

Fundamental substructures of infinite graphs

Dissertation

zur Erlangung des Doktorgrades
an der Fakultät für Mathematik, Informatik
und Naturwissenschaften
der Universität Hamburg

vorgelegt

im Fachbereich Mathematik

von

Carl Phillip Bürger

Hamburg

2020

Vorsitz der Prüfungskommission: Prof. Armin Iske
Erstgutachter: Prof. Reinhard Diestel (Betreuer)
Zweitgutachter: Dr. Max Pitz

To mum and dad.

Contents

1. Introduction	1
2. Preliminaries	5
I. Monochromatic generalised paths	6
3. Preliminaries	10
4. A high-level proof of the main result	11
4.0.1. A rough outline	11
4.0.2. The three ingredients	12
4.0.3. Combining the ingredients	14
5. Monochromatic paths covering a κ-star-linked set	19
5.0.1. Finding a monochromatic path covering a κ -star-linked set	19
5.0.2. Finding a monochromatic path covering an improved κ -star-linked set	20
6. Constructing robust paths	22
6.0.1. Constructing countable robust paths	22
6.0.2. Constructing uncountable robust paths	24
7. Constructing large monochromatic paths	27
II. Ends of digraphs	31
8. Basic theory	35
8.1. Preliminaries	39
8.2. Necklace Lemma	42

8.3. Directions	48
8.4. Limit edges and edge-directions	55
9. The topological point of view	59
9.1. Tools and terminology	63
9.2. A topology for digraphs	67
9.3. The space $ D $ as an inverse limit	72
9.4. Applications	79
10. Normal arborescences	86
10.1. Tools and terminology	91
10.2. Normal arborescences	92
10.3. Arborescences are end-faithful	96
10.4. Arborescences reflect the horizon	97
10.5. Existence of arborescences	103
III. Stars and combs	106
11. Arbitrary stars and comb	110
11.1. Tools and terminology	113
11.1.1. Star-Comb Lemma	114
11.1.2. Separations	116
11.1.3. Ends of graphs	116
11.1.4. Critical vertex sets	118
11.1.5. Normal trees	119
11.1.6. Containing vertex sets cofinally	122
11.1.7. Tree-decompositions and S -trees	123
11.1.8. Tree-decompositions and S -trees displaying sets of ends	124
11.2. Combs	126
11.2.1. Normal trees	126
11.2.2. Tree-decompositions	128
11.2.3. Critical vertex sets	130
11.2.4. Rank	131
11.2.5. Combining the duality theorems	134

11.3. Stars	143
11.3.1. Normal trees	143
11.3.2. Tree-decompositions	143
12. Dominating stars and dominated combs	148
12.1. Normal trees	151
12.2. Tree-decompositions	159
12.2.1. A duality theorem related to combs	160
12.2.2. A duality theorem related to stars	163
12.2.3. A duality theorem related to stars and combs	170
13. Undominated combs	179
13.1. Rayless trees	184
13.1.1. Proof for normally spanned graphs	186
13.1.2. Deducing our duality theorem in terms of rayless trees	189
13.2. Spanning trees reflecting the undominated ends	195
13.3. Duality theorems for undominated combs	203
14. Undominating stars	206
14.1. Tough subgraphs	209
14.2. Star-decompositions	216
14.3. Overview of all duality results	225
15. End-faithful spanning trees	229
15.1. Ranking T_κ -free graphs	232
15.2. Normally traceable graphs	234
15.3. End-faithful spanning trees	237
15.4. Trees reflecting the undominated ends	240
Appendix	243
English summary	243
Deutsche Zusammenfassung	247
Publications related to this dissertation	251
Declaration on my contributions	252
Acknowledgement	254

Eidesstattliche Versicherung	260
--	-----

1. Introduction

This dissertation is a contribution to the research on fundamental substructures of infinite directed or undirected graphs. The substructure that we focus on are of three different types and each type is the topic of one of the three parts of this dissertation.

First, we consider **paths** in infinite undirected graphs and confirm a Ramsey-type conjecture of Soukup: Every r -edge-coloured complete bipartite graph with bipartition classes of the same infinite cardinality admits a partition of its vertex set into $2r - 1$ monochromatic generalised paths.

In the second part, we propose and investigate a notion of **ends** of digraphs, for which we then develop an end space theory. While ends of undirected graphs are one of the most important concepts of infinite undirected graph theory, a similarly useful notion and theory of ends of digraphs has never been found before.

In the third part, we determine structures whose existence is complementary to the existence of two substructures that are particularly fundamental to the study of connectedness in infinite graphs: **stars and combs**. Our theorems are phrased in terms of tree-decompositions, normal spanning trees, rayless trees, ranks of rayless graphs and tangle-distinguishing separators.

I. Monochromatic generalised paths

Erdős proved (unpublished [61]) that the vertex set of every 2-edge-coloured complete graph of countably infinite order, can be partitioned into monochromatic paths of different colours, where ‘path’ means either a finite path or a one-way infinite path. Rado subsequently extended Erdős’ result to any finite number of colours [61, Theorem 2]. In the same paper, Rado then asked whether a similar result holds for all infinite complete graphs and a notion of *generalised path* that he proposed.

Soukup [68] answered Rado’s question in the affirmative and conjectured that

a similar result should hold for complete bipartite graphs: Every r -edge-coloured complete bipartite graph with bipartition classes of the same infinite cardinality admits a partition of its vertex set into $2r - 1$ monochromatic generalised paths.

In the first part of this dissertation we answer Soukup's conjecture in the affirmative. In fact, our discussion will also lead to a new, conceptually simpler closing argument for Soukup's proof of Rado's conjecture.

II. Ends of digraphs

Ends of undirected graphs are the single most important concept in infinite graph theory. They can be thought of as points at infinity to which its rays converge. Formally, an *end* of an undirected graph G is an equivalence class of its rays, where two rays are equivalent if for every finite vertex set $X \subseteq V(G)$ they have a subray in the same component of $G - X$ [20].

There is a whole branch of graph theory that is based on ends: *topological infinite graph theory* studies the topological space $|G|$ formed by an undirected graph G together with its ends. Many statements about finite undirected graphs that do not generalise verbatim to arbitrary infinite graphs extend to the space $|G|$. Examples include Nash-William's tree-packing theorem [18], Fleischner's Hamiltonicity theorem [35] and Whitney's planarity criterion [2]. In the formulation of these theorems, topological arcs and circles take the role of paths and cycles, respectively.

For directed graphs, a similarly useful notion and theory of ends has never been found. There have been a few attempts, most notably by Zuther [72], but not with very encouraging results. In this part we propose a new notion of ends of digraphs and develop a corresponding theory of their end spaces.

As for undirected graphs, the *ends* of a digraph are points at infinity to which its rays converge. Furthermore, we extend to digraphs the notion of *directions* of an undirected graph, a tangle-like description of its ends: we provide a natural one-to-one correspondence between the 'directions' of a digraph and its ends and limit edges.

Unlike for undirected graphs, some ends of digraphs are joined by *limit edges*. We introduce a topological space $|D|$ formed by a digraph D together with its ends and limit edges. This makes it possible to extend to the space $|D|$ statements

about finite digraphs that do not generalise verbatim to infinite digraphs. Furthermore, we introduce a concept of depth-first search trees in infinite digraphs, which we call *normal spanning arborescences*. We show that normal spanning arborescences capture the structure of the set of ends of the digraphs they span, both combinatorially and topologically.

III. Stars and combs

The star-comb lemma is a standard tool in infinite graph theory. Recall that a *comb* is the union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R . The last vertices of those paths are the *teeth* of this comb. Given a vertex set U , a *comb attached to U* is a comb with all its teeth in U , and a *star attached to U* is a subdivided infinite star with all its leaves in U . Then the set of teeth is the *attachment set* of the comb, and the set of leaves is the *attachment set* of the star.

Star-comb lemma. *Let U be an infinite set of vertices in a connected graph G . Then G contains either a comb attached to U or a star attached to U .*

The star-comb lemma is not primarily about the existence of one subgraph or another. Rather, it tells us something about the nature of connectedness in infinite graphs: that the way in which they link up their infinite sets of vertices can take two fundamentally different forms, a star and a comb.

Call two properties of infinite graphs *dual*, or *complementary*, in a class of infinite graphs if they partition that class. The existence of stars or combs attached to a given set U is not complementary (in the class of all infinite connected graphs containing U): an infinite complete graph, for example, contains both.

In the third part of this dissertation, we determine structures that are complementary to stars, and structures that are complementary to combs (always with respect to a fixed set U).

As stars and combs can interact with each other, this is not the end of the story. Stars and combs can be combined, positively as well as negatively. For example, a given set U might be connected in G by both a star and a comb, even with infinitely intersecting sets of leaves and teeth. We shall find complementary

structures to the existence of these substructures (again, with respect to some fixed set U).

Just like the original star-comb lemma, our results can be applied as structural tools in other contexts. We dedicate a whole chapter to one of these applications: In the last chapter of this part we make progress to a largely open problem raised by Halin, who asked for a characterisation of the class of graphs with an end-faithful spanning tree [37]. A well-studied subclass is formed by the graphs with a normal spanning tree. We determine a larger subclass, the class of normally traceable graphs, which consists of the connected graphs with a rayless tree-decomposition into normally spanned parts.

2. Preliminaries

For graph theoretic notation we follow the text book *Graph Theory* [20] by Diestel. All the graphs in Part I and in Part III will be undirected. The graphs in Part II will usually be directed, in which case we speak of *digraphs*. We usually consider digraphs without multi-edges and without loops, but possibly with inversely directed edges between distinct vertices. For a digraph D , we write $V(D)$ for the vertex set of D and we write $E(D)$ for the edge set of D . We write edges as ordered pairs (v, w) of vertices $v, w \in V(D)$, and we usually write (v, w) simply as vw .

Further preliminaries can be found in each part of this dissertation. Preliminaries for Part I can be found in Chapter 3, those for Part II in Chapters 8.1, 9.1 and 10.1 and those for Part III in Chapter 11.1.

Part I.

Monochromatic generalised paths

Throughout this part, the term *colouring* always refers to edge-colourings of graphs with finitely many colours.

In the 1970s, Erdős proved (unpublished) that the vertex set of every 2-coloured complete graph of countably infinite order, i.e., every 2-coloured K_{\aleph_0} , can be partitioned into monochromatic paths of different colours, where ‘path’ means either a finite path or a one-way infinite ray. Rado subsequently extended Erdős result to any finite number of colours [61, Theorem 2].

In the same paper, Rado then asked whether a similar result holds for all infinite complete graphs, even the uncountable ones. Clearly, it is not possible to partition such a graph into finitely many ‘usual’ paths, as graph-theoretic paths and rays are inherently countable. Hence, Rado introduced the following notion of *generalised path*: A generalised path is a graph P together with a well-order \prec on $V(P)$ (called the *path order* on P) satisfying that the set $\{w \in N(v) : w \prec v\}$ of down-neighbours of v is cofinal below v for every vertex $v \in V(P)$, i.e., for every $v' \prec v$ there is a neighbour w of v with $v' \preceq w \prec v$ (cf. Figure 2.0.1).




Figure 2.0.1.: A generalised path.

In particular, every successor element is adjacent to its predecessor in the well-order. Calling such a graph P a ‘generalised path’ is justified by the fact that between any two vertices $v \prec w$ of P there exists a *finite* path from v to w strictly increasing with respect to \prec , see e.g. [51, Observation 5.2]. If the situation is clear, we write P instead of (P, \prec) and treat P as a graph. By $\Lambda(P, \prec) = \Lambda(P)$ we denote the limit elements of the well-order (P, \prec) . When the situation is clear, we sometimes write Λ instead of $\Lambda(P)$. If necessary, the path-order \prec on $V(P)$ will be referred to as \prec_P . If $v, v' \in P$, then we denote by (v, v') and $[v, v']$ the open and closed intervals with respect to \prec , and by $[v, v + \omega)$ the ray of P starting at v compatible with the path order. Note that a one-way infinite ray can be viewed quite naturally as a generalised path of order type ω , and conversely, every generalised path of order type ω contains a spanning one-way infinite ray.

Thus, partitioning a graph into monochromatic generalised paths of order type ω is equivalent to partitioning it into monochromatic rays.

Within this part, the term *path* is used in the extended sense of a generalised path.

Elekes, Soukup, Soukup and Szentmiklòssy [51] have recently answered a special case of Rado’s question for \aleph_1 -sized complete graphs and two colours in the affirmative. Shortly after, Soukup [68] gave a complete answer to Rado’s question for any finite number of colours and complete graphs of arbitrary infinite cardinality.

Theorem 2.1 (Soukup, [68, Theorem 7.1]). *Let r be a positive integer. Every r -edge-coloured complete graph of infinite order can be partitioned into monochromatic generalised paths of different colours.*

In [68, Conjecture 8.1], Soukup conjectures that a similar result holds for complete bipartite graphs, namely that every r -coloured complete bipartite graph with bipartition classes of cardinality $\kappa \geq \aleph_0$ can be partitioned into $2r - 1$ monochromatic generalised paths, and has proven his conjecture in the countable case $\kappa = \aleph_0$ [67, Theorem 2.4.1]. If true, this bound would be best possible in the sense that there are r -colourings of $K_{\kappa, \kappa}$ for which the graph cannot be partitioned into $2r - 2$ monochromatic paths, see [67, Theorem 2.4.1].

We remark that Soukup’s conjecture is inspired by the corresponding conjecture in the finite case, due to Prokovskiy [55, Conjecture 4.5]. In contrast to the infinite case, the finite conjecture is only known for two colours [49, p. 169 (footnote)].

The main result of this part is to prove Soukup’s conjecture for all uncountable cardinalities and any (finite) number of colours.

Theorem 2.2. *Let r be a positive integer. Every r -edge-coloured complete bipartite graph with bipartition classes of the same infinite cardinality can be partitioned into $2r - 1$ monochromatic generalised paths with each colour being used at most twice.*

The first uncountable case of Theorem 2.2, where the bipartition classes have size \aleph_1 , was proved in my Master’s thesis [4]. In this part, we extend these ideas to give a proof for all uncountable cardinalities.

Our proof relies on the methods developed by Soukup in his original paper [68]. However, we re-introduce in this part the new, helpful notion of X -robust paths

from [4]—generalised paths which are resistant against the deletion of vertices from X . After introducing such paths, we will state in Chapter 4 three high-level results relying on this new notion, and then provide a proof of Theorem 2.2 from these auxiliary results. In fact, our discussion will also lead to a new, conceptually simpler closing argument for a proof of Soukup’s Theorem 2.1.

In Chapter 5 and 6, we then provide proofs of the auxiliary results. For the second of these auxiliary results, to be proved in Chapter 5, we need to strengthen a result by Soukup [68, §5] to give the statement that any edge-coloured complete bipartite graph with bipartition classes (A, B) of cardinality $\kappa > \aleph_0$ contains a monochromatic path P of order type κ in colour k (say) covering a large subset $X \subseteq A$ which itself is κ -star-linked in colour k , where it is precisely the κ -*star-linked-property* (to be defined below) which is new. We remark that while our statement is slightly stronger, our proof very much relies on Soukup’s proof [68, §5] and does not give an independent proof of Soukup’s result. A discussion how one obtains the strengthened version of Soukup’s result is provided in Chapter 7.

Finally, in Chapter 6 we prove our third auxiliary result. This part contains a crucial new idea how to directly construct an X -robust path Q of order type $\kappa > \aleph_0$ with $X \in [V(Q)]^\kappa$ from a given generalised path P with the star-linked property as above, using nothing but countable combinatorics and avoiding intricate set theoretical arguments using elementary submodels as employed in [68] and [4].

3. Preliminaries

For a natural number $n \in \mathbb{N}$ we write $[n] = \{1, 2, \dots, n\}$ and if $m \leq n$, we write $[m, n] = \{m, m + 1, \dots, n\}$. Let $G = (V, E)$ be a graph, $r \geq 1$ and $k \in [r]$. An r -edge-colouring (or simply r -colouring) of G is a map $c: E \rightarrow [r]$. A path $P \subseteq G$ is *monochromatic* (in colour k with regard to the colouring c) if P is also a path in the graph induced by the edges of colour k , i.e., if P is a path in $(V, c^{-1}(k))$. More generally, suppose that \mathcal{P} is a graph property. We say that G *has property \mathcal{P} in colour k* if $(V, c^{-1}(k))$ has property \mathcal{P} . For a vertex v of G we write $N(v, k)$ for the neighbourhood of v in $(V, c^{-1}(k))$. As a shorthand, we also write $N(v, \neq k) := N(v) \setminus N(v, k)$ for the neighbourhood of v in all colours but k . Let $A \subseteq V$. The common neighbourhood $\bigcap \{N(v) : v \in A\}$ of vertices in A is written as $N[A]$. The common neighbourhood of A in colour k is written as $N[A, k]$. For a cardinal κ , we say that A is κ -star-linked in B , if $N[F] \cap B$ has cardinality κ for every finite $F \subseteq A$. In the case where $B = V(G)$ we simply say that A is κ -star-linked.

When talking about *partitions of G* we always mean vertex partitions and we allow empty partition classes. If $A, B \subseteq V(G)$ are disjoint sets of vertices, then $G[A, B]$ denotes the bipartite graph on $A \cup B$ given by all the edges between A and B .

We write $[X]^\kappa = \{Y \subseteq X : |Y| = \kappa\}$ and $[X]^{<\kappa} = \{Y \subseteq X : |Y| < \kappa\}$, for a given set X .

4. A high-level proof of the main result

The aim of this chapter is to give an overview of the proof of Theorem 2.2. We shall start with a rough idea, inspired by Soukup's work in [68, Theorem 7.1]. After that, we present three main ingredients for our proof of Theorem 2.2: Theorem 4.0.2, Lemma 4.1 and Lemma 4.2. For the moment, we will skip the latter two and discuss them below in Chapter 5 and Chapter 6. We conclude this chapter with a proof of Theorem 2.2 and a proof of Theorem 2.1—also based on the three lemmas.

4.0.1. A rough outline

First, let us have a look at an important idea in Soukup's proof of Theorem 2.1. In [68, Lemma 4.6], Soukup provides some *conditions* which guarantee the existence of a spanning generalised path in a graph. Let us refer to these conditions by (\dagger) . Let κ be an infinite cardinal and $G = (V, E)$ the complete graph of order κ . Suppose that the edges of G are coloured with $r \geq 1$ many colours. In [68, Claim 7.1.2], Soukup shows that one can find sets $X \subseteq W \subseteq V$ and a colour $k \in [r]$, such that

- (i) $G[W \setminus X']$ satisfies (\dagger) in colour k for every $X' \subseteq X$, and
- (ii) $V \setminus W$ is covered by disjoint monochromatic paths of different colours not equal to k in the graph $G[V \setminus W, X]$.

Once such W, X and k are found, we just have to find $r - 1$ disjoint monochromatic paths of different colours $\neq k$ covering $V \setminus W$ in $G[V \setminus W, X]$ as in (2), let $X' \subseteq X$ be the vertices of X covered by these $r - 1$ paths, and apply (1) to guarantee the existence of a monochromatic path in colour k disjoint from all previous ones and covering the remaining vertices.

Whilst it is difficult to work with the conditions from (\dagger) in the bipartite setting directly, the use of (\dagger) in (1) motivates the following definition:

Definition 4.0.1. Let P be a path and $X \subseteq V(P)$. We say that P is X -robust iff for every $X' \subseteq X$ the graph $P - X'$ admits a well-order for which $P - X'$ is a path of the same order type as P .

Our strategy for the bipartite case can now be summarised as follows. Let κ be an infinite cardinal and $G = (V, E)$ the complete bipartite graph with bipartition classes of cardinality κ , where the edges of G are coloured with $r \geq 1$ many colours. Assume that we find $X \subseteq W \subseteq V$ and a colour $k \in [r]$, such that

- (1') $G[W]$ has a spanning X -robust path in colour k , and
- (2') $V \setminus W$ is covered by $2r - 2$ disjoint monochromatic paths in the graph $G[V \setminus W, X]$ in colours not equal to k with every colour appearing at most twice.

Then it is clear that we can complete a proof of Theorem 2.2 in a similar way as above.

4.0.2. The three ingredients

To prove our main theorem, we shall need the following three ingredients. The first is a special case of Soukup's [68, Thm 6.2].

Theorem 4.0.2. *Let G be an infinite bipartite graph with bipartition classes A and B . Suppose that $|A| \leq |B|$ and that $|B \setminus N(a)| < |B|$ for every vertex $a \in A$. Then for every finite edge colouring of G , there are disjoint monochromatic paths of different colours in G covering A .*

That the above theorem follows from [68, Thm 6.2] can be verified by a similar argument as in [68, p. 271, l. 17-20] which we spell out for the convenience of the reader:

Proof. Let κ be the cardinality of A and μ the cardinality of B . By [68, Thm 6.2] it suffices to show that A is (\mathcal{A}, κ) -centred, i.e. we have to find a set $\mathcal{A} = \{(A_\alpha^i)_{\alpha < \lambda_i} : i \in I\}$ for some finite set I , so that

- (i) $A_\alpha^i \subseteq A_\beta^i$ if $\alpha < \beta < \lambda_i$ and $i \in I$,
- (ii) $A \subseteq \bigcup \{A_\alpha^i : \alpha < \lambda_i\}$ for each $i \in I$, and

(iii)

$$|N_G[\bigcap_{i \in I} A_{\alpha_i}^i]| \geq \kappa$$

for all $(\alpha_i)_{i \in I} \in \prod_{i \in I} \lambda_i$.

We consider three cases. First, assume that $\text{cf}(\mu) > \kappa$. Then

$$N_G[A] = B \setminus \bigcup \{B \setminus N(a) : a \in A\}$$

still has size μ and therefore $\mathcal{A} = \{A\}$ works. Next, assume that $\text{cf}(\mu) = \kappa$. Write A as an ascending union of sets $\bigcup \{A_\alpha^1 : \alpha < \kappa\}$ each of size $< \kappa$ and let $\mathcal{A} = \{(A_\alpha^1)_{\alpha < \kappa}\}$. Then A is (\mathcal{A}, κ) -centred since each A_α^1 has size $< \text{cf}(\mu)$ and $B \setminus N(a)$ has size $< \mu$ for every $a \in A_\alpha^1$ and $\alpha < \kappa$. Finally, assume that $\text{cf}(\mu) < \kappa$. In particular, μ is a limit cardinal, so we may fix an increasing sequence $(\mu_\alpha)_{\alpha < \text{cf}(\mu)}$ of cardinals cofinal in μ . Additionally to the previously chosen sequence $(A_\alpha^1)_{\alpha < \kappa}$ define $A_\alpha^2 := \{a \in A : |B \setminus N(a)| < \mu_\alpha\}$ for $\alpha < \text{cf}(\mu)$. Let $\lambda_1 := \kappa$ and $\lambda_2 := \text{cf}(\mu)$, then $\mathcal{A} := \{(A_\alpha^i)_{\alpha < \lambda_i} : i \in \{1, 2\}\}$ satisfies (1), and since the μ_α 's are cofinal in μ , also (2). Condition (3) is true for \mathcal{A} because for all $(\alpha_1, \alpha_2) \in \lambda_1 \times \lambda_2$, both $A_{\alpha_1}^1 \cap A_{\alpha_2}^2$ and $B \setminus N(a)$ have size less than some cardinal $\gamma < \mu$. \square

The next main lemma, which is a strengthening of a similar result by Soukup [68, §5], helps to find a monochromatic path P which has some desirable additional properties.

