7.3 Proof of Theorem I

The aim of this section is to provide a complete proof of Theorem I using structural results from the previous section. Let us start by showing that the Bipartite Permutation Vertex Deletion problem can be decided in polynomial time on almost bipartite permutation graphs.

Lemma 7.10. *Let (G, k) be an instance of* Bipartite Permutation Vertex Deletion *where G is an n-vertex almost bipartite permutation graph. Then whether (G, k) is a* **yes***-instance can be decided in time $\mathcal{O}(n^6)$.*

Proof. If G is a bipartite permutation graph, (G, k) is a **yes**-instance, thus, we are done in this case, e.g. by using the linear time recognition algorithm of [SBS87]. If G is not connected, we can consider each connected component independently and, at the end, we compare k with the total number of deleted vertices over all components. Let G' be a connected r-vertex component of G such that G' is not a bipartite permutation graph (otherwise, clearly, no vertex needs to be deleted). Let $C = \{c_0, \dots, c_{m-1}\}$ be a shortest hole in G' (it exists as G' is not a bipartite permutation graph). It can be found in time $\mathcal{O}(r^6)$ as follows. We iterate over all possible four-element subsets $S = \{v_1, v_2, v_3, v_4\}$ of $V(G')$. For these S for which $G'[S]$ is an induced P_4, with consecutive vertices v_1, v_2, v_3, v_4, we construct a graph $\widetilde{G'}$ by removing the vertices from $(N(v_2) \cup N(v_3)) \setminus \{v_1, v_4\}$ (note that v_2 and v_3 also get removed). Then we find a shortest v_1-v_4-path in $\widetilde{G'}$ in time $\mathcal{O}(r^2)$.

Having found C, we may introduce the notation concerning structure around C (i.e. as in Section 7.2. By Proposition 7.9, every minimal hole cut X in G' is contained in the set $V' = V_{G'}[i-2, i+2]$ for some $i \in [\![m-1]\!]$. Therefore, we may check all the possibilities where a minimal cut is contained. For every i, we run an algorithm for finding a maximum flow in the following digraph H_i.

Digraph H_i has the vertex set $V' \times \{\mathsf{in}, \mathsf{out}\} \cup \{s, t\}$ and arc set consisting of:

- all arcs of the form $(u, \mathsf{out})(v, \mathsf{in})$, where uv is an edge of $G'[V']$,
- $s(v, \mathsf{in})$ if there exists $u \in V_{G'}[i-4, i-3]$ such that uv is an edge of G',
- $(u, \mathsf{out})t$ if there exists $v \in V_{G'}[i+3, i+4]$ such that uv is an edge of G',
- $(u, \mathsf{in})(u, \mathsf{out})$ for all $u \in V'$.

Set capacities of arcs of the form $(u, \mathsf{in})(u, \mathsf{out})$ to 1 and capacities of all the remaining arcs to ∞ (practically $|V_{G'}|$). It is readily seen that minimum (s,t)-cut in the defined network H_i corresponds to a minimum hole cut in $G'[V']$ (arc of unit capacity $(u, \mathsf{in})(u, \mathsf{out})$ naturally corresponds to the vertex u of G').

Therefore it remains to apply any classical max-flow algorithm to each H_i for $i \in [\![m-1]\!]$ and remember the smallest size $k_{G'}$ of minimal (s,t)-cuts. This can be performed in time $\mathcal{O}(m \cdot (|V'|+2) \cdot (|E_{G'[V']}| + 2|V'|)^2) = \mathcal{O}(r^6)$ [EK72]. Finally, (G,k) is a **yes**-instance if and only if the sum of remembered sizes $k_{G'}$ over the all considered connected components G' is at most k. Clearly, the total running time is $\mathcal{O}(n^6)$.□

We now propose the algorithm. Given an n-vertex graph $G = (V,E)$ and number k, we want to answer the Bipartite Permutation Vertex Deletion problem. We say that (G,k) is the *initial* instance. We split our algorithm into two parts. The first part consists of a branching algorithm for deletion to almost bipartite permutation graphs. The output of the first part is a set of instances (G', k') where G' is an almost bipartite permutation graph and $0 \leqslant k' \leqslant k$ (or **no**-answer is no such instance exists) such that the initial instance (G,k) is a **yes**-instance if and only if at least one of these instances is a **yes**-instance. In the second part, the algorithm runs an $\mathcal{O}(n^6)$-time algorithm for Bipartite Permutation Vertex Deletion for each instance (G', k') output by the first phase.