Lemma 4.1. *Let κ be an infinite cardinal and G the complete bipartite graph with bipartition classes A, B both of cardinality κ . Suppose that $c: E(G) \rightarrow [r]$ is a colouring of G with $r \geq 1$ many colours. Then there are disjoint sets $A_1, A_2 \in [A]^\kappa$, $B_1, B_2 \in [B]^\kappa$ such that (up to renaming the colours):*

- $G[A_1, B_1]$ has a spanning path P of order type κ in colour 1 all of whose limits are contained in B_1 , and
- $A_1 \sqcup A_2$ is κ -star-linked in B_2 in colour 1. (cf. Figure 4.0.1)

Our final ingredient converts the path P from above into a new path Q that has two additional properties: first, Q will be X -robust for some large X , and secondly, Q will be able to additionally cover certain highly inseparable sets of vertices.




Figure 4.0.1.: The colour 1 is indicated blue in the figure.

Definition 4.0.3 (cf. Diestel, [20, p. 354]). Let G be a graph and κ a cardinal. A set $U \subseteq V(G)$ of vertices is $<\kappa$ -inseparable if distinct vertices $v, w \in U$ cannot be separated by less than κ many vertices, i.e. v and w are contained in the same component of $G - W$ for every $W \in [V(G) \setminus \{v, w\}]^{<\kappa}$.

Lemma 4.2. Let κ be an uncountable cardinal and G a bipartite graph with bipartition classes A, B both of size κ . Suppose there are disjoint sets $A_1, A_2 \in [A]^\kappa$, $B_1, B_2 \in [B]^\kappa$ such that

- $G[A_1, B_1]$ has a spanning path P of order type κ with $\Lambda(P) \subseteq B_1$, and
- $A_1 \sqcup A_2$ is κ -star-linked in B_2 .

Then there is a set $X \in [A_2]^\kappa$ and an X -robust path Q covering $A_1 \sqcup A_2$ with $\Lambda(P) = \Lambda(Q)$. Moreover, if $C \subseteq (A \setminus A_1) \sqcup (B \setminus \Lambda)$ covers A_2 and is $<\kappa$ -inseparable in $G[A \setminus A_1, B \setminus \Lambda]$, then Q can be chosen to cover C .

4.0.3. Combining the ingredients

Our three main ingredients can be applied to yield a proof of Theorem 2.2 as follows:

Theorem 2.2. Let r be a positive integer. Every r -edge-coloured complete bipartite graph with bipartition classes of the same infinite cardinality can be partitioned into $2r - 1$ monochromatic generalised paths with each colour being used at most twice.

Proof of Theorem 2.2. Let κ be an infinite cardinal and let G be the complete bipartite graph with bipartition classes A, B both of cardinality κ . Suppose that $c: E(G) \rightarrow [r]$ is a colouring of G . Since the countable case has been solved in [67, Theorem 2.4.1] already, we may assume that κ is uncountable.

We will construct a partition $\mathcal{A} = \{A_1, \dots, A_4\}$ of A and a partition $\mathcal{B} = \{B_1, \dots, B_4\}$ of B such that, up to renaming the colours,

- (i) $G[A_1, B_1]$ has a spanning path P of order type κ in colour 1 with $\Lambda(P) \subseteq B_1$, and $|A_2| = \kappa$,
- (ii) $A_1 \sqcup A_2$ is κ -star-linked in B_2 in colour 1,
- (iii) $A_2 \sqcup A_3$ is $<\kappa$ -inseparable in $G[A_2 \sqcup A_3, B_2 \sqcup B_3]$ in colour 1, and
- (iv) $A_4 \sqcup B_4$ can be partitioned into $r - 1$ monochromatic paths P_2, \dots, P_r in $G[A_4, B_4]$ with distinct colours in $[2, r]$.




Figure 4.0.2.: The colour 1 is indicated blue in the figure.

Let us first see how to complete the proof with these partitions established:

Let C be the set of vertices with $A \cap C = A_2 \cup A_3$ and where $B \cap C$ consists of those vertices in $B \setminus (\Lambda \cup B_4)$ that send κ many edges in colour 1 to A_2 , and observe that (iii) implies that C is $<\kappa$ -inseparable in $G[A \setminus (A_1 \cup A_4), B \setminus (\Lambda \cup B_4)]$ in colour 1. Hence, by (i) and (ii), we may apply Lemma 4.2 in the subgraph of $G[A \setminus A_4, B \setminus B_4]$ induced by the edge of colour 1 in order to obtain a set $X \in [A_2]^\kappa$ and an X -robust, monochromatic path Q in colour 1 with limits $\Lambda = \Lambda(Q) = \Lambda(P)$, covering $A_1 \cup A_2 \cup A_3 \cup C \cup \Lambda$.

Next, note that since $X \subseteq A_2$, it follows by choice of C that $|X \setminus N(b, \neq 1)| = |X \cap N(b, 1)| < \kappa = |X|$ for every vertex b in $B \setminus (Q \cup B_4)$. Therefore, we may apply Theorem 4.0.2 to the bipartite graph $G[B \setminus (Q \sqcup B_4), X]$ with the edges in

colour $\neq 1$ to obtain disjoint monochromatic paths P_{r+1}, \dots, P_{2r-1} with different colours in $[2, r]$ covering $B \setminus (Q \sqcup B_4)$.

Let P_1 be the path that results by deleting $X' = X \cap (\bigcup_{i=r+1}^{2r-1} P_i)$ from Q , using that Q is X -robust. Together with the paths P_2, \dots, P_r provided by (iv), we have found a partition of G into $2r - 1$ disjoint monochromatic paths P_1, \dots, P_{2r-1} using every colour at most twice.

To complete the proof, it remains to construct the partitions \mathcal{A} and \mathcal{B} .

Claim. There are disjoint sets $A_1, A_2 \in [A]^\kappa$ and $B_1, \tilde{B}_2 \in [B]^\kappa$ such that (up to renaming the colours)

- $G[A_1, B_1]$ has a spanning path P of order type κ in colour 1 all whose limits are contained in B_1 and
- $A_1 \sqcup A_2$ is κ -star-linked in \tilde{B}_2 in colour 1.

Proof. Apply Lemma 4.1 to the graph G and the colouring c . ◇

Claim. There is a partition $\{B_2, \tilde{B}_3\}$ of \tilde{B}_2 such that

- $A_1 \sqcup A_2$ is κ -star-linked in B_2 in colour 1 and
- $G[A_2, \tilde{B}_3]$ has a perfect matching M in colour 1.

Proof. Write $A_1 \sqcup A_2$ as an ascending union of sets $\{A^\alpha : \alpha < \text{cf}(\kappa)\}$ each of size $< \kappa$. (Note that if $\kappa = \lambda^+$, we eventually have $|A^\alpha| = \lambda$ for every $\alpha < \text{cf}(\kappa)$.) Simultaneously define in $\text{cf}(\kappa)$ many steps ascending sets $\{B'_\alpha : \alpha < \text{cf}(\kappa)\}$, $\{B''_\alpha : \alpha < \text{cf}(\kappa)\}$ of vertices, and an increasing sequence of matchings $\{M_\alpha : \alpha < \text{cf}(\kappa)\}$ as follows:

To begin let B'_0, B''_0 and M_0 be the empty set. In step $\alpha > 0$, let us write $B'_{<\alpha} := \bigcup\{B'_\beta : \beta < \alpha\}$, $B''_{<\alpha} := \bigcup\{B''_\beta : \beta < \alpha\}$ and $M_{<\alpha} := \bigcup\{M_\beta : \beta < \alpha\}$. Fix a matching M_α of $A_2 \cap A^\alpha$ extending $M_{<\alpha}$ and avoiding $B''_{<\alpha}$, i.e. so that no vertex from $B''_{<\alpha}$ is incident with an edge in M_α . This is possible because $A_2 \cap A^\alpha$ is κ -star-linked. Next let B'_α consist of the matching partners of $A_2 \cap A^\alpha$ with regard to M_α , i.e. $B'_\alpha = B \cap \bigcup M_\alpha$. Finally, fix a set $B''_\alpha \subseteq \tilde{B}_2 \setminus B'_\alpha$ of size $|A^\alpha|$ extending $B''_{<\alpha}$ and so that $B''_\alpha \setminus B''_{<\alpha}$ contains $|A^\alpha|$ many vertices from $N[F, 1]$ for every finite $F \subseteq A^\alpha$ (possible because A^α is κ -star-linked).

It is straightforward to check that the sets $\tilde{B}_3 := \bigcup\{B'_\alpha : \alpha < \text{cf}(\kappa)\}$, $B_2 := \tilde{B}_2 \setminus \tilde{B}_3$ and $M := \bigcup\{M_\alpha : \alpha < \text{cf}(\kappa)\}$ are as desired. ◇

Let A_3 consist of those vertices in $A \setminus (A_1 \sqcup A_2)$ that send κ many edges in colour 1 to \tilde{B}_3 . Note that since A_2 is κ -star linked in B_2 in colour 1, it follows that $A_2 \sqcup A_3$ is $<\kappa$ -inseparable in $G[A_2 \sqcup A_3, B_2 \sqcup \tilde{B}_3]$ in colour 1.

Claim. There is a partition $\{\hat{B}_3, \tilde{B}_4\}$ of \tilde{B}_3 such that

- $A_2 \sqcup A_3$ is $<\kappa$ -inseparable in $G[A_2 \sqcup A_3, B_2 \sqcup \hat{B}_3]$ in colour 1, and
- \tilde{B}_4 has cardinality κ .

Proof. If A_3 is empty, then just take a balanced partition $\{\hat{B}_3, \tilde{B}_4\}$ of \tilde{B}_3 . Otherwise, fix a sequence $(a_\alpha)_{\alpha < \kappa}$ of vertices in A_3 such that every vertex in A_3 appears κ many times. Then fix a vertex in $N_G(a_\alpha, 1) \cap \tilde{B}_3$ for \hat{B}_3 and another in $N_G(a_\alpha) \cap \tilde{B}_3$ for \tilde{B}_4 for every $\alpha < \kappa$ (all distinct). This can be done recursively in κ many steps using that every vertex in A_3 sends κ many edges in colour 1 to \tilde{B}_3 . It is easy to check that sets \hat{B}_3 and \tilde{B}_4 that arise in this manner fulfil our requirements. \diamond

The last partition class of \mathcal{A} is just $A_4 := A \setminus (A_1 \sqcup A_2 \sqcup A_3)$. Applying Theorem 4.0.2 to the spanning subgraph of $G[A_4, \tilde{B}_4]$ induced by the colours $2, \dots, r$ (and the induced colouring) gives rise to disjoint monochromatic paths P_2, \dots, P_r of different colours in $[2, r]$. Let $B_4 := \bigcup \{B \cap P_i : i \in [2, r]\}$ and $B_3 := B \setminus (B_1 \sqcup B_2 \sqcup B_4)$.

We claim that $\mathcal{A} = \{A_1, \dots, A_4\}$ and $\mathcal{B} = \{B_1, \dots, B_4\}$ are as desired. Indeed, it is clear by construction that they are partitions of A and B respectively. From the first and second claim, it follows that (i) and (ii) is satisfied respectively. By the definition of A_4 and B_4 in the last paragraph, we have (iv). And by the third claim, since $B_3 \supseteq \hat{B}_3$, it follows that (iii) holds. \square

Finally, we demonstrate that our approach for the proof of Theorem 2.2, which itself of course relies in many ways on ideas and results from Soukup's [68], can be used to give a conceptually simple closing argument for a proof of Theorem 2.1:

Theorem 2.1 (Soukup, [68, Theorem 7.1]). *Let r be a positive integer. Every r -edge-coloured complete graph of infinite order can be partitioned into monochromatic generalised paths of different colours.*

Proof. Let κ be an infinite cardinal, G the complete graph on κ and $c: E(G) \rightarrow [r]$ a colouring for some $r \geq 1$. Since the countable case has been solved in

[61, Theorem 2], we may assume that κ is uncountable. Fix a partition $\{A, B\}$ of $V(G)$ such that both partition classes have cardinality κ . Apply Lemma 4.1—to the graph $G[A, B]$ and the colouring induced by c —in order to get disjoint sets $A_1, A_2 \in [A]^\kappa$, $B_1, B_2 \in [B]^\kappa$ and a path P as in the lemma. Let Λ be the set of limits of P and write $A' := A_1 \sqcup A_2$, $B' := V(G) \setminus A'$. Furthermore, let C consist of A' together with all those vertices in $V(G) \setminus (A' \cup \Lambda)$ that send κ many edges in colour 1 to A_2 . Apply Lemma 4.2—to the graph induced by the edges of colour 1 in $G[A', B']$ and the set C —in order to find a set $X \in [A_2]^\kappa$ and an X -robust path Q as in the lemma. Next, apply Theorem 4.0.2—to the graph induced by the edges of colour $\neq 1$ in $G[X, B' \setminus Q]$ and the colouring induced by c —in order to find paths P_2, \dots, P_r of different colours in $[2, r]$. The last path in colour 1 is $Q \setminus \bigcup_i P_i$, which is a path due to the X -robustness of Q . \square

5. Monochromatic paths covering a κ -star-linked set

In this chapter, we prove Lemma 4.1. A partial result of Soukup's [68] will assist us: It implies that an edge-coloured complete bipartite graph with bipartition classes (A, B) both of cardinality $\kappa > \aleph_0$ contains a monochromatic path P of order type κ covering a large $<\kappa$ -inseparable subset of A (cf. [68, Theorem 5.10]). By modifying the proof, we obtain a strengthened version where this $<\kappa$ -inseparable subset is even κ -star-linked, Theorem 5.0.1 below. As the main result of this chapter, we explain how to establish Lemma 4.1 as a consequence of Theorem 5.0.1. The detailed proof of Theorem 5.0.1 we defer until the end of this part.

5.0.1. Finding a monochromatic path covering a κ -star-linked set

First we remind the reader of a few concepts from Soukup's [68]: Let κ be a cardinal. Then $H_{\kappa, \kappa}$ denotes the graph $(\kappa \times \{0\} \cup \kappa \times \{1\}, E)$ where

$$\{(\alpha, i), (\beta, j)\} \in E \text{ iff } i = 1, j = 2 \text{ and } \alpha < \beta < \kappa$$

(cf. [68, p.250, l.10–13]). Furthermore, a graph $G = (V, E)$ is of *type* $H_{\kappa, \kappa}$ if there are (not necessarily disjoint) subsets $A, B \subseteq V$ with $V = A \cup B$, and enumerations $A = \{a_\xi : \xi < \kappa\}$ and $B = \{b_\xi : \xi < \kappa\}$ such that

$$\{a, b\} \in E(G) \text{ if } a = a_\xi, b = b_\zeta \text{ for some } \xi \leq \zeta < \kappa.$$

The vertex set A is called the *main class* of G and B is called the *second class* of G (cf. [68, Definition 5.3]). Informally speaking, a type $H_{\kappa, \kappa}$ graph is just a copy of $H_{\kappa, \kappa}$ where the bipartition classes are allowed to intersect.

Another concept that we need is that of a *concentrated path* [68, Definition 4.1]: Let G be a graph and $A \subseteq V(G)$. A path $P \subseteq G$ is concentrated on A if and only

if

$$N(v) \cap A \cap V(P \upharpoonright [x, v)) \neq \emptyset$$

for all $v \in \Lambda(P)$ and $x \prec_P v$.

Theorem 5.0.1. *If G is an r -edge coloured graph of type $H_{\kappa, \kappa}$ with main class A , then there is a colour $k \in [r]$ and $X \in [A]^\kappa$ which is κ -star-linked in colour k , such that X is covered by a monochromatic path of order type κ in colour k concentrated on X .*

Proof. Theorem 5.0.1 follows from Theorem 7.0.4 which is a strengthening of [68, Theorem 5.10], to be proved in our last Chapter 7 below. \square

5.0.2. Finding a monochromatic path covering an improved κ -star-linked set

We need two more lemmas before we can prove Lemma 4.1.

Lemma 5.0.2. *Let G be a bipartite graph with bipartition classes A, B . A path $P \subseteq G$ is concentrated on A if and only if all limits of P are contained in B . \square*

Lemma 5.0.3 (cf. [68], Lemma 3.4). *Let κ be an infinite cardinal, $G = (V, E)$ a graph and $A, B \subseteq V(G)$. Suppose A is κ -star-linked in B . Moreover, let $c: E(G) \rightarrow [r]$ be a colouring of G with $r \geq 1$ many colours. Then there is a partition $\{A_i: i \in [r]\}$ such that A_i is κ -star-linked in B in colour i for every $i \in [r]$.*

Proof. Take a uniform ultrafilter \mathcal{U} on B with $B \cap N[F] \in \mathcal{U}$ for every finite $F \subseteq A$ and write $A_i = \{v \in A: N(v, i) \in \mathcal{U}\}$. Then for $i \in [r]$ and $F \subseteq A_i$, we have $N[F, i] \cap B \in \mathcal{U}$ and thus $N[F, i] \cap B$ has cardinality κ . \square

We are now ready to provide the proof for Lemma 4.1 which we restate here for convenience of the reader.

Lemma 4.1. *Let κ be an infinite cardinal and G the complete bipartite graph with bipartition classes A, B both of cardinality κ . Suppose that $c: E(G) \rightarrow [r]$ is a colouring of G with $r \geq 1$ many colours. Then there are disjoint sets $A_1, A_2 \in [A]^\kappa$, $B_1, B_2 \in [B]^\kappa$ such that (up to renaming the colours):*

- $G[A_1, B_1]$ has a spanning path P of order type κ in colour 1 all of whose limits are contained in B_1 , and
- $A_1 \sqcup A_2$ is κ -star-linked in B_2 in colour 1. (cf. Figure 4.0.1)

Proof. Fix a set $A' \in [A]^\kappa$ that is κ -star-linked in as many colours as possible and let I be the set of those colours. By Lemma 5.0.3, the set I is non-empty and we may assume that colour 1 is contained in I .

Claim. There are disjoint sets $B'_1, B'_2 \subseteq B$ such that A' is κ -star-linked in B'_1 and B'_2 , in all colours in I .

Proof. Write the set A' as an ascending union of sets $\{A^\alpha : \alpha < \text{cf}(\kappa)\}$ each of size $< \kappa$. Simultaneously define in $\text{cf}(\kappa)$ steps ascending sets $\{B_\alpha^1 : \alpha < \text{cf}(\kappa)\}$ and $\{B_\alpha^2 : \alpha < \text{cf}(\kappa)\}$ such that B_α^1 and B_α^2 are disjoint and $|B_\alpha^1| = |A^\alpha| = |B_\alpha^2|$ for all α as follows.

To begin, let B_0^1 and B_0^2 be the empty set. In step $\alpha > 0$, let us write $B_{<\alpha}^1 := \bigcup\{B_\beta^1 : \beta < \alpha\}$ and $B_{<\alpha}^2 := \bigcup\{B_\beta^2 : \beta < \alpha\}$. Since A^α is κ -star-linked in all colours in I , we first find a set $B_\alpha^1 \subseteq B \setminus B_{<\alpha}^2$ of size $|A^\alpha|$ extending $B_{<\alpha}^1$ and so that $B_\alpha^1 \setminus B_{<\alpha}^1$ contains $|A^\alpha|$ many vertices from $N[F, i]$ for every finite $F \subseteq A^\alpha$ and $i \in I$. In a second step, we find a set $B_\alpha^2 \subseteq B \setminus B_\alpha^1$ of size $|A^\alpha|$ extending $B_{<\alpha}^2$ so that $B_\alpha^2 \setminus B_{<\alpha}^2$ contains $|A^\alpha|$ many vertices from $N[F, i]$ for every finite $F \subseteq A^\alpha$ and $i \in I$.

It is straightforward to check that the sets $B'_1 = \bigcup\{B_\alpha^1 : \alpha < \text{cf}(\kappa)\}$ and $B'_2 = \bigcup\{B_\alpha^2 : \alpha < \text{cf}(\kappa)\}$ are as desired. \diamond

Fix B'_1 and B'_2 as in the above claim and let $\{A'_1, A'_2\}$ be a partition of A' such that both partition classes have cardinality κ . Since $G[A'_1, B'_1]$ is complete bipartite, it is in particular of type $H_{\kappa, \kappa}$, and so we may apply Theorem 5.0.1 to $G[A'_1, B'_1]$ to find a colour $k \in [r]$ and $X \in [A'_1]^\kappa$ which is κ -star-linked in colour k such that X is covered by a monochromatic path P (say) of size κ in colour k concentrated on X .

By the maximality of I we have $k \in I$ and we may assume $k = 1$. Furthermore, by Lemma 5.0.2 we know that all limits of P are contained in B'_1 . Hence, letting $A_1 := A \cap P$, $A_2 := A'_2$, $B_1 := B \cap P$ and $B_2 := B'_2$ completes the proof. \square

6. Constructing robust paths

In this chapter we will prove Lemma 4.2. There are two major steps: First, we show how to find a ray R that is $\{x\}$ -robust for a single vertex x . Second, we will construct the path Q as a concatenation of rays each including a copy of R . The set X for which Q is X -robust will be the set of vertices x in the various copies of R .

6.0.1. Constructing countable robust paths

Consider the one-way infinite ladder on the positive integers shown in Figure 6.0.1. The well-order \leq on the positive integers together with this ladder then forms a (generalised) path R , and it is easy to see that R is $\{2\}$ -robust. Indeed, the graph $R' = R - \{2\}$ has the one-way infinite path $R' = 1436587\dots$ as a spanning subgraph. Note that additionally, the first vertices of R' and of R coincide.




Figure 6.0.1.: The fat edges indicate the path order of the one-way ladder on the positive integers.

As we work in the bipartite setting, it is of importance that generalised paths that we want to install are bipartite. Our ray R is bipartite as shown in Figure 6.0.2.

All countably infinite robust paths we construct will always consist of some finite path Q followed by a copy of R , where we denote this concatenation by $Q \frown R$.




Figure 6.0.2.: The top and bottom vertices in the ladder define the two bipartition classes of the one-way infinite ladder. The path order from Figure 6.0.1 is indicated fat again.

It is easy to see that the path $Q \frown R$ is then $\{x\}$ -robust, where x is the vertex corresponding to the vertex $2 \in V(R)$. The following lemma is our key lemma for constructing paths of that kind:

Lemma 6.0.1. *Let G be a bipartite graph with bipartition classes A, B such that A is countably infinite. Suppose further that A is \aleph_0 -star-linked and $a \in A$ is some fixed vertex. Then for any vertex $x \in A \setminus \{a\}$ there is an x -robust ray R in G starting in the vertex a and covering A . Moreover, there is a path order of $R - x$ with first vertex a .*

Proof. Fix an enumeration $(a_n)_{n \geq 1}$ of A with $a_1 = a$ and $a_2 = x$. For $n \geq 1$ fix distinct vertices $(b_n)_{n \geq 1}$ such that b_n is contained in the common neighbourhood of $\{a_n, a_{n+1}, a_{n+2}\}$ for $n \geq 1$. This can be done since A is \aleph_0 -star-linked.

Let us write $B' = \{b_n : n \geq 1\}$. Then $G[(A \setminus \{a_1\}) \cup B']$ has a copy¹ of the one-way infinite ladder L on ω as a spanning subgraph where b_1 corresponds to the vertex 1 and x corresponds to the vertex 2 of L . Let us write R' for this copy of L and endow R' with the path order induced by the path order \leq on L . By our observations at the beginning of this chapter, the ray $R = a_1 \frown R'$ is $\{x\}$ -robust and starts with a . □

¹The vertices in $A \setminus \{a_1\}$ correspond to the upper vertices in Figure 6.0.2 and the vertices in B' to the bottom vertices. The enumerations of $A \setminus \{a_1\}$ and respectively B' are the enumeration which ‘go from left to the right’.

6.0.2. Constructing uncountable robust paths

Lemma 4.2. *Let κ be an uncountable cardinal and G a bipartite graph with bipartition classes A, B both of size κ . Suppose there are disjoint sets $A_1, A_2 \in [A]^\kappa$, $B_1, B_2 \in [B]^\kappa$ such that*

- $G[A_1, B_1]$ has a spanning path P of order type κ with $\Lambda(P) \subseteq B_1$, and
- $A_1 \sqcup A_2$ is κ -star-linked in B_2 .

Then there is a set $X \in [A_2]^\kappa$ and an X -robust path Q covering $A_1 \sqcup A_2$ with $\Lambda(P) = \Lambda(Q)$. Moreover, if $C \subseteq (A \setminus A_1) \sqcup (B \setminus \Lambda)$ covers A_2 and is $<\kappa$ -inseparable in $G[A \setminus A_1, B \setminus \Lambda]$, then Q can be chosen to cover C .

Proof. Let us write λ_0 for the first vertex on P and let $\{\lambda_\alpha : 1 \leq \alpha < \kappa\}$ be the enumeration of the limits of P along the path order of P , i.e. we have $\lambda_\alpha \prec_P \lambda_\beta$ whenever $1 \leq \alpha < \beta$. Fix an enumeration $\{c_\alpha : \alpha < \kappa\}$ of C , (choose $C = A_2$ if C is not specified). Note that C has indeed cardinality κ as A_2 is included in C . We construct a sequence of pairwise disjoint paths $\mathcal{S} = (S_\alpha)_{\alpha < \kappa}$ and a sequence of distinct vertices $(x_\alpha)_{\alpha < \kappa}$ from A_2 satisfying the following:

- (i) S_α has order type ω ,
- (ii) S_α has first vertex λ_α and doesn't meet any other limits of P ,
- (iii) S_α is x_α -robust and there is a path order $\prec_{S_\alpha - x_\alpha}$ of $S_\alpha - x_\alpha$ that has first vertex λ_α ,
- (iv) $S_\alpha \cap A \cap P = A \cap P \upharpoonright [\lambda_\alpha, \lambda_\alpha + \omega)$ and
- (v) $\bigcup_{\beta \leq \alpha} S_\beta$ contains c_α .

Once \mathcal{S} is defined we obtain Q as the concatenation $Q = S_0 \frown S_1 \frown S_2 \frown \dots$ (formally, the path order is given by the lexicographic order on $\bigcup_{\alpha < \kappa} \{\alpha\} \times S_\alpha$). Indeed, conditions (1) and (2) guarantee that the limits of Q and the limits of P coincide, and so it follows from (4) that Q is indeed a generalised path. By condition (5), the path Q covers C . Finally, put $X = \{x_\alpha : \alpha < \kappa\}$.

Claim. The path Q is X -robust.

Proof. In order to see that Q is X -robust, let $X' \subseteq X$ be arbitrary. Let S'_α be the path $(S_\alpha - x_\alpha, \prec_{S_\alpha - x_\alpha})$, if S_α meets X' and $S'_\alpha = S_\alpha$ otherwise. Then $Q' = S'_0 \frown S'_1 \frown S'_2 \frown \dots$ is a path of order type κ covering $Q - X'$. \diamond

It remains to define $\mathcal{S} = (S_\alpha)_{\alpha < \kappa}$. Suppose that S_α has already been defined for $\alpha < \beta$. Write $\Sigma_\beta := \bigcup_{\alpha < \beta} S_\alpha$, a set of cardinality $< \kappa$. We first find a finite path T that

- starts in λ_β ,
- ends in a vertex $a \in A_2$ (say),
- contains c_β (unless $c_\beta \in \Sigma_\beta$ already),
- avoids Σ_β and meets P only in $P \upharpoonright [\lambda_\beta, \lambda_\beta + \omega)$.

Claim. A path T as above exists.

Proof. Let T_1 be the path of (edge-)length 1 or 0, that starts in λ_β and is followed by the successor of λ_β on P if λ_β is not already contained² in A_1 .

Let w_1 denote the last vertex on T_1 , and note that $w_1 \notin \Sigma_\beta$ by (4). Since $A_1 \sqcup A_2$ is κ -star-linked in B_2 , we may choose any $w_2 \in A_2 \setminus \Sigma_\beta$ and find a vertex $w_3 \in (B_2 \cap N[\{w_1, w_2\}]) \setminus \Sigma_\beta$ so that $T_2 := w_1 w_3 w_2$ forms a path of (edge-)length two.

If c_β is not yet covered by Σ_β , as C is $< \kappa$ -inseparable in $G[A \setminus A_1, B \setminus \Lambda]$ and $\Sigma_\beta \cup V(T_1) \cup V(T_2)$ has size $< \kappa$, we find a finite path T_3 that contains c_β , starts in the vertex w_2 and ends in a vertex $a \in A_2$ and avoids

$$A_1 \cup \Lambda(P) \cup \Sigma_\beta \cup V(T_1) \cup V(T_2).$$

Otherwise, we put $T_3 = \emptyset$ and $a = w_2$. Then T can be chosen as $T_1 \hat{\ } T_2 \hat{\ } T_3$. \diamond

To complete the proof, we now find a path R of order type ω such that it

- starts in the vertex a and avoids T everywhere else,
- is $\{x_\beta\}$ -robust for a vertex $x_\beta \in A_2 \setminus \{a\}$ and there is a path order of $R - x_\beta$ that starts with a ,
- avoids Σ_β and meets P precisely in $(A \cap P \upharpoonright [\lambda_\beta, \lambda_\beta + \omega)) \setminus V(T)$.

Claim. A path R as above exists.

²Since all limits of P are contained in B_1 , we have $\lambda_\beta \notin A_1$ as soon as $\beta \geq 1$. In the case where $\beta = 0$, we might have $\lambda_\beta \in A_1$.

Proof. Choose $x_\beta \in A_2 \setminus (\Sigma_\beta \cup V(T))$ arbitrary. Apply Lemma 6.0.1 inside the bipartite graph $G[A', B']$ with the vertex a and the vertex $x = x_\beta$, where

$$A' = \{a, x_\beta\} \cup ((A \cap P \upharpoonright [\lambda_\beta, \lambda_\beta + \omega)) \setminus V(T))$$

is countable, and $B' = B_2 \setminus (\Sigma_\beta \cup V(T))$. ◇

Letting $S_\beta = T \frown R$ completes the construction of S_β and thereby our proof is complete. □

7. Constructing large monochromatic paths

The following lemma of Soukup is the main tool of constructing large generalised paths. To state the lemma, we need the following definition.

Definition 7.0.1 ([68, Definition 4.4]). Suppose that $G = (V, E)$ is a graph and $A \subseteq V$. We say that A satisfies \spadesuit_κ if for each $\lambda < \kappa$ there are κ many disjoint paths concentrated on A each of order type λ .

Moreover, if we have a fixed edge-colouring $c: E \rightarrow [r]$ in mind, we write $\spadesuit_{\kappa,i}$ for “ \spadesuit_κ in colour i ”.

Lemma 7.0.2 ([68, Lemma 4.6]). *Suppose that $G = (V, E)$ is a graph, κ an infinite cardinal, and $A \in [V]^\kappa$. If*

- (i) *A is $<\kappa$ -inseparable and if κ is uncountable, then*
- (ii) *A satisfies \spadesuit_κ , and*
- (iii) *there is a nice sequence of elementary submodels $(M_\alpha)_{\alpha < cf(\kappa)}$ for $\{A, G\}$ covering A so that there is $x_\beta \in A \setminus M_\beta$, $y_\beta \in V \setminus M_\beta$ with $x_\beta y_\beta \in E$ and*

$$|N_G(y_\beta) \cap A \cap M_\beta \setminus M_\alpha| \geq \omega$$

for all $\alpha < \beta < cf(\kappa)$,

then A is covered by a generalised path P concentrated on A .

Recall that Soukup considers for fixed κ the following statements:

- (IH) $_{\kappa,r}$ Let H be a graph of type $H_{\kappa,\kappa}$ with main class A and second class B . Then for every r -colouring of H , there is a colour k and an $X \in [A]^\kappa$ so that X satisfies all three conditions of Lemma 7.0.2 in colour k .
- (IH) $_\kappa$ The statement (IH) $_{\kappa,r}$ holds for every $r \geq 1$.

Soukup's main result is then

Theorem 7.0.3 ([68, Theorem 5.10]). *(IH) $_{\kappa}$ holds for all κ . In particular, if G is a graph of type $H_{\kappa,\kappa}$ with a finite-edge colouring, then we can find a monochromatic path of size κ concentrated on the main class of G .*

We now strengthen Soukup's results as follows, and consider the statements:

(IH) $'_{\kappa,r}$ The statement (IH) $_{\kappa,r}$ with the additional requirement that X is also κ -star-linked in colour k .

(IH) $'_{\kappa}$ The statement (IH) $'_{\kappa,r}$ holds for every $r \geq 1$.

The corresponding version of theorem [68, Theorem 5.10] then reads:

Theorem 7.0.4. *(IH) $'_{\kappa}$ holds for all κ . In particular, if G is a graph of type $H_{\kappa,\kappa}$ with main class A with an r -edge colouring, then there is a colour $k \in [r]$ and $X \in [A]^{\kappa}$ which is κ -star-linked in colour k , such that X is covered by a monochromatic path of size κ in colour k concentrated on X .*

The proof of Theorem 7.0.4 relies on the following lemma.

Lemma 7.0.5 (cf. [68, Lemma 5.9]). *Let κ be an infinite cardinal. Suppose that c is an r -edge colouring of a graph $G = (V, E)$ of type $H_{\kappa,\kappa}$ with main class A and second class B . Let $I \subsetneq [r]$, $X \in [A]^{\kappa}$ and suppose that X is κ -star linked in all colours $i \in I$. If (IH) $_{\lambda}$ holds for all $\lambda < \kappa$ then either*

- (a) *there is an $i \in I$ such that X satisfies $\spadesuit_{\kappa,i}$, or*
- (b) *there is $\tilde{X} \in [X]^{<\kappa}$ and a partition $\{X_j : j \in [r] \setminus I\}$ of $X \setminus \tilde{X}$ such that X_j is κ -star-linked in B in colour j for each $j \in [r] \setminus I$.*

Proof of Lemma 7.0.5. Follow the proof of [68, Lemma 5.9, p. 261] and in the last line apply the following Claim A instead of [68, Claim 5.9.3].

Claim A (cf. [68, Claim 5.9.3]). *Suppose that c is an r -edge colouring of a graph $G = (V, E)$ of type $H_{\kappa,\kappa}$ with main class A and second class B . Let $I \subseteq r$ and $X \subseteq A$. If for each finite subset $F \subseteq A$ we have*

$$|B \setminus \bigcup \{N(x, i) : x \in F, i \in I\}| = \kappa,$$

then there is a partition $\{X_j : j \in r \setminus I\}$ of X such that X_j is κ -star-linked in B in colour j for each $j \in [r] \setminus I$.

Proof. Take a uniform ultrafilter U on B so that $B \setminus \bigcup \{N(x, i) : x \in F, i \in I\} \in U$ for all finite subsets $F \subseteq A$. Define $X_j = \{x \in X \setminus \tilde{X} : N(x, j) \in U\}$ for each colour j and note that $\{X_j : j \in [r] \setminus I\}$ partitions X . Since ultrafilters are closed under finite intersections, it follows that

$$N[F, j] \in U$$

for all finite subsets $F \subseteq X_j$ and $j \in [r] \setminus I$, and since the filter U is uniform, we have $|N[F, j]| = \kappa$ and therefore that X_j is κ -star-linked in B for each such j . \diamond

Indeed, by applying Claim A to the set $X \setminus (X^* \cup \tilde{A})$ (defined in Soukup's proof), we readily obtain the stronger conclusion that the X_j are not only $<\kappa$ -inseparable, but even κ -star-linked. \square

Proof of Theorem 7.0.4. We prove $(\text{IH})'_{\kappa, r}$ by induction on κ and r .

Note that $(\text{IH})'_\omega$ holds by [68, Lemma 3.4], so we may suppose that κ is uncountable. Also, $(\text{IH})'_{\kappa, 1}$ holds: From [68, Observation 5.7], we know that for any graph G of type $H_{\kappa, \kappa}$, the main class of G satisfies all conditions of Lemma 7.0.2 (and so $(\text{IH})_{\kappa, 1}$ holds). However, it is clear that the main class A is automatically κ -star-linked in G , and hence we have $(\text{IH})'_{\kappa, 1}$.

Now fix an r -edge colouring with $r > 1$ of a graph G of type $H_{\kappa, \kappa}$ with main class A and second class B . As in the six line argument in Soukup's proof of [68, Theorem 5.10] (Theorem 7.0.3 above), we may assume by the induction assumption $(\text{IH})'_{\kappa, r-1}$ that every $X \in [A]^\kappa$ satisfies condition (3) in Lemma 7.0.2 for each colour in $[r]$.

Next, Soukup fixes a maximal $I \subseteq [r]$ with the property that there is a set $X \in [A]^\kappa$ such that X is $<\kappa$ -inseparable in all colours $i \in I$ and he fixes such I and X . Instead, we now fix I maximal with the property that there is a set $X \in [A]^\kappa$ such that X is κ -star-linked in all colours $i \in I$. Then fix such I and X . Note that $I \neq \emptyset$ by Lemma 7.0.5.

Claim B (cf. [68, Claim 5.10.1]). *There is $k \in I$ such that $\spadesuit_{\kappa, k}$ holds for X .*

Proof. Suppose that X fails $\spadesuit_{\kappa, i}$ for all $i \in I$. If $I \subsetneq [r]$, then apply Lemma 7.0.5 in G to the set X and the set of colours I . As X fails $\spadesuit_{\kappa, i}$ for all $i \in I$, condition (b) of Lemma 7.0.5 must hold (note that by induction assumption, $(\text{IH})'_\lambda$ and hence $(\text{IH})_\lambda$ hold for all $\lambda < \kappa$, so we may apply Lemma 7.0.5): However, this

means there is a colour $j \in [r] \setminus I$ and a set $X_j \in [X]^\kappa$ such that X_j is κ -star-linked in colour j as well. But the fact that X_j is then κ -star-linked in all colours $i \in I \cup \{j\}$ contradicts the maximality of I .

Therefore, $I = [r]$ must hold. From this, however, we may obtain a contradiction precisely as in the second half of the proof of [68, Claim 5.10.1]. \diamond

Hence, the “in particular” part of the theorem, and hence Theorem 5.0.1 follows by applying Lemma 7.0.2 to the set X provided by $(\text{IH})'_\kappa$. The proof is complete. \square

Part II.

Ends of digraphs

Ends of graphs are one of the most important concepts in infinite graph theory. They can be thought of as points at infinity to which its rays converge. Formally, an *end* of a graph G is an equivalence class of its rays, where two rays are equivalent if for every finite vertex set $X \subseteq V(G)$ they have a tail in the same component of $G - X$. For example, infinite complete graphs or grids have one end, while the binary tree has continuum many ends, one for every rooted ray [20]. The concept of ends was introduced in 1931 by Freudenthal [33], who defined ends for certain topological spaces. In 1964, Halin [37] introduced ends for infinite undirected graphs, taking his cue directly from Carathéodory's *Primenden* of regions in the complex plane [14].

There is a natural topology on the set of ends of a graph G , which makes it into the *end space* $\Omega(G)$. Polat [58, 59] studied the topological properties of this space. Diestel and Kühn [26] extended this topological space to the space $|G|$ formed by the graph G together with its ends. Many well known theorems of finite graph theory extend to this space $|G|$, while they do not generalise verbatim to infinite graphs. Examples include Nash-William's tree-packing theorem [18], Fleischner's Hamiltonicity theorem [35], and Whitney's planarity criterion [2]. In the formulation of these theorems, topological arcs and circles take the role of paths and cycles, respectively.

For directed graphs, a similarly useful notion and theory of ends has never been found. There have been a few attempts, most notably by Zuther [72], but not with very encouraging results. In this part we propose a new notion of ends of digraphs and develop a corresponding theory of their end spaces. Let us give a brief overview of the part.

In the first chapter of this part we lay the foundation for the whole part by extending to digraphs a number of techniques that are important in the study of ends of graphs.

As our main result we show that the one-to-one correspondence between the directions and the ends of a graph has an analogue for digraphs. A *direction* of a graph G is a map f , with domain the set of finite vertex sets X of G , that maps every such X to a component of $G - X$ so that $f(X) \supseteq f(Y)$ whenever $X \subseteq Y$. Every end ω of G naturally defines a direction f_ω which maps every finite vertex set $X \subseteq V(G)$ to the unique component of $G - X$ in which every ray representing ω has a tail. It is straightforward to show that f_ω is indeed a direction of G .

Conversely, Diestel and Kühn [25] proved that for every direction f of G there is a (unique) end ω of G that defines f in that $f_\omega = f$.

For a digraph D we will adapt the definition of a direction by first replacing every occurrence of the word ‘component’ with ‘strong component’. These directions of D will correspond bijectively to the ends of D . However, as there may be edges between distinct strong components of D , there will be another type of direction: one that maps finite vertex sets $X \subseteq V(D)$ to the set of edges between two distinct strong components of $D - X$. These latter directions of digraphs will correspond bijectively to its *limit edges*—additional edges between distinct ends, or between ends and vertices, of a digraph.

In the course of proving that the ends and limit edges of a digraph correspond to its two types of directions in this way, we extend to digraphs a number of fundamental tools and techniques for ends of graphs, such as the star-comb lemma [20, Lemma 8.2.2] and Schmidt’s ranking of rayless graphs [64].

In the second chapter of this part we will define a topology on the space $|D|$ formed by the digraph D together with its ends and limit edges. To illustrate the typical use of this space $|D|$, we extend to it two statements about finite digraphs that do not generalise verbatim to infinite digraphs. The first statement is the characterisation of Eulerian digraphs by the condition that the in-degree of every vertex equals its out-degree. The second statement is the characterisation of strongly connected digraphs by the existence of a closed Hamilton walk, see [1]. In the course of our proofs we extend to the space $|D|$ a number of techniques that have become standard in proofs of statements about $|G|$, such as the jumping arc lemma or the fact that $|G|$ is an inverse limit of finite contraction minors of G .

In the third chapter of this part we consider normal spanning trees, one of the most important structural tools in infinite graph theory. Here a rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T . (A T -path in G is a non-trivial path that meets T exactly in its endvertices.) In finite graphs, normal spanning trees are precisely the depth-first search trees [20].

As a directed analogue of normal spanning trees we introduce and study *normal spanning arborescences* of digraphs. These are generalisations of depth-first search trees to infinite digraphs, which promise to be as powerful for a structural analysis of digraphs as normal spanning trees are for graphs. We show that normal

spanning arborescences capture the structure of the set of ends of the digraphs they span, both combinatorially and topologically. Furthermore, we provide a Jung-type [43] criterion for the existence of normal spanning arborescences in digraphs.

8. Basic theory

In order to state the main results of this first chapter of this part more formally, we need a few definitions.

A *directed ray* is an infinite directed path that has a first vertex (but no last vertex). The directed subrays of a directed ray are its *tails*. For the sake of readability we shall omit the word ‘directed’ in ‘directed path’ and ‘directed ray’ if there is no danger of confusion. We call a ray in a digraph D *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. We call two solid rays in a digraph D *equivalent* if for every finite vertex set $X \subseteq V(D)$ they have a tail in the same strong component of $D - X$. The classes of this equivalence relation are the *ends* of D . The set of ends of D is denoted by $\Omega(D)$. In Chapter 9 we will equip $\Omega(D)$ with a topology and we will call $\Omega(D)$ together with this topology *the end space of D* . Note that two solid rays R and R' in D represent the same end if and only if D contains infinitely many disjoint paths from R to R' and infinitely many disjoint paths from R' to R .

For example, the digraph D in Figure 8.0.1 has two ends, which are shown as small dots on the right. Both the upper ray R and the lower ray R' are solid in D because the vertex set of any tail of R or R' is strongly connected in D . Deleting finitely many vertices of D always results in precisely two infinite strong components (and finitely many finite strong components) spanned by the vertex sets of tails of R or R' .




Figure 8.0.1.: A digraph with two ends (depicted as small dots) linked by a limit edge (depicted as a dashed line). Every undirected edge in the figure represents a pair of inversely directed edges.

Similarly to ends of graphs, the ends of a digraph can be thought of as points at infinity to which its rays converge. We will make this formal in Chapter 9, but roughly one can think of this as follows. For a finite vertex set $X \subseteq V(D)$ and an end $\omega \in \Omega(D)$ we write $C(X, \omega)$ for the unique strong component of $D - X$ that contains a tail of every ray that represents ω ; the end ω is then said to *live* in that strong component. In our topological space the strong components of the form $C(X, \omega)$ together with all the ends that live in them will essentially form the basic open neighbourhoods around ω .