Let us start with the first part. We say that $X \subseteq V$ is a *forbidden set* if $G[X]$ is isomorphic to one of the graphs: $K_3, T_2, X_2, X_3, C_5, C_6, C_7, C_8, C_9$. We define the following rule.

Rule : Given an instance (G,k), $k \geqslant 1$, and a minimal forbidden set X, branch into $|X|$ instances, $(G - v, k-1)$ for each $v \in X$.

Starting with the initial instance, the algorithm applies the rule exhaustively. In other words, the algorithm forms a branching tree with leaves corresponding to instances (G', k') where $k' = 0$ or G' is an almost bipartite permutation graph. Clearly, as at least one vertex from each forbidden set must be removed from G, the initial instance is a **yes**-instance if and only if at least one of the leaves is a **yes**-instance.

The algorithm continues to the second part only with leaves (G', k') such that G' is an almost bipartite permutation graph (as otherwise, the leaf is a **no**-instance). It runs the algorithm described in Lemma 7.10 to find if G' can be transformed into a bipartite permutation graph by using at most k' vertex deletions. It either finds a **yes**-instance or concludes after checking all the instances that there is no solution; that is, the initial instance is a **no**-instance.

We note that such a branching into a bounded number of smaller instances is a standard technique, see e.g., [vtHV13] for more details.

We now analyze the running time of the whole algorithm. In the first part, observe that the branching tree has depth at most k and has at most 9^k leaves, as k decreases by one whenever the algorithm branches and each of the listed forbidden subgraphs has at most nine vertices. Therefore the total number of nodes in the branching tree is $\mathcal{O}(9^k)$. Moreover, in each node (G'', k'') the algorithm works in time $\mathcal{O}(n^9)$ as it checks if G'' contains a forbidden set of size at most 9. So, the first part works in time $\mathcal{O}(9^k \cdot n^9)$. In the second part, the algorithm does a work $\mathcal{O}(n^6)$ in each of at most 9^k leaves, by Lemma 7.10. Thus, the second part works in time $\mathcal{O}(9^k \cdot n^6)$. We conclude that the total running time of our algorithm for Bipartite Permutation Vertex Deletion is $\mathcal{O}(9^k \cdot n^9)$.

7.4 Proof of Theorem J

In this section, we provide a proof of Theorem J. The idea of the algorithm is very similar to the FPT algorithm described in Lemma 7.10.

Let $G = (V,E)$ be a graph and let $Y \subseteq V$ be a subset of vertices of G such that $G - Y$ is a bipartite permutation graph. We want to construct a set $Z \subseteq V$ in polynomial time such that $G - Z$ is a bipartite permutation graph and $|Z| \leqslant 9|Y|$. We construct Z as follows. We start with $Z = \emptyset$. Then, as long as $G - Z$ contains a set X isomorphic to one of $K_3, T_2, X_2, X_3, C_5, C_6, C_7, C_8, C_9$ we add all vertices of X to Z. Observe that $Y \cap X \neq \emptyset$ and $|X| \leqslant 9$.

After this step $G - Z$ is an almost bipartite permutation graph. Note that $|Z| \leqslant 9|Z \cap Y|$. We find a shortest hole $C = \{c_0, \ldots, c_{m-1}\}$ in $G - Z$ and find a minimum hole cut X as described in Section 7.3. Since $Y - Z$ is a hole cut in $G - Z$ we have $|X| \leqslant |Y - Z|$. We add X to Z. Observe that $G - Z$ is a bipartite permutation graph.

Since $K_3, T_2, X_2, X_3, C_5, C_6, C_7, C_8, C_9$ have at most 9 vertices, we have that $|Z| \leqslant 9|Y|$. This implies that the above algorithm is a 9-approximation algorithm. It runs in polynomial time because finding small forbidden subgraphs can be done in polynomial time and finding a minimum hole cut in an almost bipartite permutation graph can be done in polynomial time.