Given an infinite vertex set $U \subseteq V(D)$, we say that an end ω is *in the closure* of U in D if $C(X, \omega)$ meets U for every finite vertex set $X \subseteq V(D)$. (It will turn out that an end is in the closure of U in D if and only if it is in the topological closure of U .)

For undirected graphs G one often needs to know whether an end ω is in the closure of a given vertex set U , i.e., whether U meets $C(X, \omega)$ for every finite vertex set $X \subseteq V(G)$. This is equivalent to G containing a comb with all its teeth in U . Recall that a *comb* is the union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R . The last vertices of those paths are the *teeth* of this comb. A standard tool in this context is the star-comb lemma [20, Lemma 8.2.2] which states that a connected graph contains for a given set U of vertices either a comb with all its teeth in U or an infinite subdivided star with all its leaves in U . In this chapter we will prove a directed version of the star-comb lemma.

Call two statements A and B *complementary* if the negation of A is equivalent to B . For a graph G , the statement that G has an end in the closure of $U \subseteq V(G)$ is complementary to the statement that G has a U -rank, see [5]. For $U = V(G)$, the U -rank is known as Schmidt's ranking of rayless graphs [20, 64]. It is a standard technique to prove statements about rayless graphs by transfinite induction on Schmidt's rank. For example Bruhn, Diestel, Georgakopoulos, and Sprüssel [3] employed this technique to prove the unfriendly partition conjecture for countable rayless graphs.

The directed analogue of a comb with all its teeth in U will be a 'necklace' attached to U . The *symmetric ray* is the digraph obtained from an undirected ray by replacing each of its edges by its two orientations as separate directed edges. A *necklace* is an inflated symmetric ray with finite branch sets. (An inflated H

is obtained from a digraph H by subdividing some edges of H finitely often and then replacing the ‘old’ vertices by strongly connected digraphs. The *branch sets* of the inflated H are these strongly connected digraphs. See Section 8.1 for the formal definition of inflated, and of branch sets.) Figure 8.0.2 shows an example of a necklace. Given a set U of vertices in a digraph D , a necklace $N \subseteq D$ is




Figure 8.0.2.: A necklace up to the fourth branch set. Every undirected edge in the figure represents a pair of inversely directed edges.

attached to U if infinitely many of the branch sets of N contain a vertex from U . We will see that the statement that D has an end in the closure of U is equivalent to the statement that D contains a necklace attached to U as a subdigraph.

We extend Schmidt’s result that a graph is rayless if and only if it has a rank. See Section 8.2 for the definition of ‘ U -rank’ in digraphs.

Lemma 8.1 (Necklace Lemma). *Let D be any digraph and U any set of vertices in D . Then the following statements are complementary:*

- (i) D has a necklace attached to U ;
- (ii) D has a U -rank.

Let us now define a directed analogue of the directions of undirected infinite graphs. Consider any digraph D , and write $\mathcal{X}(D)$ for the set of finite vertex sets in D . A (*vertex-*)*direction* of D is a map f with domain $\mathcal{X}(D)$ that sends every $X \in \mathcal{X}(D)$ to a strong component of $D - X$ so that $f(X) \supseteq f(Y)$ whenever $X \subseteq Y$. Ends of digraphs define vertex-directions in the same way as ends of graphs do; for every end $\omega \in \Omega(D)$ we write f_ω for the vertex-direction that maps every $X \in \mathcal{X}(D)$ to the strong component $C(X, \omega)$ of $D - X$. We will show that this correspondence between ends and vertex-directions is bijective:

Theorem 8.2. *Let D be any infinite digraph. The map $\omega \mapsto f_\omega$ with domain $\Omega(D)$ is a bijection between the ends and the vertex-directions of D .*

While most of the concepts that we investigate have undirected counterparts, there is one important exception: limit edges. If ω and η are distinct ends of




Figure 8.0.3.: A digraph with one end (depicted as a small dot) and a limit edge (depicted as a dashed line) from the lower vertex to the end. Every undirected edge in the figure represents a pair of inversely directed edges.

a digraph, there exists a finite vertex set $X \in \mathcal{X}(D)$ such that ω and η live in distinct strong components of $D - X$. Let us say that such a vertex set X *separates* ω and η . For two distinct ends $\omega, \eta \in \Omega(D)$ we call the pair (ω, η) a *limit edge* from ω to η if D has an edge from $C(X, \omega)$ to $C(X, \eta)$ for every finite vertex set X that separates ω and η .

Similarly, for a vertex $v \in V(D)$ and an end $\omega \in \Omega(D)$ we call the pair (v, ω) a *limit edge from v to ω* if D has an edge from v to $C(X, \omega)$ for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. And we call the pair (ω, v) a *limit edge from ω to v* if D has an edge from $C(X, \omega)$ to v for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. We write $\Lambda(D)$ for the set of limit edges of D .

The digraph in Figure 8.0.1 has a limit edge from the lower end to the upper end, and the digraph in Figure 8.0.3 has a limit edge from the lower vertex to the unique end. Let us enumerate from left to right the vertical edges e_0, e_1, \dots of the digraph D in Figure 8.0.1. We may think of the e_n as converging towards the unique limit edge. This will be made precise in Chapter 9.

Every limit edge $\omega\eta$ between two ends naturally defines a map $f_{\omega\eta}$ with domain $\mathcal{X}(D)$ as follows. If $X \in \mathcal{X}(D)$ separates ω and η , then $f_{\omega\eta}$ maps X to the set of edges between $C(X, \omega)$ and $C(X, \eta)$; otherwise $f_{\omega\eta}$ maps X to the strong component of $D - X$ in which both ends live. The map $f_{\omega\eta}$ is consistent in that $f_{\omega\eta}(X) \supseteq f_{\omega\eta}(Y)$ whenever $X \subseteq Y$.¹

This gives rise to a second type of direction of a digraph D , as follows. Given $X \in \mathcal{X}(D)$, a non-empty set of edges is a *bundle of $D - X$* if it is the set of all the

¹Here, as later in this context, we do not distinguish rigorously between a strong component and its set of edges. Thus if Y separates ω and η but $X \subseteq Y$ does not, the expression $f_{\omega\eta}(X) \supseteq f_{\omega\eta}(Y)$ means that the strong component $f_{\omega\eta}(X)$ of $D - X$ contains all the edges from the edge set $f_{\omega\eta}(Y)$.

edges from C to C' , or from v to C , or from C to v , for strong components C and C' of $D - X$ and a vertex $v \in X$. A *direction* of D is a map f with domain $\mathcal{X}(D)$ that maps every $X \in \mathcal{X}(D)$ to a strong component of $D - X$ or to a bundle of $D - X$ so that $f(X) \supseteq f(Y)$ whenever $X \subseteq Y$. We call a direction of D an *edge-direction* of D if there is some $X \in \mathcal{X}(D)$ such that $f(X)$ is a bundle of $D - X$, in other words, if it is not a vertex-direction. Hence f_λ is an edge-direction for limit edges λ between two ends, and for limit edges λ between vertices and ends an edge-direction f_λ can be defined analogously. Our next theorem states that every edge-direction can be described in this way:

Theorem 8.3. *Let D be any infinite digraph. The map $\lambda \mapsto f_\lambda$ with domain $\Lambda(D)$ is a bijection between the limit edges and the edge-directions of D .*

This chapter is organised as follows. In Section 8.1 we provide the basic terminology that we use throughout this part. In Section 8.2 we prove the necklace lemma and discuss some basic properties of ends of digraphs. In Section 8.3 we prove Theorem 8.2. Finally, in Section 8.4 we investigate limit edges and prove Theorem 8.3.

8.1. Preliminaries

For the sake of readability, we sometimes omit curly brackets of singletons, i.e., we write x instead of $\{x\}$ for a set x . Furthermore, we omit the word ‘directed’—for example in ‘directed path’—if there is no danger of confusion.

Throughout this paper D is an infinite digraph without multi-edges and without loops, but which may have inversely directed edges between distinct vertices. For a digraph D , we write $V(D)$ for the vertex set of D , we write $E(D)$ for the edge set of D and $\mathcal{X}(D)$ for the set of finite vertex sets of D . We write edges as ordered pairs (v, w) of vertices $v, w \in V(D)$, and we usually write (v, w) simply as vw . The *reverse* of an edge vw is the edge wv . More generally, the *reverse* of a digraph D is the digraph on $V(D)$ where we replace every edge of D by its reverse, i.e., the reverse of D has the edge set $\{vw \mid wv \in E(D)\}$. A *symmetric path* is a digraph obtained from an undirected path by replacing each of its edges by its two orientations as separate directed edges. Similarly, a *symmetric ray* is a digraph obtained from an undirected ray by replacing each of its edges by its two

orientations as separate directed edges. Hence the reverse of any symmetric path or symmetric ray is a symmetric path or symmetric ray, respectively.

The directed subrays of a ray are its *tails*. Call a ray *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$.

Two solid rays in D are *equivalent*, if they have a tail in the same strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. We call the equivalence classes of this relation the *ends* of D and we write $\Omega(D)$ for the set of ends of D .

Similarly, the reverse subrays of a reverse ray are its *tails*. We call a reverse ray *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. With a slight abuse of notation, we say that a reverse ray R *represents* an end ω if there is a solid ray R' in D that represents ω such that R and R' have a tail in the same strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$.

For a finite vertex set $X \subseteq V(D)$ and a strong component C of $D - X$ an end ω is said to *live in* C if one (equivalent every) solid ray in D that represents ω has a tail in C . We write $C(X, \omega)$ for the strong component of $D - X$ in which ω lives. For two ends ω and η of D a finite set $X \subseteq V(D)$ is said to *separate* ω and η if $C(X, \omega) \neq C(X, \eta)$, i.e., if ω and η live in distinct strong components of $D - X$.

Given sets $A, B \subseteq V(D)$ of vertices a *path from* A *to* B , or A - B path is a path that meets A precisely in its first vertex and B precisely in its last vertex. We say that a vertex v can *reach* a vertex w in D if there is a v - w path in D . A set W of vertices is *strongly connected* in D if every vertex of W can reach every other vertex of W in $D[W]$.

Let H be a fixed digraph. A *subdivision* of H is a digraph that is obtained from H by replacing every edge vw of H by a path P_{vw} with first vertex v and last vertex w so that the paths P_{vw} are internally disjoint and do not meet $V(H) \setminus \{v, w\}$. We call the paths P_{vw} *subdividing* paths. If D is a subdivision of H , then the original vertices of H are the *branch vertices* of D and the new vertices its *subdividing vertices*.

An *inflated* H is any digraph that arises from a subdivision H' of H as follows. Replace every branch vertex v of H' by a strongly connected digraph H_v so that the H_v are disjoint and do not meet any subdividing vertex; here replacing means that we first delete v from H' and then add $V(H_v)$ to the vertex set and $E(H_v)$ to the edge set. Then replace every subdividing path P_{vw} that starts in v and

ends in w by an H_v - H_w path that coincides with P_{vw} on inner vertices. We call the vertex sets $V(H_v)$ the *branch sets* of the inflated H . A *necklace* is an inflated symmetric ray with finite branch sets; the branch sets of a necklace are its *beads*. (See Figure 8.0.2 for an example of a necklace.)

A vertex set $Y \subseteq V(D)$ *separates* A and B in D with $A, B \subseteq V(D)$ if every A - B path meets Y , or if every B - A path meets Y . For two vertices v and w of D we say that $Y \subseteq V(D) \setminus \{v, w\}$ *separates* v and w in D , if it separates $\{v\}$ and $\{w\}$ in D . A *separation* of D is an ordered pair (A, B) of vertex sets A and B with $V(D) = A \cup B$ for which there is no edge from $B \setminus A$ to $A \setminus B$. The set $A \cap B$ is the *separator* of (A, B) and the vertex sets A and B are the two *sides* of the separation (A, B) . Note that the separator of a separation indeed separates its two sides. The size of the separator of a separation (A, B) is the *order* of (A, B) . Separations of finite order are also called *finite order separations*. There is a natural way to compare separations, namely one defines $(A_1, B_1) \leq (A_2, B_2)$ if $A_1 \subseteq A_2$ and $B_2 \subseteq B_1$. Regarding to this partial order $(A_1 \cup A_2, B_1 \cap B_2)$ is the supremum and $(A_1 \cap A_2, B_1 \cup B_2)$ is the infimum of two separations (A_1, B_1) and (A_2, B_2) . More generally, if $((A_i, B_i))_{i \in I}$ is a family of separations, then

$$\left(\bigcup_{i \in I} A_i, \bigcap_{i \in I} B_i \right) \text{ and } \left(\bigcap_{i \in I} A_i, \bigcup_{i \in I} B_i \right)$$

is its supremum and infimum, respectively.

For vertex sets $A, B \subseteq V(D)$ let $E(A, B)$ be the set of edges from A to B , i.e., $E(A, B) = (A \times B) \cap E(D)$. Given a subdigraph $H \subseteq D$, a *bundle* of H is a non-empty edge set of the form $E(C, C')$, $E(v, C)$, or $E(C, v)$ for strong components C and C' of H and a vertex $v \in V(D) \setminus V(H)$. We say that $E(C, C')$ is a bundle, *between strong components* and $E(v, C)$ and $E(C, v)$ are bundles *between a vertex and a strong component*. In this paper we consider only bundles of subdigraphs H with $H = D - X$ for some $X \in \mathcal{X}(D)$.

Now, consider a vertex $v \in V(D)$, two ends $\omega, \eta \in \Omega(D)$ and a finite vertex set $X \subseteq V(D)$. If X separates ω and η we write $E(X, \omega\eta)$ as short for the edge set $E(C(X, \omega), C(X, \eta))$. Similarly, if $v \in C'$ for a strong component $C' \neq C(X, \omega)$ of $D - X$ we write $E(X, v\omega)$ and $E(X, \omega v)$ as short for the edge set $E(C', C(X, \omega))$ and $E(C(X, \omega), C')$, respectively. If $v \in X$ we write $E(X, v\omega)$ and $E(X, \omega v)$ as short for $E(v, C(X, \omega))$ and $E(C(X, \omega), v)$, respectively. Note that $E(X, \omega\eta)$, $E(X, v\omega)$ and $E(X, \omega v)$ each are bundles if they are non-empty.

An *arborescence* is a rooted oriented tree T that contains for every vertex $v \in V(T)$ a directed path from the root to v . The vertices of any arborescence are partially ordered as $v \leq_T w$ if T contains a directed path from v to w . We write $[v]_T$ for the up-closure of v in T .

A *directed star* is an arborescence whose underlying tree is an undirected star that is centred in the root of the arborescence. A *directed comb* is the union of a ray with infinitely many finite paths (possibly trivial) that have precisely their first vertex on R . Hence the underlying graph of a directed comb is an undirected comb. The *teeth* of a directed comb or reverse directed comb are the teeth of the underlying comb. The ray from the definition of a comb is the *spine* of the comb.

8.2. Necklace Lemma

This section is dedicated to the necklace lemma. We begin with our directed version of the star-comb lemma, which motivates the necklace lemma. Then we continue with the definition of the U -rank, in fact we will define the U -rank in a slightly more general setting by considering not only one set U but finitely many. Finally, we prove the necklace lemma and provide two of its applications.

The star-comb lemma [20] for undirected graphs is a standard tool in infinite graph theory and reads as follows:

Lemma 8.2.1 (Star-Comb Lemma). *Let U be an infinite set of vertices in a connected undirected graph G . Then G contains a comb with all its teeth in U or a subdivided infinite star with all its leaves in U .*

Let us see how to translate the star-comb lemma to digraphs. Given a set U of vertices in a digraph, a comb *attached* to U is a comb with all its teeth in U and a star *attached* to U is a subdivided infinite star with all its leaves in U . The set of teeth is the *attachment set* of the comb and the set of leaves is the *attachment set* of the star. We adapt the notions of ‘attached to’ and ‘attachment sets’ to reverse combs or reverse stars, respectively.

Lemma 8.2.2 (Directed Star-Comb Lemma). *Let D be any strongly connected digraph and let $U \subseteq V(D)$ be infinite. Then D contains a star or comb attached to U and a reverse star or reverse comb attached to U sharing their attachment sets.*

Proof. Since D is strongly connected we find a spanning arborescence T . Applying the star-comb lemma in the undirected tree underlying T yields a comb or a star attached to U . Let U' be the attachment set in either case.

Again using that D is strongly connected we find a reverse spanning arborescence T' . Applying the star-comb lemma a second time, now in the undirected tree underlying T' yields a reverse comb or a reverse star attached to U' . Thinning out the teeth or leaves of the comb or star, respectively, completes the proof. \square

The star-comb lemma fundamentally describes how an infinite set of vertices can be connected in an infinite graph, namely through stars and combs. Similarly, the directed star-comb lemma describes the nature of strongly connectedness in infinite digraphs. Indeed, adding a single path from the first vertex of the reverse comb's spine or centre of the reverse star to the first vertex of the comb's spine or centre of the star, respectively, yields a strongly connected digraph that contains U . We shall use the directed star-comb lemma in the proof of one of our main results of Chapter 9.

As noted in the introduction the star-comb lemma is often used in order to find an end of a given undirected graph G in the closure of an infinite set $U \subseteq V(G)$ of vertices. This is usually done in situations where G contains no infinite subdivided star with all its leaves in U ; for example if the graph is locally finite. Then the star-comb lemma in G applied to U always returns a comb with all its teeth in U and the end represented by the comb's spine is contained in the closure of U .

The directed star-comb lemma however does not manage the task of finding an end of a digraph in the closure of an infinite set of vertices. Consider for example the digraph D that is obtained from the digraph in Figure 8.0.1 by subdividing each vertical edge once. We write U for the set of subdividing vertices. As D contains neither an infinite star nor an infinite reverse star, the directed star-comb lemma applied to U returns a comb attached to U and a reverse comb attached to U sharing their attachment sets. Therefore we would expect that the ends that are represented by the spines are contained in the closure of U . But U does not have any end in its closure because the subdividing vertices all lie in singleton strong components of D .

The necklace lemma will perform the task of finding an end in the closure of a given set of vertices. Before we state it, we need to introduce the \mathcal{U} -rank for

digraphs: For this, consider a finite set \mathcal{U} and think of \mathcal{U} as consisting of infinite sets of vertices. We define in a transfinite recursion the class of digraphs that have a \mathcal{U} -rank. A digraph D has \mathcal{U} -rank 0 if there is a set $U \in \mathcal{U}$ such that $U \cap V(D)$ is finite. It has \mathcal{U} -rank α if it has no \mathcal{U} -rank $< \alpha$ and there is some $X \in \mathcal{X}(D)$ such that every strong component of $D - X$ has a \mathcal{U} -rank $< \alpha$. In the case $U = V(D)$ we call the U -rank of D the *rank of D* (provided that D has a U -rank). Note that if $U \supseteq V(D)$ for a digraph, then its U -rank equals its rank.

We remark that our notion of ranking extends the notion of Schmidt's *ranking of rayless graphs*, in that the rank of a given undirected graph G is precisely the rank of the digraph obtained from G by replacing every edge by its two orientations as separate directed edges, see [64] or Chapter 8.5 of [20] for Schmidt's rank. More generally, for a set U , our U -rank of digraphs extends the notion of the U -rank of graphs, in that an undirected graph G has a U -rank if and only if the digraph that is obtained from G by replacing every edge by its two orientations as separate directed edges has a U -rank; see [5] for the definition of the U -rank of an undirected graph.

Before we prove the necklace lemma, we provide two basic lemmas for the \mathcal{U} -rank of digraphs:

Lemma 8.2.3. *Let D be a digraph and let \mathcal{U} be a finite set. If D has \mathcal{U} -rank α and $H \subseteq D$, then H has some \mathcal{U} -rank $\leq \alpha$.*

Proof. We prove the statement by transfinite induction on the \mathcal{U} -rank of D . Clearly, if D has \mathcal{U} -rank 0, then so does every subdigraph. Let D be a digraph with \mathcal{U} -rank α and $H \subseteq D$. We find a finite vertex set $X \subseteq V(D)$ such that every strong component of $D - X$ has \mathcal{U} -rank less than α . As every strong component of $H - X$ is contained in a strong component of $D - X$, every strong component of $H - X$ has a \mathcal{U} -rank less than α by the induction hypothesis. Hence H has a \mathcal{U} -rank $\leq \alpha$. \square

Lemma 8.2.4. *Let D be any digraph and let \mathcal{U} be a finite set. If D has a \mathcal{U} -rank $\alpha > 0$ and $X \subseteq V(D)$ is a finite vertex set such that every strong component of $D - X$ has a \mathcal{U} -rank $< \alpha$, then infinitely many strong components of $D - X$ meet every set in \mathcal{U} .*

Proof. Suppose for a contradiction that the set \mathcal{C} of strong components of $D - X$ that meet every set in \mathcal{U} is finite. We find for every $C \in \mathcal{C}$ a finite vertex set

$X_C \subseteq V(D)$ witnessing that C has a \mathcal{U} -rank $< \alpha$. Let Y be the union of X with all the finite vertex sets X_C . Then Y witnesses that D has a \mathcal{U} -rank $< \alpha$ contradicting our assumption that D has \mathcal{U} -rank α . \square

Given a set \mathcal{U} , a necklace $N \subseteq D$ is *attached to \mathcal{U}* if infinitely many beads of N meet every set in \mathcal{U} .

Lemma 8.1 (Necklace Lemma). *Let D be any digraph and \mathcal{U} a finite set of vertex sets of D . Then the following statements are complementary:*

- (i) D has a necklace attached to \mathcal{U} ;
- (ii) D has a \mathcal{U} -rank.

Proof. Let us start by showing that not both statements hold at the same time. Suppose for a contradiction there is a digraph D that has a \mathcal{U} -rank and contains a necklace attached to \mathcal{U} as a subdigraph. Then, by Lemma 8.2.3, every necklace $N \subseteq D$ has a \mathcal{U} -rank. But deleting finitely many vertices from any necklace attached to \mathcal{U} leaves a strong component that is a necklace attached to \mathcal{U} by its own. Hence choosing a necklace $N \subseteq D$ attached to \mathcal{U} with minimal \mathcal{U} -rank results in a contradiction.

In order to prove that at least one of (i) and (ii) holds, let us assume that D has no \mathcal{U} -rank. Then for every $X \in \mathcal{X}(D)$, the digraph $D - X$ has a strong component that has no \mathcal{U} -rank. In particular, every such strong component contains a vertex—in fact infinitely many—from every set in \mathcal{U} .

We will recursively construct an ascending sequence $(H_n)_{n \in \mathbb{N}}$ of inflated symmetric paths with finite branch sets, so that H_n extends H_{n-1} , by adding an inflated vertex Y_n that meets every set in \mathcal{U} . In order to make the construction work, we will make sure that Y_n is contained in a strong component of $D - X_n$ that has no \mathcal{U} -rank, where $X_n = H_n \setminus Y_n$. The overall union of the H_n then gives a necklace attached to \mathcal{U} .

Let $H_0 = Y_0$ be a finite strongly connected vertex set that is included in a strong component of $D = D - \emptyset$, that has no \mathcal{U} -rank, and that meets every set in \mathcal{U} . Now, suppose that $n \in \mathbb{N}$ and that H_n and Y_n have already been defined. Let C be the strong component of $D - X_n$ that includes Y_n . As C has no \mathcal{U} -rank, the digraph $C - Y_n$ has a strong component C' that has no \mathcal{U} -rank. Let P be a path in C from Y_n to C' and Q a path from C' to Y_n . Let $Y_{n+1} \subseteq C'$ be a

strongly connected vertex set that contains the last vertex of P , the first vertex of Q and one vertex of every set in \mathcal{U} . We define H_{n+1} to be the union of H_n , P , Q and Y_{n+1} . \square

As our first application of the necklace lemma we describe the connection between Zuther's notion of ends from [72], which we call pre-ends, with our notion of ends. Two rays or reverse rays $R_1, R_2 \subseteq D$ are *equivalent*, if there are infinitely many disjoint paths from R_1 to R_2 and infinitely many disjoint paths from R_2 to R_1 . We call the equivalence classes of this relation the *pre-ends* of D .

Lemma 8.2.5. *Let D be any digraph and γ a pre-end of D . Then γ includes an end ω of D if and only if γ is represented both by a ray and a reverse ray. Moreover, ω is the unique end of D included in γ .*

Proof. Consider any pre-end γ of D . For the forward implication suppose that γ includes an end ω of D . Then there is a ray R that is solid in D and that represents γ . It suffices to find a necklace that is attached to $U := V(R)$. Indeed, every necklace N contains a ray and a reverse ray and if N is attached to R then these rays must be equivalent to R .

So suppose for a contradiction that there is no such necklace. Then by the necklace lemma applied to U in D , the digraph D has a U -rank, say α . Let $X \subseteq V(D)$ be a finite vertex set that witnesses that the U -rank of D is α . As $U \subseteq V(D)$ is infinite, we have $\alpha > 0$. Now, it follows by Lemma 8.2.4 that the ray R meets infinitely many strong components of $D - X$. We conclude that R has no tail in any strong component of $D - X$ contradicting that R is solid in D .

For the backward implication we assume that γ is represented by a ray and a reverse ray. We prove that every ray R that represents γ is solid in D . So let R be any ray that represents γ and let R' be a reverse ray that represents γ . As R and R' are equivalent we find a path system \mathcal{P} that consists of infinitely many pairwise disjoint paths from R to R' and infinitely many pairwise disjoint paths from R' to R .

The subdigraph H of D that consists of R , R' and all the paths in \mathcal{P} has exactly one infinite strong component and finitely many finite strong components (possibly none). Moreover, deleting finitely many vertices from H results again in exactly one infinite strong component and finitely many finite strong components.

Consequently, R has a tail that is contained in a strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$.

For the ‘moreover’ part note that the above argument shows that any ray that represents γ has a tail in the same strong component of $D - X$ as the reverse ray R' , for every finite vertex set $X \subseteq V(D)$. Consequently, any two rays that represent γ have a tail in the same strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. \square

Our second application of the necklace lemma demonstrates how the rank can be used to prove statements about digraphs that have no end. A set of vertices of a digraph D is *acyclic* in D if its induced subdigraph does not contain a directed cycle. The *dichromatic number* [52] of a digraph D is the smallest cardinal κ so that D admits a vertex partition into κ partition classes that are acyclic in D . As a consequence of the necklace lemma we obtain a sufficient condition for D to have a countable dichromatic number:

Theorem 8.2.6. *If D is a digraph that contains no necklace as a subdigraph, then the dichromatic number of D is countable.*

Proof. By the necklace lemma, the statement that D contains no necklace as a subdigraph is equivalent to the statement that D has a rank. Therefore we can apply induction on the rank of D . The vertex set of a finite digraph clearly has a partition into finitely many singleton—and thus acyclic—partition classes, which settles the base case. Now assume that D has rank $\alpha > 0$ and that the statement is true for all ordinals $< \alpha$. We find a finite vertex set $X \subseteq V(D)$ such that every strong component of $D - X$ has some rank $< \alpha$. Hence the induction hypothesis yields a partition $\{V_n(C) \mid n \in \mathbb{N}\}$ of every strong component C of $D - X$ into acyclic partition classes. For every $n \in \mathbb{N}$, let V_n consist of the union of all the sets $V_n(C)$ with C a strong component of $D - X$. Note that V_n is acyclic in D . Combining a partition of X into singleton partition classes with the partition $\{V_n \mid n \in \mathbb{N}\}$ of $V(D - X)$ completes the induction step. \square

8.3. Directions

In this section we will prove the main result of this chapter. To state it properly we need two definitions. A *direction* of a digraph D is a map f with domain $\mathcal{X}(D)$ that sends every $X \in \mathcal{X}(D)$ to a strong component or a bundle of $D - X$ so that $f(X) \supseteq f(Y)$ whenever $X \subseteq Y$. We call a direction f of D a *vertex-direction* if $f(X)$ is a strong component of $D - X$ for every $X \in \mathcal{X}(D)$.

Every end of D naturally defines a direction f_ω which maps every finite vertex set $X \subseteq V(D)$ to the unique strong component of $D - X$ in which every ray that represents ω has a tail. Now, our main theorem reads as follows:

Theorem 8.2. *Let D be any infinite digraph. The map $\omega \mapsto f_\omega$ with domain $\Omega(D)$ is a bijection between the ends and the vertex-directions of D .*

The proof of this needs some preparation. Let D be any digraph and let \mathcal{U} be a set of vertex sets of D . We say that an end ω of D is contained in the *closure* of \mathcal{U} if $C(X, \omega)$ meets every vertex set in $U \in \mathcal{U}$ for every finite vertex set $X \subseteq V(D)$. In Chapter 9 we will define a topology on the space $|D|$ formed by D together with its ends and limit edges and in this topology an end ω will be in the closure of \mathcal{U} if and only if it is in the topological closure of every set in \mathcal{U} . Note that an end ω is contained in the closure of the vertex set of a ray R if and only if R represents ω .

Similarly, we say that a vertex-direction f of D is contained in the *closure* of \mathcal{U} , if $f(X)$ meets every $U \in \mathcal{U}$ for every $X \in \mathcal{X}(D)$. Note that if f is contained in the closure of \mathcal{U} , then $f(X)$ meets every $U \in \mathcal{U}$ in an infinite vertex set. The following lemma describes the connection between ends in the closure of \mathcal{U} , vertex-directions in the closure of \mathcal{U} and necklaces attached to \mathcal{U} :

Lemma 8.3.1. *Let D be any digraph, and let \mathcal{U} be a finite set of vertex sets of D . Then the following assertions are equivalent:*

- (i) D has an end in the closure of \mathcal{U} ;
- (ii) D has a vertex-direction in the closure of \mathcal{U} ;
- (iii) D has a necklace attached to \mathcal{U} .

Proof. (i) \rightarrow (ii): Let ω be any end in the closure of \mathcal{U} . It is straightforward to check that f_ω is a vertex-direction in the closure of \mathcal{U} .

(ii)→(iii): Suppose that f is a vertex-direction in the closure of \mathcal{U} . We need to find a necklace attached to \mathcal{U} . By the necklace lemma we may equivalently show that D has no \mathcal{U} -rank. Suppose for a contradiction that D has a \mathcal{U} -rank α . By Lemma 8.2.3 subdigraphs of digraphs that have a \mathcal{U} -rank have a \mathcal{U} -rank and thus we may choose X' such that $f(X')$ has the smallest \mathcal{U} -rank among all $f(X)$ with $X \in \mathcal{X}(D)$. Note that $f(X)$ has \mathcal{U} -rank ≥ 1 for every $X \in \mathcal{X}(D)$. Indeed, if $f(X) \cap U$ is finite for some $U \in \mathcal{U}$, then

$$f(X \cup (f(X) \cap U)) \cap U = \emptyset$$

contradicting that f is a vertex-direction in the closure of \mathcal{U} . Hence we find a finite vertex set $X'' \subseteq f(X')$ such that all strong components of $f(X') - X''$ have \mathcal{U} -rank less than that of $f(X')$. But then $X' \cup X''$ would have been a better choice for X' .

(iii)→(i): Given a necklace N attached to U , let $R \subseteq N$ be a ray. Then R is solid in D . It is straightforward to show that the end that is represented by R is contained in the closure of \mathcal{U} . \square

Let D be any digraph and let f be any vertex-direction of D . We think of a separation (A, B) of D as pointing towards its side B . Now, if (A, B) is a finite order separation of D , then $f(A \cap B)$ is either included in $B \setminus A$ or $A \setminus B$. In the first case we say that (A, B) *points towards* f and in the second case we say that (A, B) *points away from* f . If two separations point towards or away from f , then the same is true for their supremum or infimum, respectively:

Lemma 8.3.2. *Let D be any digraph and let f be a vertex-direction of D . Suppose that (A_1, B_1) and (A_2, B_2) are finite order separations of D .*

- (i) *If (A_1, B_1) and (A_2, B_2) point towards f , then $(A_1 \cup A_2, B_1 \cap B_2)$ points towards f .*
- (ii) *If (A_1, B_1) and (A_2, B_2) point away from f , then $(A_1 \cap A_2, B_1 \cup B_2)$ points away from f .*

Proof. (i) We have to show that for $(A, B) := (A_1 \cup A_2, B_1 \cap B_2)$ and $X = A \cap B$ the strong component $f(X)$ is included in $B \setminus A$. For this let us consider the auxiliary separation $(A, B') := (A, X' \cup B)$, where $X' := \bigcup_{i=1,2} A_i \cap B_i$ (cf. Figure 8.3.1).




Figure 8.3.1.: The separations (A, B) and (A, B') from the proof of Lemma 8.3.2.

Recall that the separator of a separation separates its two sides. Hence the vertex set $B \setminus A$ is partitioned into the strong components of $D - X$ that it meets.

First, we observe that (A, B') points towards f , a fact that we verify as follows: Since $A_i \cap B_i \subseteq X'$ and because (A_i, B_i) points towards f we have

$$f(X') \subseteq f(A_i \cap B_i) \subseteq B_i$$

for $i = 1, 2$. Hence $f(X') \subseteq B_1 \cap B_2 = B$. Now, $f(X')$ avoids X' because it is a strong component of $D - X'$, giving $f(X') \subseteq B \setminus X'$. As $B \setminus X' = B' \setminus A$, we conclude that $f(X')$ is included in $B' \setminus A$.

Second, we observe that the strong components of $D - X$ that partition $B \setminus A$ are exactly the strong components of $D - X'$ that meet $B' \setminus A$; the reason for this is that B' is obtained from B by adding only vertices from $A \setminus B$.

Finally, we employ the two observations in order to prove that (A, B) points towards f . Indeed, since $X \subseteq X'$ and because f is a vertex-direction, we have that $f(X') \subseteq f(X)$. Now, the first observation says that $f(X')$ is included in $B' \setminus A$. Together with the second observation we obtain $f(X') = f(X)$. So the equation $B' \setminus A = B \setminus A$ yields $f(X) \subseteq B \setminus A$ as desired.

(ii) Apply (i) to the reverse of D . □

Recall, that for a given undirected graph G a vertex v is said to *dominate* an end ω of G if there is an infinite v - R fan in G for some (equivalently every) ray R that represents ω . Equivalently v dominates ω if v is contained in $C(X, \omega)$ for every finite vertex set $X \subseteq V(G) \setminus \{v\}$. An end $\omega \in \Omega(G)$ is *dominated* if some vertex of G dominates it. Ends not dominated by any vertex of G are *undominated*, see [20]. The main case distinction in the proof of Diestel and Kühn's theorem [25, Theorem 2.2], which states that the ends of an undirected

graph correspond bijectively to its directions, essentially distinguishes between directions that correspond to dominated ends and those that correspond to undominated ends. Our plan is to make a similar case distinction for which we need a concept of domination for ends of digraphs.

Let D be any digraph. For a vertex $a \in V(D)$ and $B \subseteq V(D)$ a set of a - B paths in D is called an a - B fan if any two of the paths meet precisely in a . Similarly, a set of B - a paths in D is called an a - B reverse fan if any two of the paths meet precisely in a . We say that a vertex $v \in V(D)$ dominates a ray $R \subseteq D$ if there is an infinite v - R fan in D . The vertex v dominates an end $\omega \in \Omega(D)$ if it dominates some (equivalently every) ray that represents ω . Similarly, a vertex $v \in V(D)$ reverse dominates a ray $R \subseteq D$ if D contains a v - R reverse fan. The vertex v reverse dominates an end $\omega \in \Omega(D)$ if it reverse dominates some (equivalently every) ray that represents ω . An end of D is dominated or reverse dominated if some vertex dominates or reverse dominates it, respectively.

Now, we translate the concept of domination and reverse domination to vertex-directions of digraphs. A vertex $v \in V(D)$ dominates a vertex-direction f in D , if $v \in A$ for every finite order separation (A, B) of D that points away from f . If f is dominated by some vertex, then it is dominated. Similarly, v reverse dominates f if $v \in B$ for every finite order separation (A, B) of D that points towards f . If f is reverse dominated by some vertex, then f is reverse dominated. The following proposition shows that our translation of the concept forwards and reverse domination to vertex-directions of digraphs is accurate:

Proposition 8.3.3. *Let D be any digraph and ω an end of D . A vertex (reverse) dominates ω if and only if it (reverse) dominates f_ω .*

Proof. We prove the statement in its ‘dominates’ version; for the ‘reverse dominates’ version consider the reverse of D . First, suppose that $v \in V(D)$ dominates ω and let (A, B) be a finite order separation pointing away from f_ω . Every ray R that represents ω has a tail in $f_\omega(A \cap B)$; in particular in $D[A]$. As D contains an infinite v - R fan and the separator of (A, B) is finite, it follows that v is contained in A as well.

For the backward implication suppose that $v \in V(D)$ dominates f_ω . Given a ray R that represents ω , with $v \notin R$ say, we need to find an infinite v - R fan in D . For this, we show that every finite v - R fan F in D can be extended by one

additional v - R path; then an infinite such fan can be constructed recursively in countably many steps. Let H be the union of the paths in F and let X consist of $V(H - v)$ together with the vertices of some finite initial segment of R that contains all the vertices that H meets on R . We may view the strong components of $D - X$ partially ordered by $C_1 \leq C_2$ if there is a path in $D - X$ from C_1 to C_2 . Let C be the strong component of $D - X$ that contains v and let $[C]$ be the set of all the strong components of $D - X$ that are $\geq C$. If $C(X, \omega)$ is contained in $[C]$, then it is easy to find a v - R path in D that extends our fan F . We claim that this is always the case: Otherwise consider the finite order separation (A, B) with $A := V(D) \setminus \bigcup [C]$ and $B := X \cup \bigcup [C]$. On the one hand, (A, B) points away from f_ω . On the other hand, we have $v \notin A$ contradicting that v dominates f_ω . \square

Lemma 8.3.4. *Let D be any strongly connected digraph and let f be any vertex-direction of D . Then the following assertions are complementary:*

- (i) f is (reverse) dominated;
- (ii) there is a strictly descending (ascending) sequence $((A_i, B_i))_{i \in \mathbb{N}}$ of finite order separations in D with pairwise disjoint separators all pointing away from (towards) f .

Moreover, a vertex-direction f as in (ii) is the unique vertex-direction in the closure of U for any vertex set U consisting of one vertex of $f(A_i \cap B_i)$ for every $i \in \mathbb{N}$.

Proof. We prove the case where f is dominated and that the sequence in (ii) is descending; the proof of the case where f is reverse dominated and the sequence in (ii) is ascending can then be obtained by considering the reverse of D . To begin, we will show that not both assertions can hold at the same time. Suppose that $((A_i, B_i))_{i \in \mathbb{N}}$ is as in (ii). We show that for every $v \in V(D)$ there is a separation (A, B) of D pointing away from f with $v \in B \setminus A$. We claim that $(A, B) := (A_j, B_j)$ can be taken for $j \in \mathbb{N}$ large enough, a fact that we verify as follows:

As D is strongly connected there is path from $B_0 \setminus A_0$ to v . Let j be the length of a shortest path P from $B_0 \setminus A_0$ to v . Then v is contained in $B_j \setminus A_j$, because otherwise P would contain $j + 1$ vertices—one from each of the separators $B_i \cap A_i$ with $i \leq j$ —contradicting that P has length $\leq j$.

Next, we assume that f is not dominated and construct a sequence $((A_i, B_i))_{i \in \mathbb{N}}$ as in (ii). Let (A_0, B_0) be any finite order separation with non-empty separator pointing away from f . To see that such a separation exist consider any non-empty finite vertex set $X \subseteq V(D)$. We may view the strong components of $D - X$ partially ordered by $C_1 \leq C_2$ if there is a path in $D - X$ from C_1 to C_2 . Let $\lceil f(X) \rceil$ be the down-closure of all the strong components $\leq f(X)$. Then we can take $A_0 := \lceil f(X) \rceil \cup X$ and $B_0 := V(D) \setminus \lceil f(X) \rceil$.

Now, assume that (A_n, B_n) has already been defined. Since no vertex in X_n dominates f we find for every $x \in A_n \cap B_n$ a separation (A_x, B_x) pointing away from f such that $x \in B_x \setminus A_x$. Letting (A_{n+1}, B_{n+1}) be the infimum of all the (A_x, B_x) and (A_n, B_n) completes the construction. Indeed, (A_{n+1}, B_{n+1}) points away from f by Lemma 8.3.2 and its separator is disjoint from all the previous ones as $A_{n+1} \cap B_{n+1} \subseteq A_n \setminus B_n$.

For the ‘moreover’ part let us write $X_i := A_i \cap B_i$ for every $i \in \mathbb{N}$. We first show that f is a vertex-direction in the closure of U . Given $X \in \mathcal{X}(D)$ we need to show that $f(X)$ meets U . With a distance argument as above one finds j such that all the vertices of X are contained in $B_j \setminus A_j$. Then $f(X_j)$ is included in $f(X)$ because $f(X_j) = f(X \cup X_j)$. In particular $f(X)$ contains the vertex from U that was picked from $f(X_j)$.

Finally, we prove that $f = f'$ for every vertex-direction f' that is in the closure of U . Given f' it suffices to show that $f(X_i) = f'(X_i)$ for every $i \in \mathbb{N}$: then $f(X) = f'(X)$ for every $X \in \mathcal{X}(D)$ since we have

$$f(X_j) = f(X \cup X_j) \subseteq f(X) \quad \text{and} \quad f'(X_j) = f'(X \cup X_j) \subseteq f'(X)$$

for j large enough. We verify that $f(X_i) = f'(X_i)$ for every $i \in \mathbb{N}$ as follows: First note that the sequence $(f(X_i))_{i \in \mathbb{N}}$ is descending, because

$$f(X_{i+1}) = f(X_i \cup X_{i+1}) \subseteq f(X_i).$$

Hence $f(X_i)$ contains all but finitely many vertices from U for every $i \in \mathbb{N}$. In particular $f(X_i)$ is the only strong component of $D - X_i$ that contains infinitely many vertices from U . As a consequence we have $f(X_i) = f'(X_i)$ for every $i \in \mathbb{N}$. \square

Proof of Theorem 8.2. It is straightforward to show that the map $\omega \mapsto f_\omega$ with domain $\Omega(D)$ and codomain the set of vertex-directions of D is injective; we prove

that it is onto. So given a vertex-direction f of D we need to find an end $\omega \in \Omega(D)$ such that $f_\omega = f$. Let S_1^* be the set of all the vertices that dominate f and S_2^* the set of all the vertices that reverse dominate f . We split the proof into three cases:

First, assume that both vertex sets $S_1^* \cap f(X)$ and $S_2^* \cap f(X)$ are non-empty for every $X \in \mathcal{X}(D)$. Then f is a vertex-direction in the closure of \mathcal{U} for the set $\mathcal{U} := \{S_1^*, S_2^*\}$. By Lemma 8.3.1 we find an end ω in the closure of \mathcal{U} and we claim that $f_\omega = f$. Indeed, given a finite vertex set $X \in \mathcal{X}(D)$ we need to show that $C(X, \omega) = f(X)$. We may view the strong components of $D - X$ partially ordered by $C_1 \leq C_2$ if there is a path in $D - X$ from C_1 to C_2 . By the order-extension-principle we choose a linear extension of \leq . Let

$$(A_1, B_1) := \left(\bigcup \mathcal{A}_1 \cup X, X \cup \bigcup \mathcal{B}_1 \right),$$

where \mathcal{A}_1 consists of all the strong components of $D - X$ strictly smaller than $f(X)$ and \mathcal{B}_1 of all the others. Then (A_1, B_1) points towards f . Since S_2^* consists of the vertices reverse dominating f we have $S_2^* \subseteq B_1$. Since ω is in the closure of \mathcal{U} we have $C(X, \omega) \in \mathcal{B}_1$. Similarly, let

$$(A_2, B_2) := \left(\bigcup \mathcal{A}_2 \cup X, X \cup \bigcup \mathcal{B}_2 \right),$$

where \mathcal{A}_2 consists of all the strong components of $D - X$ smaller or equal to $f(X)$ and \mathcal{B}_2 of all the others. Analogously to the argumentation for $C(X, \omega) \in \mathcal{B}_1$, one finds out that $C(X, \omega) \in \mathcal{A}_2$; together $C(X, \omega) \in \mathcal{B}_1 \cap \mathcal{A}_2$. Now, $f(X) = C(X, \omega)$ follows from the fact that $f(X)$ is the only element in the intersection $\mathcal{B}_1 \cap \mathcal{A}_2$.

Second, suppose that $S_1^* \cap f(X)$ is empty for some $X \in \mathcal{X}(D)$. If even S_1^* is empty and D strongly connected, then Lemma 8.3.4 and Lemma 8.3.1 do the rest: with U as in the ‘moreover’ part of Lemma 8.3.4, we have that f is the unique vertex-direction in the closure of U and Lemma 8.3.1 yields an end ω in the closure of U ; by uniqueness $f_\omega = f$.