Bibliography

[AKKW16] Saeed Akhoondian Amiri, Ken-ichi Kawarabayashi, Stephan Kreutzer, and Paul Wollan. The Erdős-Pósa property for directed graphs. 2016. Available at: arxiv.org/abs/1603.02504.

[BDK+22] Łukasz Bożyk, Jan Derbisz, Tomasz Krawczyk, Jana Novotná, and Karolina Okrasa. Vertex deletion into bipartite permutation graphs. *Algorithmica*, 84(8):2271–2291, Aug 2022.

[BDOP22a] Łukasz Bożyk, Oscar Defrain, Karolina Okrasa, and Michał Pilipczuk. On objects dual to tree-cut decompositions. *Journal of Combinatorial Theory, Series B*, 157:401–428, 2022.

[BDOP22b] Łukasz Bożyk, Oscar Defrain, Karolina Okrasa, and Michał Pilipczuk. On digraphs without onion star immersions. Preprint, 2022. Available at: arxiv.org/abs/2211.15477.

[BFL+16] Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. (Meta) Kernelization. *J. ACM*, 63(5):44:1–44:69, 2016.

[BFR72] Kirby A. Baker, Peter C. Fishburn, and Fred S. Roberts. Partial orders of dimension 2. *Networks*, 2(1):11–28, 1972.

[Boj09] Mikołaj Bojańczyk. Factorization forests. In *13th International Conference on Developments in Language Theory, DLT 2009*, volume 5583 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2009.

[BP20] Marthe Bonamy and Michał Pilipczuk. Graphs of bounded cliquewidth are polynomially χ-bounded. *Advances in Combinatorics*, 2020:8, 2020.

[BP22a] Łukasz Bożyk and Michał Pilipczuk. On the Erdős-Pósa property for immersions and topological minors in tournaments. *Discrete Mathematics & Theoretical Computer Science*, 24(1), 2022.

[BP22b] Łukasz Bożyk and Michał Pilipczuk. Polynomial kernel for immersion hitting in tournaments. In *30th Annual European Symposium on Algorithms (ESA 2022)*, volume 244 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:17, 2022.

[BPP18] Florian Barbero, Christophe Paul, and Michał Pilipczuk. Exploring the complexity of layout parameters in tournaments and semicomplete digraphs. *ACM Trans. Algorithms*, 14(3):38:1–38:31, 2018.

[Cai96] Leizhen Cai. Fixed-parameter tractability of graph modification problems for hereditary properties. *Information Processing Letters*, 58(4):171–176, 1996.

[Cao16] Yixin Cao. Linear recognition of almost interval graphs. In *27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1096–1115. SIAM, 2016.

[Cao17] Yixin Cao. Unit interval editing is fixed-parameter tractable. *Information and Computation*, 253:109–126, 2017.

[CM16] Yixin Cao and Dániel Marx. Chordal editing is fixed-parameter tractable. *Algorithmica*, 75(1):118–137, 2016.

[COS12] Maria Chudnovsky, Alexandra Ovetsky Fradkin, and Paul D. Seymour. Tournament immersion and cutwidth. *Journal of Combinatorial Theory, Series B*, 102(1):93–101, 2012.

[Cou90] Bruno Courcelle. The monadic second-order logic of graphs. I. recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990.

[CRST06] Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The strong perfect graph theorem. *Annals of Mathematics*, pages 51–229, 2006.

[CS11] Maria Chudnovsky and Paul D. Seymour. A well-quasi-order for tournaments. *Journal of Combinatorial Theory, Series B*, 101(1):47–53, 2011.

[Die16] Reinhard Diestel. *Graph Theory*, volume 173. Springer-Verlag, Heidelberg, 2016.

[Din76] Efim A. Dinits. On the structure of a family of minimal weighted cuts in a graph. *Studies in Discrete Optimization*, 1976.

[DO21] Reinhard Diestel and Sang-il Oum. Tangle-tree duality in abstract separation systems. *Advances in Mathematics*, 377:107470, 2021.

[EH63] Pál Erdős and Haim Hanani. On a limit theorem in combinatorial analysis. *Publ. Math. Debrecen*, 10:10–13, 1963.

[EK72] Jack Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.

[EP65] Pál Erdős and Lajos Pósa. On independent circuits contained in a graph. *Canadian Journal of Mathematics*, 17:347–352, 1965.