In the following we will argue that we may assume S_1^* to be empty and D to be strongly connected. Fix $X' \in \mathcal{X}(D)$ with $S_1^* \cap f(X') = \emptyset$. Let $D' = f(X')$ and let f' be the vertex-direction of D' induced by f , i.e., f' sends a finite vertex set $X \subseteq V(D')$ to $f(X \cup X')$. Then the set of all the vertices that dominate f' is empty: If (A, B) is a finite order separation of D that points away from f , then $(A \cap V(D'), B \cap V(D'))$ is a finite order separation of D' that points away from f' .

As a consequence, any vertex from D' that dominates f' also dominates f , which means there is none.

Now, consider the end ω' of D' with $f_{\omega'} = f'$ and the unique end ω of D that contains ω' as a subset (of rays). We claim $f_{\omega} = f$, a fact that we verify as follows. First observe that for $X \in \mathcal{X}(D)$ with $X' \subseteq X$ we have

$$f_{\omega}(X) = f_{\omega'}(X \cap V(D')) = f'(X \cap V(D')) = f(X).$$

Now let X be an arbitrary finite vertex set of D . Since f and f_{ω} are vertex-directions we have that $f(X \cup X') \subseteq f(X)$ and $f_{\omega}(X \cup X') \subseteq f_{\omega}(X)$. Furthermore, by our observation we have $f(X \cup X') = f_{\omega}(X \cup X')$. Hence also $f(X) = f_{\omega}(X)$ using that both $f(X)$ and $f_{\omega}(X)$ are strong components of $D - X$.

Finally, the proof of the last case, that $S_2^* \cap f(X)$ is empty for some $X \in \mathcal{X}(D)$, is analogue to the proof of the second case. \square

8.4. Limit edges and edge-directions

In this section, we investigate limit edges of digraphs. Recall that, for two distinct ends $\omega, \eta \in \Omega(D)$, we call the pair (ω, η) a *limit edge* from ω to η , if D has an edge from $C(X, \omega)$ to $C(X, \eta)$ for every finite vertex set $X \subseteq V(D)$ that separates ω and η . For a vertex $v \in V(D)$ and an end $\omega \in \Omega(D)$ we call the pair (v, ω) a *limit edge from v to ω* if D has an edge from v to $C(X, \omega)$ for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. Similarly, we call the pair (ω, v) a *limit edge from ω to v* if D has an edge from $C(X, \omega)$ to v for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. We write $\Lambda(D)$ for the set of limit edges of D . As we do for ‘ordinary’ edges of a digraph, we will suppress the brackets and the comma in our notation of limit edges. For example we write $\omega\eta$ instead of (ω, η) for a limit edge between ends ω and η .

We begin this section with two propositions (Proposition 8.4.1 and Proposition 8.4.2) saying that limit edges are witnessed by subdigraphs that are essentially the digraphs in Figure 8.0.1 or Figure 8.0.3. Subsequently we prove Theorem 8.3.

Let D be any digraph and let $\omega \in \Omega(D)$. With a slight abuse of notation, we say that a necklace $N \subseteq D$ *represents* an end ω of D if one (equivalently every) ray in N represents ω . Note that for every end ω there is a necklace that represents ω . Indeed, apply the necklace lemma to any ray that represents ω .

Proposition 8.4.1. *For a digraph D and two distinct ends ω and η of D the following assertions are equivalent:*

- (i) D has a limit edge from ω to η ;
- (ii) there are necklaces $N_\omega \subseteq D$ and $N_\eta \subseteq D$ that represent ω and η respectively such that every bead of N_ω sends an edge to a bead of N_η .

Moreover, the necklaces may be chosen disjoint from each other and such that the n th bead of N_ω sends an edge to the n th bead of N_η .

Proof. We begin with the forward implication (i) \rightarrow (ii). By possibly deleting a finite vertex set of D that separate ω and η , we may assume that ω and η live in distinct strong components of D . Given a necklace N let us write $N[n, m]$ for the inflated symmetric path from the n th bead to the m th bead of N and $N[n]$ for the inflated symmetric path from the first bead to the n th bead of N . First, let us fix auxiliary necklaces $N'_\omega \subseteq D$ and $N'_\eta \subseteq D$ that represent ω and η , respectively.

We inductively construct sequences $(N_\alpha^n)_{n \in \mathbb{N}}$ of necklaces, for $\alpha \in \{\omega, \eta\}$, so that $N_\alpha^n[n-1] = N_\alpha^{n-1}[n-1]$ and the n th bead of N_ω^n sends an edge to the n th bead of N_η^n . Furthermore, we will make sure that $N'_\alpha[n] \subseteq N_\alpha^n[n]$.

Then the unions $\bigcup\{N_\alpha^n[n] \mid n \in \mathbb{N}\}$ define necklaces N_α , for $\alpha \in \{\omega, \eta\}$, as desired. Indeed, as N_α includes N'_α it also represents α . Note, that our construction yields the ‘moreover’ part. Let $n \in \mathbb{N}$ and suppose that N_ω^n and N_η^n have already been constructed. Let X be the union of $N_\omega^n[n]$, $N_\eta^n[n]$ and the two paths between the n th bead and the $(n+1)$ th bead of N_ω^n and N_η^n , respectively. So X might be empty for $n = 0$. Note that by our assumption ω and η live in distinct strong components of D , so in particular they also live in distinct strong components of $D - X$. As D has a limit edge from ω to η we find an edge e from $C(X, \omega)$ to $C(X, \eta)$. Fix a finite strongly connected vertex set $Y_\alpha \subseteq C(X, \alpha)$ that includes $N_\alpha^n[n+1, m]$ for a suitable $m \geq n+1$ and the endvertex of e in $C(X, \alpha)$ but that avoids the rest of N_α^n for $\alpha \in \{\omega, \eta\}$. Replacing the inflated symmetric subpath $N_\alpha^n[n+1, m]$ by Y_α and declaring Y_α as the $(n+1)$ th bead of N_α^{n+1} for $\alpha \in \{\omega, \eta\}$ yields necklaces N_ω^{n+1} and N_η^{n+1} that are as desired.

Now, let us prove the backward implication (ii) \rightarrow (i). As every finite vertex set X meets only finitely many beads of N_ω and N_η there are beads of N_ω and N_η that are included in $C(X, \omega)$ and $C(X, \eta)$, respectively. Hence, if X separates ω and η , there is an edge from $C(X, \omega)$ to $C(X, \eta)$. \square

There is a natural partial order on the set of ends, where $\omega \leq \eta$ if for every two rays R_ω and R_η that represent ω and η , respectively, there are infinitely many pairwise disjoint paths from R_ω to R_η . By Proposition 8.4.1 we have that $\omega \leq \eta$, whenever $\omega\eta$ is a limit edge for ends ω and η . The converse of this is in general false, for example in the digraph that is obtained from the digraph in Figure 8.0.1 by subdividing every vertical edge once.

Proposition 8.4.2. *For a digraph D , a vertex v and an end ω of D the following assertions are equivalent:*

- (i) *D has a limit edge from v to ω (from ω to v);*
- (ii) *there is a necklace $N \subseteq D$ that represents ω such that v sends (receives) an edge to (from) every bead of N .*

Proof. We consider the case that v sends an edge to every bead of N ; for the other case consider the reverse of D .

For the forward implication (i) \rightarrow (ii) a similar recursive construction as in the proof of Proposition 8.4.1 yields a necklace N as desired.

Now, let us prove the backward implication (ii) \rightarrow (i). As every finite vertex set X hits only finitely many beads of N , there is one bead that is contained in $C(X, \omega)$. Therefore there is an edge from v to $C(X, \omega)$ whenever $v \notin C(X, \omega)$. \square

As a consequence of this proposition, every vertex $v \in V(D)$ for which D has a limit edge from v to an end $\omega \in \Omega(D)$ dominates ω . The converse of this is in general false, for example in the digraph that is obtained from the digraph in Figure 8.0.3 by subdividing every edge once. Similarly, if ωv is a limit edge between an end ω and a vertex v , then v reverse dominates ω ; the converse is again false in general.

Now, let us turn to our second type of directions. We call a direction f of D an *edge-direction*, if there is some $X \in \mathcal{X}(D)$ such that $f(X)$ is a bundle of $D - X$, i.e., if f is not a vertex-direction. Recall that every end defines a vertex-direction. Similarly, every limit edge λ defines an edge-direction as follows.

We say that a limit edge $\lambda = \omega\eta$ *lives* in the bundle defined by $E(X, \lambda)$ if $X \in \mathcal{X}(D)$ separates ω and η . If $X \in \mathcal{X}(D)$ does not separate ω and η , we say that $\lambda = \omega\eta$ *lives* in the strong component $C(X, \omega) = C(X, \eta)$ of $D - X$. We use similar notations for limit edges of the form $\lambda = v\omega$ or $\lambda = \omega v$ with $v \in V(D)$ and

$\omega \in \Omega(D)$: We say that a limit edge λ *lives in* the bundle $E(X, \lambda)$ if $v \notin C(X, \omega)$ and we say that λ *lives in* the strong component $C(X, \omega)$ of $D - X$, if $v \in C(X, \omega)$.

The edge-direction f_λ defined by λ is the edge-direction that sends every finite vertex set $X \subseteq V(D)$ to the bundle or strong component of $D - X$ in which λ lives. Our next theorem states that there is a one-to-one correspondence between the edge-directions of a digraph and its limit edges:

Theorem 8.3. *Let D be any infinite digraph. The map $\lambda \mapsto f_\lambda$ with domain $\Lambda(D)$ is a bijection between the limit edges and the edge-directions of D .*

Proof. It is straightforward to show that the map given in (ii) is injective; we prove onto. So let f be any edge-direction of D . First suppose that $f(X)$ is always a strong component or a bundle between strong components for every $X \in \mathcal{X}(D)$. Then f defines two vertex-directions f_1 and f_2 as follows. If $f(X) = E(C_1, C_2)$ is a bundle then let $f_1(X) = C_1$ and $f_2(X) = C_2$. Otherwise, $f(X)$ is a strong component and we put $f_1(X) = f_2(X) = f(X)$. Now, the inverse of the function from Theorem 8.2 returns ends ω and η for f_1 and f_2 , respectively. We conclude that $\omega\eta$ is a limit edge and that $f = f_{\omega\eta}$.

Now, suppose that f maps some finite vertex set X' to a bundle between a vertex $v \in X'$ and a strong component of $D - X'$. Then also $f(\{v\})$ is a bundle between v and a strong component. We consider the case where $f(\{v\})$ is of the form $E(v, C_v)$ for some strong component C_v of $D - v$; the other case is analogue.

Let us define a vertex-direction f' of D . First, for every $X \in \mathcal{X}(D)$ with $v \in X$ we have that $f(X)$ is a bundle of the form $E(v, C)$ for a strong component C of $D - X$ and we put $f'(X) = C$. Second, if $v \notin X$ for some $X \in \mathcal{X}(D)$ we have that $f(X)$ is either a strong component C' of $D - X$ or a bundle $E(C, C')$ with $v \in C$. We then put $f'(X) = C'$. It is straightforward to check that f' is indeed a vertex-direction. Finally, the inverse of the map from Theorem 8.2 applied to f' returns an end ω . By the definition of f' we have that $v\omega$ is a limit edge of D and a close look to the definitions involved points out that $f = f_{v\omega}$. \square

9. The topological point of view

In 2004, Diestel and Kühn [26] introduced a topological framework for infinite graphs which makes it possible to extend theorems about finite graphs to infinite graphs that do not generalise verbatim. The main point is to consider not only the graph itself but the graph together with its ends, and to equip both together with a suitable topology. For locally finite graphs G , this space $|G|$ coincides with the Freudenthal compactification of G [20, 30].

Diestel and Kühn's approach has become standard and has led to several results found by various authors. Examples include Nash-Williams's tree-packing theorem [18], Fleischner's Hamiltonicity theorem [35], and Whitney's planarity criterion [2]. In the formulation of these theorems, topological arcs and circles take the role of paths and cycles, respectively.

To illustrate this, consider Euler's theorem that a connected finite graph contains an Euler tour if and only if every vertex has even degree. This statement fails for infinite graphs, since a closed walk in a connected infinite graph cannot traverse all its infinitely many edges. Diestel and Kühn [26] extended Euler's Theorem to the space $|G|$ for locally finite graphs G , as follows. A *topological Euler tour* of $|G|$ is a continuous map $\sigma: S^1 \rightarrow |G|$ such that every inner point of an edge of G is the image of exactly one point of S^1 . Hence a topological Euler tour 'traverses' every edge exactly once. Diestel and Kühn showed that for a connected locally finite graph G the space $|G|$ admits a topological Euler tour if and only if every finite cut in G is even. Note that in a finite graph every vertex has even degree if and only if all finite cuts are even, but even for locally finite infinite graphs the latter statement is stronger. Theorem 9.3 below is a directed analogue of the Diestel-Kühn theorem about topological Euler tours.

In Chapter 8 we introduced the concept of *ends* and *limit edges* of a digraph. A *directed ray* is an infinite directed path that has a first vertex (but no last vertex). The directed subrays of a directed ray are its *tails*. For the sake of readability we shall omit the word 'directed' in 'directed path' and 'directed ray' if there is

no danger of confusion. We call a ray in a digraph *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. We call two solid rays in a digraph D *equivalent* if for every finite vertex set $X \subseteq V(D)$ they have a tail in the same strong component of $D - X$. The classes of this equivalence relation are the *ends* of D . The ends of a digraph can be thought of as points at infinity to which its solid rays converge.

For limit edges the situation is similar. Informally, they are additional edges that naturally arise between distinct ends of a digraph, as follows. Unlike graphs, digraphs may have two rays R and R' that represent distinct ends and yet there may be a set E of infinitely many independent edges from R to R' . In this situation there will be a limit edge from the end represented by R to the end represented by R' , and the edges in E can be thought of as converging towards this limit edge. The precise definition of limit edges can be found in Section 8.1.

We begin this chapter by introducing a topology, which we call DTOP , on the space $|D|$ formed by the digraph D together with its ends and limit edges. In this topology rays and edges will converge to ends and limit edges, respectively. As our first main result we characterise those digraphs for which $|D|$ with DTOP compactifies D .

For graphs G , a necessary condition for $|G|$ to be compact is that $G - X$ has only finitely many components for every finite vertex set $X \subseteq V(G)$: if $G - X$ has infinitely many components, then these components together with all ends living in them will form a disjoint family of open sets, and combining this family with a suitable cover of the finite graph $G[X]$ yields an open cover of $|G|$ that has no finite subcover. In [24], Diestel proves that this necessary condition is also sufficient. In analogy to this, let us call a digraph D *solid* if $D - X$ has only finitely many strong components for every finite vertex set $X \subseteq V(D)$.

Our first main result is that Diestel's characterisation carries over to digraphs:

Theorem 9.1. *The space $|D|$ is compact if and only if D is solid.*

We remark that for every digraph D the space $|D|$ is Hausdorff and that D is dense in $|D|$. Hence if D is solid, the space $|D|$ is a Hausdorff compactification of D .

A common way to generalise statements about finite graphs to infinite graphs is to use so-called compactness arguments. These can be phrased in terms of inverse

limits. For example, the ray given by König's infinity lemma [20, Lemma 8.1.2] exists because the inverse limit of compact discrete spaces is non-empty. In order to make this technique applicable for the space $|D|$ we provide the following:

Theorem 9.2. *For a solid digraph D the space $|D|$ is the inverse limit of finite contraction minors of D .*

For the precise statement of this theorem see Section 9.3.

Recall that our motivation for introducing a topology on a digraph D , together with its ends and limit edges, was to extend to the space $|D|$ theorems about finite digraphs that would be either false, or trivial, or undefined for D itself. As a proof of concept, we prove two such applications for $|D|$.

For our first application recall that a finite digraph with a connected underlying graph contains an Euler tour if and only if the in-degree equals the out-degree at every vertex [1]. This statement fails for infinite digraphs, since a closed walk can only traverse finitely many edges.

As in the case of undirected graphs, however, there is a natural topological notion of Euler tours of $|D|$. Call a continuous map $\alpha: [0, 1] \rightarrow |D|$ that respects the direction of the edges of $|D|$ a *topological path in $|D|$* , which is *closed* if $\alpha(0) = \alpha(1)$. Call a closed topological path an *Euler tour* if it traverses every edge exactly once, and call $|D|$ *Eulerian* if it admits an Euler tour. See Section 9.4 for precise definitions.

If $|D|$ is Eulerian, then the in-degree equals the out-degree at every vertex. The converse of this fails in general. For example, the digraph D on \mathbb{Z} with edges $n(n+1)$ for every $n \in \mathbb{Z}$ has in- and out-degrees 1 at every vertex, but $|D|$ has no Euler tour.

A *cut* of a digraph D is an ordered pair (V_1, V_2) of non-empty sets $V_1, V_2 \subseteq V(D)$ such that $V_1 \cup V_2 = V(D)$ and $V_1 \cap V_2 = \emptyset$. The sets V_1 and V_2 are the *sides* of the cut, and its *size* is the cardinality of the set of edges from V_1 to V_2 . We call a cut (V_1, V_2) *balanced* if its size equals that of (V_2, V_1) . Note that in a finite digraph the in-degree at every vertex equals the out-degree if and only if all finite cuts are balanced, but our \mathbb{Z} -example shows that for infinite digraphs, even locally finite ones, the latter statement is stronger.

Any unbalanced finite cut in a digraph D is an obstruction that prevents $|D|$ from being Eulerian: by the directed jumping arc lemma (Lemma 9.4.1), any Euler

tour enters a side of a finite cut as often as it leaves it. A second obstruction is a vertex of infinite in- or out-degree, as an Euler tour that traverses a vertex infinitely often forces the tour to converge to that vertex.

As our first application we show that there are no further obstructions. A digraph is *locally finite* if all of its vertices have finite in- and out-degree.

Theorem 9.3. *For a digraph D with a connected underlying graph the following assertions are equivalent:*

- (i) $|D|$ is Eulerian;
- (ii) D is locally finite and every finite cut of D is balanced.

In our second application we characterise the digraphs that are strongly connected. It is easy to see that a finite digraph is strongly connected if and only if it contains a closed directed walk that contains all its vertices. We obtain the following characterisation of strongly connected infinite digraphs:

Theorem 9.4. *For a countable solid digraph D the following assertions are equivalent:*

- (i) D is strongly connected;
- (ii) there is a closed topological path in $|D|$ that contains all the vertices of D .

We remark that the requirements ‘countable’ and ‘solid’ for D are necessary. Indeed, any closed topological path in $|D|$ that traverses uncountably many vertices also traverses uncountably many edges of D . This gives rise to uncountably many disjoint open intervals in $[0, 1]$, which is impossible. Furthermore, recall that the image of a compact space under a continuous map is compact. In particular the image of every topological path that contains all the vertices of D is compact in $|D|$. Hence the closure of $V(D)$ is compact, which implies that D is solid. Indeed, if $D - X$ has infinitely many strong components for some finite vertex set $X \subseteq V(D)$, then these strong components together with all the ends that live in them will form a disjoint family of open sets, and combining this family with a suitable cover of X yields an open cover of the closure of $V(D)$ that has no finite subcover.

This chapter is organised as follows. In Section 9.1 we collect together the results that we need from Chapter 8, or from general topology. In Section 9.2 we formally

define the topological space $|D|$ for a given digraph D and prove Theorem 9.1. In Section 9.3 we define an inverse system for a given digraph D and show that the inverse limit of this system coincides with $|D|$ if D is solid (Theorem 9.2). Finally, in Section 9.4 we prove our two applications of our framework, Theorem 9.3 and Theorem 9.4.

9.1. Tools and terminology

In this section we provide further tools and terminology that we use throughout this chapter. We also collect together some tools and terminology from Chapter 8 that become particularly important in this chapter.

Let D be any digraph. We write edges as order pairs (v, w) with $v, w \in V(D)$, and usually we write (v, w) simply as vw ; except if D is a multi-digraph in which case we write edges of D as triples (e, v, w) . The vertex v is the *tail* of vw and the vertex w its *head*.

A direction of a digraph D is a map f with domain $\mathcal{X}(D)$ that sends every $X \in \mathcal{X}(D)$ to a strong component or a bundle of $D - X$ so that $f(X) \supseteq f(Y)$ whenever $X \subseteq Y$.¹ We call a direction f on D a *vertex-direction* if $f(X)$ is a strong component of $D - X$ for every $X \in \mathcal{X}(D)$, and we call it an *edge-direction* otherwise, i.e., if $f(X)$ is a bundle of $D - X$ for some $X \in \mathcal{X}(D)$. Every end ω of a digraph D defines a direction f_ω on D in that it maps $X \in \mathcal{X}(D)$ to $C(X, \omega)$. The ends of D correspond bijectively to its vertex-directions:

Theorem 8.2. *Let D be any infinite digraph. The map $\omega \mapsto f_\omega$ with domain $\Omega(D)$ is a bijection between the ends and the vertex-directions of D .*

We denote by $\Lambda(D)$ the set of all the limit edges of D and we use the usual definitions for edges accordingly; for example we will speak of the head and the tail of a limit edge. Every limit edge λ defines an edge-direction as follows. We say that a limit edge $\lambda = \omega\eta$ *lives* in the bundle defined by $E(X, \lambda)$ if $X \in \mathcal{X}(D)$ separates ω and η . If $X \in \mathcal{X}(D)$ does not separate ω and η , we say that $\lambda = \omega\eta$ *lives* in

¹Here, as later in this context, we do not distinguish rigorously between a strong component and its set of edges. Thus if Y separates ω and η but $X \subseteq Y$ does not, the expression $f_{\omega\eta}(X) \supseteq f_{\omega\eta}(Y)$ means that the strong component $f_{\omega\eta}(X)$ of $D - X$ contains all the edges from the edge set $f_{\omega\eta}(Y)$.

the strong component $C(X, \omega) = C(X, \eta)$ of $D - X$. We use similar notations for limit edges of the form $\lambda = v\omega$ or $\lambda = \omega v$ with $v \in V(D)$ and $\omega \in \Omega(D)$: We say that the limit edge λ *lives in* the bundle $E(X, \lambda)$ if $x \notin C(X, \omega)$ and we say that λ *lives in* the strong component $C(X, \omega)$ of $D - X$, if $v \in C(X, \omega)$. The edge-direction f_λ defined by λ is the edge-direction that maps every finite vertex set $X \in \mathcal{X}(D)$ to the bundle or strong component of $D - X$ in which λ lives. The limit edges of any digraph correspond bijectively to its edge-directions:

Theorem 8.3. *Let D be any infinite digraph. The map $\lambda \mapsto f_\lambda$ with domain $\Lambda(D)$ is a bijection between the limit edges and the edge-directions of D .*

With a slight abuse of notation, we say that a necklace $N \subseteq D$ *represents* an end ω of D if one (equivalently every) ray in N represents ω . For limit edges we have the following:

Proposition 8.4.1. *For a digraph D and two distinct ends ω and η of D the following assertions are equivalent:*

- (i) D has a limit edge from ω to η ;
- (ii) there are necklaces $N_\omega \subseteq D$ and $N_\eta \subseteq D$ that represent ω and η respectively such that every bead of N_ω sends an edge to a bead of N_η .

Moreover, the necklaces may be chosen disjoint from each other and such that the n th bead of N_ω sends an edge to the n th bead of N_η .

Proposition 8.4.2. *For a digraph D , a vertex v and an end ω of D the following assertions are equivalent:*

- (i) D has a limit edge from v to ω (from ω to v);
- (ii) there is a necklace $N \subseteq D$ that represents ω such that v sends (receives) an edge to (from) every bead of N .

For a digraph D and a set \mathcal{U} we say that a necklace $N \subseteq D$ is *attached to* \mathcal{U} if infinitely many beads of N meet every set of \mathcal{U} . In Chapter 8 we introduced an ordinal rank function that can be used to find out whether a digraph D contains for a given set \mathcal{U} a necklace attached to \mathcal{U} . For this, consider a finite set \mathcal{U} and think of \mathcal{U} as consisting of infinite sets of vertices. We define in a transfinite recursion the class of digraphs that have a \mathcal{U} -rank. A digraph D has \mathcal{U} -rank 0 if

there is a set $U \in \mathcal{U}$ such that $U \cap V(D)$ is finite. It has \mathcal{U} -rank α if it has no \mathcal{U} -rank $< \alpha$ and there is some $X \in \mathcal{X}(D)$ such that every strong component of $D - X$ has a \mathcal{U} -rank $< \alpha$. In the case $U = V(D)$ we call the U -rank of D the *rank of D* (provided that D has a U -rank). Note that more generally if $U \supseteq V(D)$ for a digraph, then its U -rank equals its rank.

Lemma 8.1 (Necklace Lemma). *Let D be any digraph and \mathcal{U} a finite set of vertex sets of D . Then exactly one of the statements is true:*

- (i) D has a necklace attached to \mathcal{U} ;
- (ii) D has a \mathcal{U} -rank.

Given a set U of vertices in a digraph, a comb *attached* to U is a comb with all its teeth in U and a star *attached* to U is a subdivided infinite star with all its leaves in U . The set of teeth is the *attachment set* of the comb and the set of leaves is the *attachment set* of the star. We adapt the notions of ‘attached to’ and ‘attachment sets’ to reverse combs or reverse stars, respectively.

Lemma 9.1.1 (Directed Star-Comb Lemma). *Let D be any strongly connected digraph and let $U \subseteq V(D)$ be infinite. Then D contains a star or comb attached to U and a reverse star or reverse comb attached to U sharing their attachment sets.*

In the second part of this section, we list the tools and terminology about inverse limits that we need. Here we follow the textbook of Zalesskii and Ribes [50].

Let (I, \leq) be a *directed* partially ordered set, i.e., I is partially ordered by \leq and for any two elements $i, j \in I$ there exist an element $k \in I$ such that $i, j \leq k$. A collection $\{X_i \mid i \in I\}$ of topological spaces together with continuous maps $f_{ji}: X_j \rightarrow X_i$, for all $i \leq j$, is called *inverse system* if $f_{ki} = f_{ji} \circ f_{kj}$ whenever $i \leq j \leq k$ and f_{ii} is the identity on X_i , for all $i \in I$. We denote such an inverse system by $\{X_i, f_{ij}, I\}$. The continuous maps $f_{ji}: X_j \rightarrow X_i$ are called *bonding maps*. The *inverse limit* $\varprojlim (X_i)_{i \in I}$ is the subspace of the product space $\prod_{i \in I} X_i$ that consists of all the $(x_i)_{i \in I}$ with $f_{ji}(x_j) = x_i$ for all $i \leq j$. In this setup, we write f_i for the projection from $\varprojlim (X_i)_{i \in I}$ to X_i . If all the X_i are Hausdorff, the inverse limit is closed in the product space. Therefore, by Tychonoff’s theorem, if all the X_i are in addition compact, then the inverse limit is compact. For a topological

space Y together with continuous maps $\varphi_i: Y \rightarrow X_i$, for all $i \in I$, the collection of maps $\{\varphi_i \mid i \in I\}$ is called *compatible* if $\varphi_i = f_{ji} \circ \varphi_j$ for all $i \leq j$. The inverse limit of an inverse system is (up to unique homeomorphism) characterised by the following *universal property*:

For every topological space Y together with compatible maps $\varphi_i: Y \rightarrow X_i$, for $i \in I$, there is a unique continuous map $\Phi: Y \rightarrow \varprojlim (X_i)_{i \in I}$ with $\varphi_i = f_i \circ \Phi$ for all $i \in I$.

In this situation, we say that the map Φ is *induced* by the maps φ_i . For a topological space Y together compatible maps $\varphi_i: Y \rightarrow X_i$, for $i \in I$, the collection of maps $\{\varphi_i \mid i \in I\}$ is called *eventually injective* if for every two distinct $y, y' \in Y$ there is some $i \in I$ with $\varphi_i(y) \neq \varphi_i(y')$, see [20, Lemma 8.8.4].

Lemma 9.1.2 (Lifting Lemma). *Let $\{X_i, f_{ij}, I\}$ be any inverse system and let Y be a topological space together with eventually injective maps $\varphi_i: Y \rightarrow X_i$, for $i \in I$. Then the unique continuous map $\Phi: Y \rightarrow \varprojlim (X_i)_{i \in I}$ given by the universal property of the inverse limit is injective.*

We need the following two results [50, Lemma 1.1.7] and [50, Lemma 1.1.9]:

Lemma 9.1.3. *Let $\{X_i, f_{ij}, I\}$ be an inverse system of topological spaces over a directed set I , and let $\varphi_i: X \rightarrow X_i$ be surjections from the space X onto the spaces X_i ($i \in I$). Then either $\varprojlim (X_i)_{i \in I} = \emptyset$ or the induced mapping $\Phi: X \rightarrow \varprojlim (X_i)_{i \in I}$ maps X onto a dense subset of $\varprojlim (X_i)_{i \in I}$.*

For a partially ordered set I a subset $I' \subseteq I$ is called *cofinal* if for all $i \in I$ there is an $i' \in I'$ with $i \leq i'$.

Lemma 9.1.4. *Let $\{X_i, f_{ij}, I\}$ be an inverse system of compact topological spaces over a directed poset I and assume that I' is a cofinal subset of I . Then*

$$\varprojlim (X_i)_{i \in I} \cong \varprojlim (X_{i'})_{i' \in I'}.$$

Finally, we need the following theorem [32, Theorem 3.1.13] from basic topology:

Lemma 9.1.5. *Every continuous one-to-one mapping of a compact space onto a Hausdorff space is a homeomorphism.*

9.2. A topology for digraphs

In this section we define a topology on the space $|D|$ formed by a digraph D together with its ends and limit edges that we call DTOP.

In this topological space, topological arcs and circles take the role of paths and cycles, respectively. This makes it possible to extend to the space $|D|$ statements about finite digraphs. As an important cornerstone we characterise, in this section, those digraphs D for which $|D|$ is compact, see Theorem 9.1.

Consider a digraph $D = (V, E)$ with its set $\Omega = \Omega(D)$ of ends and its set $\Lambda = \Lambda(D)$ of limit edges. The ground set $|D|$ of our topological space is defined as follows. Take $V \cup \Omega$ together with a copy $[0, 1]_e$ of the unit interval for every edge $e \in E \cup \Lambda$. Now, identify every vertex or end x with the copy of 0 in $[0, 1]_e$ for which x is the tail of e and with the copy of 1 in $[0, 1]_f$ for which x is the head of f , for all $e, f \in E \cup \Lambda$.

For inner points $z_e \in [0, 1]_e$ and $z_f \in [0, 1]_f$ of edges $e, f \in E \cup \Lambda$ we say that z_e *corresponds* to z_f if both correspond to the same point of the unit interval. For $e \in E \cup \Lambda$ the point set obtained from $[0, 1]_e$ in $|D|$ is an *edge of $|D|$* . The vertex or end that was identified with the copy of 0 is the *tail* of the edge of $|D|$ and the vertex or end that was identified with the copy of 1 its *head*.

We define the topological space DTOP on $|D|$ by specifying the basic open sets. For a vertex v we take the collection of uniform stars of radius ε around v as basic open neighbourhoods. For inner points z of edges $[0, 1]_e$ with $e \in E$ we keep the open balls around z of radius ε as basic open sets (considered as subsets of $[0, 1]_e$). Here we make the convention that for edges e (possibly limit edges) the ε of open balls $B_\varepsilon(z)$ of radius ε around points $z \in e$ is implicitly chosen small enough to guarantee $B_\varepsilon(z) \subseteq e$.

Neighbourhoods $\hat{C}_\varepsilon(X, \omega)$ of an end ω are of the following form: Given $X \in \mathcal{X}(D)$ let $\hat{C}_\varepsilon(X, \omega)$ be the union of (see Figure 9.2.1)

- the point set of $C(X, \omega)$,
- the set of all ends and points of limit edges that live in $C(X, \omega)$ and
- half-open partial edges $(\varepsilon, y]_e$ respectively $[y, \varepsilon)_e$ for every edge $e \in E \cup \Lambda$ for which y is contained or lives in $C(X, \omega)$.




Figure 9.2.1.: A basic open neighbourhood of the form $\hat{C}_\varepsilon(X, \omega)$.

Neighbourhoods $\hat{E}_{\varepsilon, z}(X, \omega\eta)$ of an inner point z of a limit edge $\omega\eta$ between ends are of the following form: Given $X \in \mathcal{X}(D)$ that separates ω and η let $\hat{E}_{\varepsilon, z}(X, \omega\eta)$ be the union of (see Figure 9.2.2)

- the open balls of radius ε around points z_e of edges $e \in E(X, \omega\eta)$ and with z_e corresponding to z and
- the open balls of radius ε around points z_λ of limit edges λ that live in the bundle $E(X, \omega\eta)$ and with z_λ corresponding to z .

Similarly, for an inner point z of a limit edge $v\omega$ between a vertex v and an end ω we define the open neighbourhoods $\hat{E}_{\varepsilon, z}(X, v\omega)$ as follows. Given $X \in \mathcal{X}(D)$ with $v \in X$ let $\hat{E}_{\varepsilon, z}(X, v\omega)$ be the union of

- the open balls of radius ε around points z_e of edges $e \in E(X, v\omega)$ and with z_e corresponding to z and
- the open balls of radius ε around points z_λ of limit edges λ that live in the bundle $E(X, v\omega)$ and with z_λ corresponding to z .

Open sets $\hat{E}_{\varepsilon, z}(X, \omega v)$ for a limit edge ωv between an end ω and a vertex v are defined analogously.




Figure 9.2.2.: A basic open neighbourhood of the form $\hat{E}_{\epsilon, z}(X, \omega\eta)$.

We view a digraph D as a subspace of $|D|$, namely the subspace that is formed by all the (equivalence classes of) vertices and inner points of edges of D . If there is no danger of confusion we will not distinguish between the digraph D and the topological space D . Furthermore, we call the subspace $\Omega(D)$ of $|D|$ the *end space* of D . The end space of an undirected graph G coincides with the end space of the digraph obtained from G by replacing every edge by its two orientations as separate directed edges.

One of the key definitions in Chapter 8, was that an end ω of D is said to be in the closure of \mathcal{U} , for a set of vertex sets \mathcal{U} , if for all $X \in \mathcal{X}(D)$ every $U \in \mathcal{U}$ has a vertex in $C(X, \omega)$. Now that DTOP is at hand this is tantamount to $\omega \in \bar{U}$ for every $U \in \mathcal{U}$. We therefore obtain an extension of [10, Lemma 4.1]:

Lemma 9.2.1. *Let D be any digraph, and let \mathcal{U} be a finite set of vertex sets of D . Then the following assertions are equivalent:*

- (i) D has an end in the closure of \mathcal{U} ;
- (ii) D has a vertex-direction in the closure of \mathcal{U} ;
- (iii) D has a necklace attached to \mathcal{U} ;
- (iv) D has an end in $\bigcap \{ \bar{U} \mid U \in \mathcal{U} \}$.

Recall that we call a digraph D *solid* if $D - X$ has finitely many strong components for every $X \in \mathcal{X}(D)$. The main result of this section reads as follows:

Theorem 9.1. *The space $|D|$ is compact if and only if D is solid.*

Proof. We prove the forward implication with contraposition. If D is not solid let X be a finite vertex set such that $D - X$ has infinitely many strong components. We obtain an open cover \mathcal{O} of $|D|$ that has no finite subcover as follows. Fix for every strong component C of $D - X$ a vertex $u_c \in C$ and denote by U the set of all the vertices u_c . It is straightforward to check that every point in $|D| \setminus U$ has a basic open neighbourhood that avoids U ; this shows that U is closed in $|D|$. Let \mathcal{O} consist of the uniform stars of radius $\frac{1}{2}$ around each u_c and the open set $|D| \setminus U$. Then, \mathcal{O} is the desired open cover.

Now, let us prove the backward implication. For this, let D be any solid digraph and let \mathcal{O} be an open cover of $|D|$. We may assume that \mathcal{O} consists of basic open sets. For every $X \in \mathcal{X}(D)$ and every strong component C of $D - X$, we let \hat{C} be the union of the point set of C , the set of all the ends that live in C and the point set of all the limit edges that live in C . For a bundle F of $D - X$, let \hat{F} consist of the inner points of edges in F and all the inner points of limit edges that live in F . A strong component C of $D - X$ is *bad for X* if \hat{C} is not covered by any cover set in \mathcal{O} . A bundle F of $D - X$ is *bad for X* if \hat{F} is not covered by finitely many cover sets in \mathcal{O} . A bad strong component for X or a bad bundle for X is a *bad set for X* .

If there is no bad set for some $X \in \mathcal{X}(D)$, we find a finite subcover as follows. For every strong component C of $D - X$ fix a cover set from \mathcal{O} that covers \hat{C} . And for every bundle of $D - X$ fix finitely many cover sets from \mathcal{O} that cover \hat{F} . Note that our assumption that D is solid ensures that there are only finitely many strong components and bundles of $D - X$. Therefore, we have fixed only finitely many cover sets in total. Combining these with a finite subcover of $D[X]$, which exists because $D[X]$ is a finite digraph, yields a finite subcover of $|D|$. Note that all the edges between vertices $x \in X$ and strong components of $D - X$ are covered, as they are bundles.

So let us assume for a contradiction that there is a bad set for every $X \in \mathcal{X}(D)$. We will find a bad set for every $X \in \mathcal{X}(D)$ in a consistent way, i.e., for every two vertex sets $X, Y \in \mathcal{X}(D)$ the bad set of X contains that of Y whenever $X \subseteq Y$. In

other words the bad sets will give rise to a direction f and we will then conclude that $f(X)$ is covered by finitely many sets in \mathcal{O} for some $X \in \mathcal{X}(D)$, contradicting that $f(X)$ is bad.

Given $X \in \mathcal{X}(D)$, let B_X be the union of all the sets \hat{B} for which B is bad for X . It is straightforward to see that $\{B_X : X \in \mathcal{X}\}$ is a filter base on $|D|$ and we denote by \mathcal{B} some ultrafilter that extends it. On the one hand, \mathcal{B} contains for every $X \in \mathcal{X}(D)$ at most one set \hat{C} or \hat{F} with C a strong component of $D - X$ or F a bundle of $D - X$, respectively, because intersections of filter sets are non-empty. On the other hand, there is at least one strong component C or bundle F of $D - X$ such that \hat{C} or \hat{F} is contained in \mathcal{B} : Otherwise, \mathcal{B} contains $|D| \setminus \hat{C}$ and $|D| \setminus \hat{F}$ for every strong component of $D - X$ respectively every bundle F of $D - X$. As \mathcal{B} does not contain the point set of $D[X]$ we have $|D| \setminus D[X] \in \mathcal{B}$. But, the intersection of all the $|D| \setminus \hat{C}$ and $|D| \setminus \hat{F}$ with $|D| \setminus D[X]$ is empty. Consequently, \mathcal{B} contains for every $X \in \mathcal{X}(D)$ exactly one set of the form \hat{C} or \hat{F} with C a strong component of $D - X$ or F a bundle of $D - X$, respectively. As intersections of filter sets are non-empty, these bundles and strong components form a direction f . Note, that for every $X \in \mathcal{X}(D)$ the set $f(X)$ is bad for X as it is the superset of B_Y for some $Y \in \mathcal{X}(D)$.

In order to arrive at a contradiction we consider three cases. First, if f is a vertex-direction, then by Theorem 8.2, we have that f corresponds to an end ω which is covered by some cover set $O \in \mathcal{O}$. As O is a basic open set it is of the form $\hat{C}_\varepsilon(X, \omega)$ for some $X \in \mathcal{X}(D)$. This contradicts that $f(X)$ is bad.

Second, suppose that f is an edge-direction and that f corresponds to a limit edge $\omega\eta$ between ends in the sense of Theorem 8.3. This limit edge $\omega\eta$ is covered by a finite subset $\mathcal{O}' \subseteq \mathcal{O}$, as it is homeomorphic to the unit interval. Since each cover set $O \in \mathcal{O}'$ is basic open it comes by its definition together with a finite vertex set $X_O \in \mathcal{X}(D)$. Let $\mathcal{X}' := \{X_O \mid O \in \mathcal{O}'\}$ and let X be large enough so that it contains $\bigcup \mathcal{X}'$ and so that it separates ω and η . To get a contradiction, we show that $\widehat{f(X)}$ is covered by \mathcal{O}' . Consider a point $z \in \widehat{f(X)}$ and let z' be its corresponding point on $\omega\eta$. Then z' is covered by some $O \in \mathcal{O}'$. Since $X_O \subseteq X$ we have $f(X) \subseteq f(X')$ and therefore \mathcal{O}' also contains z .

Finally, the case that f is an edge-direction and f corresponds to a limit edge between an end and a vertex is analogue to the second case. \square

We complete this section by listing a few more properties that are equivalent to the assertion that $|D|$ is compact.

Corollary 9.2.2. *The following statements are equivalent for any digraph D :*

- (i) $|D|$ is compact;
- (ii) every closed set of vertices is finite;
- (iii) D has no U -rank for any infinite vertex set U ;
- (iv) for every infinite set U of vertices there is a necklace attached to U ;
- (v) D is solid.

Proof. (i)→(ii): If $U \subseteq V(D)$ is closed and infinite, then any open cover that consists of $|D| - U$ and pairwise disjoint open neighbourhoods for the vertices in U has no finite subcover.

(ii)→(iii): Suppose that there is an infinite vertex set U for which D has a U -rank α . We may choose U so that α is minimal. Let $X \in \mathcal{X}(D)$ witness that D has U -rank α . By the choice of U , all the strong components of $D - X$ contain only finitely many vertices of U . Hence, U is closed in $|D|$, as every point in $|D|$ has an open neighbourhood that avoids U .

(iii)→(iv) This is immediate by the necklace lemma.

(iv)→(v) If D is not solid, say $D - X$ has infinitely many strong components for $X \in \mathcal{X}(D)$; then let U be a vertex set that contains exactly one vertex of every strong component of $D - X$. Clearly, there is no necklace attached to U .

(v)→(i) Theorem 9.1. □

9.3. The space $|D|$ as an inverse limit

In this section we show that the space $|D|$ for a solid digraph D can be obtained as an inverse limit of finite contraction minors of D , Theorem 9.2. We begin by defining an inverse system of finite digraphs for any digraph. Then, we show that every digraph embeds in the inverse limit of its inverse system. This gives a compactification for arbitrary digraphs, Theorem 9.3.1.

Let us introduce an inverse system for a given digraph D . For this, we define a directed partially ordered set (\mathcal{P}, \leq) as follows. We call a finite partition P of $V(D)$ *admissible* if any two partition classes of P can be separated in D by a finite

vertex set. We denote by $\mathcal{P} := \mathcal{P}(D)$ the set of all the admissible partitions of D . For any two partitions P_1 and P_2 of the vertex set of D we write $P_1 \leq P_2$ and say that P_2 is *finer* than P_1 if every partition class of P_2 is a subset of a partition class of P_1 .

We claim that the set of admissible partitions is a directed partially ordered set. Indeed, the relation \leq is easily seen to be a partial order on the set of all the partitions of $V(D)$. In particular, it restricts to a partial order on the set of all the admissible partitions. To see that \mathcal{P} is directed, let $P, P' \in \mathcal{P}$ be admissible partitions, and let P'' be the partition that consists of all the non-empty sets of the form $p \cap p'$ with $p \in P$ and $p' \in P'$. Clearly, P'' is finer than both P and P' . To see that P'' is admissible, let any two distinct partition classes of P'' be given, say $p_1 \cap p'_1$ and $p_2 \cap p'_2$ with $p_1, p_2 \in P$ and $p'_1, p'_2 \in P'$. As these partition classes are distinct, we have $p_1 \neq p_2$ or $p'_1 \neq p'_2$, say $p_1 \neq p_2$. Since P is admissible D has a finite vertex set that separates p_1 and p_2 , which in particular separates $p_1 \cap p'_1$ and $p_2 \cap p'_2$.