[FLMS12] Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Planar $\mathcal{F}$-Deletion: Approximation, kernelization and optimal FPT algorithms. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012*, pages 470–479. IEEE Computer Society, 2012.

[FLSZ19] Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: Theory of Parameterized Preprocessing*. Cambridge University Press, 2019.

[FP19] Fedor V. Fomin and Michał Pilipczuk. On width measures and topological problems on semi-complete digraphs. *Journal of Combinatorial Theory, Series B*, 138:78–165, 2019.

[FST11] Fedor V. Fomin, Saket Saurabh, and Dimitrios M. Thilikos. Strengthening Erdős-Pósa property for minor-closed graph classes. *J. Graph Theory*, 66(3):235–240, 2011.

[Gal67] Tibor Gallai. Transitiv orientierbare Graphen. *Acta Mathematica Academiae Scientiarum Hungarica*, 18(1-2):25–66, 1967.

[GKKK20] Archontia C. Giannopoulou, Ken-ichi Kawarabayashi, Stephan Kreutzer, and O-joung Kwon. The directed flat wall theorem. In *2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020*, pages 239–258. SIAM, 2020.

[GKKK22] Archontia C. Giannopoulou, Ken-ichi Kawarabayashi, Stephan Kreutzer, and O-joung Kwon. Directed tangle tree-decompositions and applications. In *2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*, pages 377–405. SIAM, 2022.

[GPR+21] Archontia C. Giannopoulou, Michał Pilipczuk, Jean-Florent Raymond, Dimitrios M. Thilikos, and Marcin Wrochna. Linear kernels for edge deletion problems to immersion-closed graph classes. *SIAM Journal on Discrete Mathematics*, 35(1):105–151, 2021.

[GRU83] Martin Charles Golumbic, Doron Rotem, and Jorge Urrutia. Comparability graphs and intersection graphs. *Discrete Mathematics*, 43(1):37–46, 1983.

[Hie73] Carl Hierholzer. Ueber die Möglichkeit, einen Linienzug ohne Wiederholung und ohne Unterbrechung zu umfahren. *Mathematische Annalen*, 6:30–32, 1873.

[HvtHJ+13] Pinar Heggernes, Pim van 't Hof, Bart M. P. Jansen, Stefan Kratsch, and Yngve Villanger. Parameterized complexity of vertex deletion into perfect graph classes. *Theoretical Computer Science*, 511:172–180, 2013.

[JMPW23] Gwenaël Joret, Piotr Micek, Michał Pilipczuk, and Bartosz Walczak. Cliquewidth and dimension. Preprint, 2023. Available at: arxiv.org/abs/2308.11950.

[JPS+14] Gwenaël Joret, Christophe Paul, Ignasi Sau, Saket Saurabh, and Stéphan Thomassé. Hitting and harvesting pumpkins. *SIAM Journal on Discrete Mathematics*, 28(3):1363–1390, 2014.

[KK15] Ken-ichi Kawarabayashi and Stephan Kreutzer. The directed grid theorem. In *47th Annual ACM on Symposium on Theory of Computing, STOC 2015*, pages 655–664. ACM, 2015.

[KS15] Ilhee Kim and Paul D. Seymour. Tournament minors. *Journal of Combinatorial Theory, Series B*, 112:138–153, 2015.

[KST54] Tamás Kővári, Vera T. Sós, and Pál Turán. On a problem of Zarankiewicz. In *Colloquium Mathematicum*, volume 3, pages 50–57. Polish Academy of Sciences, 1954.

[Kuf08] Manfred Kufleitner. The height of factorization forests. In *33rd International Symposium Mathematical Foundations of Computer Science 2008, MFCS 2008*, volume 5162 of *Lecture Notes in Computer Science*, pages 443–454. Springer, 2008.

[Kur30] Kazimierz Kuratowski. Sur le problème des courbes gauches en topologie. *Fundamenta Mathematicae*, 15(1):271–283, 1930.

[LB62] Cornelis G. Lekkerkerker and Jan Ch. Boland. Representation of a finite graph by a set of intervals on the real line. *Fundamenta Mathematicae*, 51(1):45–64, 1962.

[LY80] John M. Lewis and Mihalis Yannakakis. The node-deletion problem for hereditary properties is NP-complete. *Journal of Computer and System Sciences*, 20(2):219–230, 1980.

[Man08] Federico Mancini. Graph modification problems related to graph classes. *PhD degree dissertation, University of Bergen Norway*, 2, 2008.

[Mar10] Dániel Marx. Chordal deletion is fixed-parameter tractable. *Algorithmica*, 57(4):747–768, 2010.

[Maz13] Frédéric Mazoit. A simple proof of the tree-width duality theorem. Technical Report, 2013. Available at: https://hal.archives-ouvertes.fr/hal-00859912/.

[Men27] Karl Menger. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10:96–1159, 1927.

[MNS17] Kurt Mehlhorn, Adrian Neumann, and Jens M. Schmidt. Certifying 3-edge-connectivity. *Algorithmica*, 77(2):309–335, 2017.

[NOP+21] Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. Rankwidth meets stability. In *2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021*, pages 2014–2033. SIAM, 2021.

[NORS21] Jaroslav Nešetřil, Patrice Ossona de Mendez, Roman Rabinovich, and Sebastian Siebertz. Classes of graphs with low complexity: The case of classes with bounded linear rankwidth. *European Journal of Combinatorics*, 91:103223, 2021. Colorings and structural graph theory in context (a tribute to Xuding Zhu).

[OS13] Alexandra Ovetsky Fradkin and Paul D. Seymour. Tournament pathwidth and topological containment. *Journal of Combinatorial Theory, Series B*, 103(3):374–384, 2013.

[OS15] Alexandra Ovetsky Fradkin and Paul D. Seymour. Edge-disjoint paths in digraphs with bounded independence number. *Journal of Combinatorial Theory, Series B*, 110:19–46, 2015.

[Ove11] Alexandra Ovetsky Fradkin. *Forbidden structures and algorithms in graphs and digraphs.* PhD thesis, Princeton University, 2011.

[PLE71] Amir Pnueli, Abraham Lempel, and Shimon Even. Transitive orientation of graphs and identification of permutation graphs. *Canadian Journal of Mathematics*, 23(1):160–175, 1971.

[Ray] Jean-Florent Raymond. Dynamic Erdős-Pósa listing. Available at: perso.limos.fr/~jfraymon/Erdős-Pósa/.

[Ray18] Jean-Florent Raymond. Hitting minors, subdivisions, and immersions in tournaments. *Discrete Mathematics & Theoretical Computer Science*, 20(1), 2018.

[RRST96] Bruce A. Reed, Neil Robertson, Paul D. Seymour, and Robin Thomas. Packing directed circuits. *Combinatorica*, 16(4):535–554, 1996.

[RS86] Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.

[RS91] Neil Robertson and Paul D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52(2):153–190, 1991.

[RS03] Neil Robertson and Paul D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.

[RS04] Neil Robertson and Paul D. Seymour. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004. Special Issue Dedicated to Professor W.T. Tutte.

[RS10] Neil Robertson and Paul D. Seymour. Graph minors XXIII. Nash-Williams' immersion conjecture. *Journal of Combinatorial Theory, Series B*, 100(2):181–205, 2010.

[SBS87] Jeremy P. Spinrad, Andreas Brandstädt, and Lorna Stewart. Bipartite permutation graphs. *Discrete Applied Mathematics*, 18(3):279–292, 1987.

[Sim90] Imre Simon. Factorization forests of finite height. *Theoretical Computer Science*, 72(1):65–94, 1990.

[ST93] Paul D. Seymour and Robin Thomas. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B*, 58(1):22–33, 1993.

[Tho82] Andrew Thomason. On finite Ramsey numbers. *European Journal on Combinatorics*, 3(3):263–273, 1982.

[VHPT09] Yngve Villanger, Pinar Heggernes, Christophe Paul, and Jan Arne Telle. Interval completion is fixed parameter tractable. *SIAM Journal on Computing*, 38(5):2007–2020, 2009.

[vtHV13] Pim van 't Hof and Yngve Villanger. Proper interval vertex deletion. *Algorithmica*, 65(4):845–867, 2013.

[Wag37] Klaus Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114:570–590, 1937.

[Wol15] Paul Wollan. The structure of graphs not admitting a fixed immersion. *Journal of Combinatorial Theory, Series B*, 110:47–66, 2015.