Let us proceed by defining the topological spaces associated with the admissible partitions of D . Every admissible partition P of D gives rise to a finite (multi-) digraph D/P by contracting each partition class and replacing all the edges between two partition classes by a single edge whenever there are infinitely many. Formally, declare P to be the vertex set of D/P . Given distinct partition classes $p_1, p_2 \in P$ we define an edge (e, p_1, p_2) of D/P for every edge e in D from p_1 to p_2 if there are finitely many such edges. And if there are infinitely many edges from p_1 to p_2 we just define a single edge $(p_1 p_2, p_1, p_2)$. We call the latter type of edges *quotient edges*. Endowing D/P with the 1-complex topology turns it into a compact Hausdorff space, i.e., basic open sets are uniform ε stars around vertices and open subintervals of edges. In other words, D/P is defined as our topological space from the previous section (for finite D) with the only difference that multi-edges are taken into account. We will usually not distinguish between the finite (multi-) digraphs D/P and the topological space $|D/P|$. Now, let us turn to the final ingredient of our inverse system for D : bonding maps. We define for every two distinct admissible partitions $P \leq P'$ of D a bonding map $f_{P'P}: D/P' \rightarrow D/P$ as follows. Vertices $p' \in P'$ of D/P' get mapped to the unique vertex $p \in P$ of D/P with $p' \subseteq p$. Edges get mapped according to their endvertices: For edges (e', p'_1, p'_2) of D/P' we consider two cases: First, if $p'_1, p'_2 \subseteq p$ for a partition class

$p \in P$, then (e', p'_1, p'_2) gets mapped to the vertex p of D/P . Second, if $p'_1 \subseteq p_1$ and $p'_2 \subseteq p_2$ for two distinct partition classes $p_1, p_2 \in P$, then there is at least one edge from p_1 to p_2 in D/P . If (e', p'_1, p'_2) is a quotient edge in D/P' , then also $(p_1 p_2, p_1, p_2)$ is a quotient edge in D/P and we map (e', p'_1, p'_2) to $(p_1 p_2, p_1, p_2)$. If (e', p'_1, p'_2) is not a quotient edge in D/P' and $(p_1 p_2, p_1, p_2)$ is a quotient edge in D/P , we map (e', p'_1, p'_2) to $(p_1 p_2, p_1, p_2)$. Finally, if (e', p'_1, p'_2) is not a quotient edge and there is no quotient edge between p_1 and p_2 , then (e', p_1, p_2) is an edge in D/P and we map (e', p'_1, p'_2) to (e', p_1, p_2) .

It is straightforward to check that the $f_{P'P}$ are continuous and that we have $f_{P''P} = f_{P'P} \circ f_{P''P'}$ for all admissible partitions $P \leq P' \leq P''$. The bonding maps turn $\{D/P, f_{P'P}, \mathcal{P}\}$ in an inverses system and we denote its inverse limit by $\varprojlim_{P \in \mathcal{P}} (D/P)$. Note that $\varprojlim_{P \in \mathcal{P}} (D/P)$ is non-empty, as the collection of points that consists for every $P \in \mathcal{P}$ of the vertex of D/P that contains a fixed vertex of D is an element of $\varprojlim_{P \in \mathcal{P}} (D/P)$.

Our next goal is to find an embedding from D to the inverse limit $\varprojlim_{P \in \mathcal{P}} (D/P)$ witnessing that the inverse limit is a Hausdorff compactification of D . We obtain this embedding by defining continuous maps $\varphi_P: D \rightarrow D/P$, one for every admissible partition $P \in \mathcal{P}$. Once the φ_P are defined, the universal property of the inverse limit gives rise to the desired embedding.

So let us define φ_P for a given admissible partition $P \in \mathcal{P}$. For a vertex v of D let $\varphi_P(v)$ be the partition class of P that contains v . For an inner point z of an edge vw of D , consider the partition classes that contain v and w , respectively. If these coincide, map z to the partition class that contains v and w . Otherwise the partition classes that contain v respectively w differ and there is an edge e from $\varphi_P(v)$ to $\varphi_P(w)$ in D/P , which is either a copy of vw considered as an edge of D/P or a quotient edge. Map z to its corresponding point on e . It is straightforward to see that φ_P is continuous for every $P \in \mathcal{P}$ and that $f_{P'P} \circ \varphi_{P'} = \varphi_P$ for every admissible partitions $P \leq P'$, i.e., the collection of maps $\{\varphi_P \mid P \in \mathcal{P}\}$ is compatible. Hence, by the universal property of the inverse limit, the φ_P induce a map $\Phi: D \rightarrow \varprojlim_{P \in \mathcal{P}} (D/P)$ with $\varphi_P = f_P \circ \Phi$, for every $P \in \mathcal{P}$.

Theorem 9.3.1. *For every digraph D the space $\varprojlim(D/P)_{P \in \mathcal{P}}$ is a Hausdorff compactification of D , in particular, the map*

$$\Phi: D \rightarrow \varprojlim(D/P)_{P \in \mathcal{P}}$$

is an embedding and its image is dense in $\varprojlim(D/P)_{P \in \mathcal{P}}$.

Proof. We have to show that $\varprojlim(D/P)_{P \in \mathcal{P}}$ is compact and Hausdorff, that the image of Φ is dense in $\varprojlim(D/P)_{P \in \mathcal{P}}$ and that Φ is an embedding i.e., it is a homeomorphism onto its image. The inverse limit $\varprojlim(D/P)_{P \in \mathcal{P}}$ is compact and Hausdorff because all the topological spaces D/P are compact and Hausdorff. As every φ_P is surjective the image of Φ is dense in $\varprojlim(D/P)_{P \in \mathcal{P}}$, by Lemma 9.1.3. In order to show that Φ is a homeomorphism onto its image, note first that the collection of maps $\{\varphi_P \mid P \in \mathcal{P}\}$ is eventually injective. Hence Φ is injective by the lifting lemma.

It remains to show that the inverse of Φ is continuous, for which we equivalently show that Φ is open, i.e., the image under Φ of open sets in D is open in $\Phi(D)$. It suffices to show this on a base for the open sets in D . We prove that Φ is open for the base \mathcal{B} given by the open uniform stars around vertices and the open subintervals of edges. Our goal is to find for every $B \in \mathcal{B}$ an open set O such that $\Phi(B) = O \cap \Phi(D)$. First consider the case where $B = B_\varepsilon(v)$ is an open ball of radius ε around a vertex v . Then let P be any admissible partition in which $\{v\}$ is a singleton partition class. In D/P we have that $\varphi_P(B)$ is an open ball of radius ε around the vertex $\varphi_P(v)$. We claim that $O := f_P^{-1}(\varphi_P(B))$ is the desired open set. Clearly, $\Phi(B) \subseteq O \cap \Phi(D)$, we prove the converse inclusion. For this let $x \in O \cap \Phi(D)$ be given. Let $d \in D$ be the preimage of x under Φ . We have to show that $d \in B$. If $d \notin B$, then $\varphi_P(d) \notin \varphi_P(B)$, contradicting the fact that $x \in O$.

Second let B be an open subinterval of an edge e in D , say with end points v and w . Then let P be any admissible partition in which $\{v\}$ and $\{w\}$ are singleton partition classes. A similar argument as above shows that $O := f_P^{-1}(\varphi_P(B))$ is as desired. \square

For example consider the directed ray R . Note first that every admissible partition of R has exactly one infinite partition class. One can check that $\varprojlim(R/P)_{P \in \mathcal{P}}$ is homeomorphic to the space where one adds a single point ω at infinity to R and where a neighbourhood base of ω is given by the tails of R together with ω .

We now extend the maps φ_P to maps $\hat{\varphi}_P: |D| \rightarrow D/P$. For this we define how $\hat{\varphi}_P$ behaves on ends and on inner points of limit edges; the values of $\hat{\varphi}_P$ on D are then given by the values of φ_P on D . For an end ω of D all the rays that represent ω have a tail in the same partition class p of P . The reason for this is that any two partition classes of P can be separated by a finite vertex set. Here we map ω to p .

Now, consider an inner point z of a limit edge λ . Note that we have already defined the images of the two endpoints of λ . If these images coincide, then map z to the unique image of the endpoints of λ . Otherwise, Proposition 8.4.1 or Proposition 8.4.2 gives rise to a quotient edge λ' between the partition classes of the endpoints of λ . In this case we map z to the corresponding point on λ' . This completes the definition of $\hat{\varphi}_P$.

Lemma 9.3.2. *The map $\hat{\varphi}_P: |D| \rightarrow D/P$ is continuous for every $P \in \mathcal{P}$.*

Proof. In order to prove that $\hat{\varphi}_P$ is continuous, we show that the preimage of every open ball with radius ε around a vertex of D/P is open in $|D|$ and that the preimage of every open subinterval of an edge in D/P is open in $|D|$. As these open sets form a base of the topology of D/P , the map $\hat{\varphi}_P$ is continuous.

Consider an open ball $B_\varepsilon(p)$ of radius ε around a vertex $p \in P$ in D/P . To see that $\hat{\varphi}_P^{-1}(B_\varepsilon(p))$ is open in $|D|$ we will define for every $y \in \hat{\varphi}_P^{-1}(p)$ an open set O_y in $|D|$ such that $\hat{\varphi}_P(O_y) \subseteq B_\varepsilon(p)$; in other words, the union of the open sets O_y is included in $\hat{\varphi}_P^{-1}(B_\varepsilon(p))$. A closer look on the definition of the O_y will show that this latter inclusion is in fact an equality.

So let $y \in |D|$ with $\hat{\varphi}_P(y) = p$ be given. To begin, if y is a vertex of D let O_y be the open ball in $|D|$ of radius ε around y . If y is an inner point of an edge e of D , then the whole edge e is mapped to p and we choose O_y to be the interior of e . If y is an end or an inner point of a limit edge, we fix a finite vertex set X that separates p from every other partition class in P . Note, that a strong component of $D - X$ is either contained in p or is disjoint from p . If y is an end, let O_y be the basic open neighbourhood $\hat{C}_\varepsilon(X, y)$. Note that, by the choice of X , the strong component $C(X, y)$ is included in the partition class p . If y is an inner point of a limit edge λ , and X separates the endpoints of this limit edge, then let $O_y = \hat{E}_{\varepsilon', y}(X, \lambda)$ with $\varepsilon' < \varepsilon$ small enough to fit into λ , i.e., such that $B_{\varepsilon'}(y) \subseteq \lambda$ where the $B_{\varepsilon'}(y)$ is considered in the space $[0, 1]_\lambda$; otherwise let

$C(X, \omega)$ be the strong component of $D - X$ that contains both endpoints of λ and let $O_y = \hat{C}_\varepsilon(X, \omega)$.

Clearly, the union of the O_y is included in $\hat{\varphi}_P^{-1}(B_\varepsilon(p))$. Moreover, for every $z \neq p$ in $B_\varepsilon(p)$ the set $\hat{\varphi}_P^{-1}(z)$ is a set of inner points of edges (possibly limit edges). Each such inner point is contained in an ε -neighbourhood of the endpoint e that is mapped to p , for e the edge that contains the inner point. Hence each of these inner points is contained in at least one of the open sets O_y .

Now, consider an open subinterval $B_\varepsilon(z)$ of radius ε around z for an inner point z of an edge (e, p, p') of D/P . If (e, p, p') is not a quotient edge of D/P , then e is an edge of D and the preimage of $B_\varepsilon(z)$ is an open subinterval of e considered as an edge in $|D|$, namely around the point $\hat{\varphi}_P^{-1}(z)$ of radius ε . So suppose that (e, p, p') is a quotient edge. We will find for every point $y \in \hat{\varphi}_P^{-1}(z)$ an open neighbourhood O_y of y in $|D|$ with $O_y \subseteq \hat{\varphi}_P^{-1}(B_\varepsilon(z))$. A similar argument for every point in $B_\varepsilon(z)$ shows that $\hat{\varphi}_P^{-1}(B_\varepsilon(z))$ is the union of open subsets in $|D|$.

Note that all the points in $\hat{\varphi}_P^{-1}(z)$ are inner points of edges in $|D|$ (possibly limit edges). Let $y \in \hat{\varphi}_P^{-1}(z)$ be given. First, if y is an inner point of an edge of D , then let O_y be the open subinterval of radius ε around y . Second, suppose that y is an inner point of a limit edge whose end points are ends, say ω and η , and with $\hat{\varphi}_P(\omega) = p$ and $\hat{\varphi}_P(\eta) = p'$. Fix finite vertex sets $X_p, X_{p'} \subseteq V(D)$ that separate p respectively p' from every other partition class in P . Note, that $X_p \cup X_{p'}$ separates ω and η . Now, every edge that is contained in or lives in $E(X_p \cup X_{p'}, \omega\eta)$ is mapped to (e, p, p') ; thus the basic open neighbourhood $\hat{E}_{\varepsilon, y}(X_p \cup X_{p'}, \omega\eta)$ is mapped to $B_\varepsilon(z)$. Finally, suppose that y is an inner point of a limit edge λ between a vertex v and an end ω , say with $\hat{\varphi}_P(v) = p$ and $\hat{\varphi}_P(\omega) = p'$; the other case is analogue. Let $X_{p'}$ be a finite vertex set of D that separates the partition class p' from every other partition class in P . Then $\hat{E}_{\varepsilon, y}(X_{p'} \cup \{v\}, \lambda)$ is mapped to $B_\varepsilon(z)$. \square

We are now ready to prove the main result of this section:

Theorem 9.2. *Let D be a solid digraph. The map induced by the $\hat{\varphi}_P: |D| \rightarrow D/P$*

$$\hat{\Phi}: |D| \rightarrow \varinjlim_{P \in \mathcal{P}} (D/P)$$

is a homeomorphism.

Proof. It is straightforward to show that the $\hat{\varphi}_P$ are compatible. Let us show that the collection of maps $\{\hat{\varphi}_P \mid P \in \mathcal{P}\}$ is eventually injective, that is to say for

every two points $x, y \in |D|$ there is a $P \in \mathcal{P}$ such that $\hat{\varphi}_P(x) \neq \hat{\varphi}_P(y)$. Such an admissible partition is easily defined if at least one of the points x and y lies in D . So suppose x and y are ends or inner points of limit edges. If x and y are both ends choose an $X \in \mathcal{X}(D)$ that separates x and y . Then the admissible partition P_X given by the strong components of $D - X$ and all the vertices in X as singletons, is the desired partition. Similarly, if x is an end and y is an inner point of a limit edge of the form $\omega\eta$ for two ends of D , then choose an $X \in \mathcal{X}(D)$ that separates all the ends in $\{x, \omega, \eta\}$ simultaneously. Again the admissible partition given by the strong components of $D - X$ and all the vertices in X as singletons is as desired. The other cases are analogue and we leave the details to the reader.

By the lifting lemma and Lemma 9.3.2 the $\hat{\varphi}_P$ induce a continuous injective map $\hat{\Phi}: |D| \rightarrow \varprojlim_{P \in \mathcal{P}} (D/P)$. By Lemma 9.1.3 we have that the image of the map $\hat{\Phi}$ is dense in $\varprojlim_{P \in \mathcal{P}} (D/P)$. Moreover, as D is solid, we have that $|D|$ is compact by Theorem 9.1 so the image of $\hat{\Phi}$ is closed; hence it is all of $\varprojlim_{P \in \mathcal{P}} (D/P)$. The statement now follows from Lemma 9.1.5. \square

Corollary 9.3.3. *For a solid digraph D the topology DTOP is the finest compact topology on $|D|$.*

Proof. Let D be solid and \mathcal{T} a compact topology on $|D|$ that is finer than DTOP . Our goal is to show that $\mathcal{T} = \text{DTOP}$. For this it suffices to find a homeomorphism from $(|D|, \mathcal{T})$ to $\varprojlim_{P \in \mathcal{P}} (D/P)$ since the latter is homeomorphic $(|D|, \text{DTOP})$. Consider the collection of maps $\{\hat{\varphi}_P \mid P \in \mathcal{P}\}$ from the text before Lemma 9.3.2. As \mathcal{T} is finer than DTOP these maps are continuous for $(|D|, \mathcal{T})$. Clearly, the collection of these maps is still compatible and eventually injective. Therefore, we obtain an injective continuous map $\Phi': (|D|, \mathcal{T}) \rightarrow \varprojlim_{P \in \mathcal{P}} (D/P)$, by Lemma 9.1.2. As the $\hat{\varphi}_P$ are still surjective, we have by Lemma 9.1.3 that the image of Φ' is dense. As $(|D|, \mathcal{T})$ is compact the image of Φ' is also closed and it follows that Φ' is surjective. Using Lemma 9.1.5 we conclude that the map Φ' with domain $(|D|, \mathcal{T})$ is a homeomorphism. \square

In the proof of Theorem 9.2 we used those admissible partitions that arise by deleting a finite vertex set from a solid digraph to ensure that the map Φ that is induced by the $\hat{\varphi}_P$ is injective. Next, we show that these admissible partitions capture the whole inverse system for a solid digraph.

To make this formal, let D be a solid digraph and $X \in \mathcal{X}(D)$. We denote by P_X the admissible partition where each vertex in X is a singleton partition class and the other partition classes consist of the strong components of $D - X$. We claim that $\mathcal{P}_X := \{P_X \mid X \in \mathcal{X}(D)\}$ is cofinal in the set of admissible partitions of D , that is for every admissible partition P there is an X such that $P \leq P_X$. Indeed, given $P \in \mathcal{P}$ we have $P \leq P_X$ for any finite set $X \in \mathcal{X}(D)$ that separates any two partition classes in P .

Now, $\{D/P_X, f_{P_X P_{X'}}, \mathcal{P}_X\}$ is an inverse system by itself and by Lemma 9.1.4 we have that

$$\varprojlim_{P \in \mathcal{P}} (D/P) \cong \varprojlim_{X \in \mathcal{X}} (D/P_X).$$

If D is countable one can simplify the directed system even further: Fix an enumeration v_0, v_1, \dots of the vertex set of D and write X_n for the set of the first n vertices. Then the set of all the P_{X_n} is cofinal in \mathcal{P}_X and therefore it is also cofinal in the set of all the admissible partitions of D .

Corollary 9.3.4. *Let D be a countable solid digraph and let X_n consist of the first n vertices of D with regard to a any fixed enumeration of $V(D)$. Then $|D| \cong \varprojlim_{n \in \mathbb{N}} (D/P_{X_n})$. \square*

9.4. Applications

In this last section we prove two statements about finite digraphs that naturally generalise to the space $|D|$, but do not generalise verbatim to infinite digraphs, Theorem 9.3 and Theorem 9.4. We begin this section by introducing all the definitions needed. We then provide an important tool that describes how (topological) paths in $|D|$ can pass through cuts in D , the directed jumping arc lemma. Finally, we prove our two main results of this section, Theorem 9.3 and Theorem 9.4.

A continuous function $\alpha: [0, 1] \rightarrow |D|$ is called a *local homeomorphism on the edges of $|D|$* if for every $x \in [0, 1]$ that is mapped to an inner point of an edge $e \in E \cup \Lambda$, there is a neighbourhood (a, b) of x such that α restricts on (a, b) to a homeomorphism to the interior of e , i.e., $\alpha \upharpoonright (a, b) \cong \mathring{e}$. Note that by the continuity of α any such homeomorphism $\alpha \upharpoonright (a, b)$ extends to a homeomorphism $\alpha \upharpoonright [a, b] \cong e$. If in addition α respects the orientation of the edges in $|D|$, that

is if $[a, b] \subseteq [0, 1]$ is mapped to an edge $e \in E \cup \Lambda$ we have that $x \leq y$ implies $\alpha(x) \leq \alpha(y)$ for all the $x, y \in [a, b]$, then we call α a *directed topological path* in $|D|$. (Here $\alpha(x) \leq \alpha(y)$ refers to \leq in $[0, 1]_e$.)

We think of directed topological paths in $|D|$ as generalised directed walks in D . Here the edges of a directed walk in D are directed along the walk. Indeed, every directed walk in D defines via a suitable parametrisation a directed topological path in $|D|$. If the image of α contains a vertex or an end x we simply say that α contains x . We say that a directed topological path α *traverses an edge* $e \in E \cup \Lambda$ of $|D|$ if α restricts on a subinterval of $[0, 1]$ to a homeomorphism on e . The points $\alpha(0)$ and $\alpha(1)$ are called the *endpoints* of α and we say that α *connects* $\alpha(0)$ to $\alpha(1)$. A directed topological path whose endpoints coincide is *closed*.

Next, let us gain some understanding of how directed topological paths in $|D|$ can pass through cuts of D , see also the jumping arc lemma [20, Lemma 8.6.3].

Lemma 9.4.1 (Directed Jumping Arc Lemma). *Let D be any digraph and let $\{V_1, V_2\}$ be any bipartition of $V(D)$.*

- (i) If $\overline{V_1} \cap \overline{V_2} = \emptyset$, then every directed topological path in $|D|$ from V_1 to V_2 traverses an edge of $|D|$ with tail in $\overline{V_1}$ and head in $\overline{V_2}$.
- (ii) If $\overline{V_1} \cap \overline{V_2} \neq \emptyset$, there will be a directed topological path in $|D|$ from V_1 to V_2 that traverses none of the edges between $\overline{V_1}$ and $\overline{V_2}$ if both $D[V_1]$ and $D[V_2]$ are solid.

Proof. (i) Suppose that $\overline{V_1} \cap \overline{V_2}$ is empty. Then every end of D is either contained in $\overline{V_1}$ or $\overline{V_2}$. First, we show that every edge of $|D|$ that has both of its endpoints in $\overline{D[V_i]}$ is contained in $\overline{D[V_i]}$, for $i = 1, 2$. For edges of D this is trivial. So consider a limit edge λ with both endpoints in $\overline{D[V_i]}$. All but finitely many vertices of a subdigraph obtained by Proposition 8.4.1 or Proposition 8.4.2 applied to λ are contained in $D[V_i]$, otherwise this gives an end in $\overline{V_1} \cap \overline{V_2}$. Consequently, for every inner point $z \in \lambda$ there is a sequence of inner points of edges in $\overline{D[V_i]}$ that converge to z , giving $z \in \overline{D[V_i]}$.

From this first observation, we now know that $|D| \setminus (\overline{D[V_1]} \cup \overline{D[V_2]})$ consists only of inner points of edges (possibly limit edges) between $\overline{V_1}$ and $\overline{V_2}$. Now, consider a directed topological path α that connects a point in V_1 to a point in V_2 . As $[0, 1]$ is connected and α is continuous there is a point $x \in [0, 1]$ with

$\alpha(x) \in |D| \setminus (\overline{D[V_1]} \cup \overline{D[V_2]})$. Hence the preimage of $|D| \setminus (\overline{D[V_1]} \cup \overline{D[V_2]})$ is non-empty and a union of pairwise disjoint intervals (a, b) each of which is mapped homeomorphic to an open edge between $\overline{V_1}$ and $\overline{V_2}$. The usual relation \leq on the reals defines a linear order on these intervals. Among these intervals, choose (a, b) minimal. That this is possible can be seen as follows: If not, we find a strictly decreasing sequence $(a_0, b_0) \geq (a_1, b_1) \geq \dots$ of intervals with $\alpha \upharpoonright [a_i, b_i] \cong e_i$ for some edges e_i between $\overline{V_1}$ and $\overline{V_2}$. Then $(a_i)_{i \in \mathbb{N}}$ and $(b_i)_{i \in \mathbb{N}}$ converge to some $c \in [0, 1]$ and using that α is continuous, we get $\alpha(c) \in \overline{V_1} \cap \overline{V_2}$, a contradiction.

We claim that the image of (a, b) under α is an edge from $\overline{V_1}$ to $\overline{V_2}$. To see this, it suffices to show that $\alpha(a) \in \overline{V_1}$. So suppose for a contradiction that $\alpha(a) \in \overline{V_2}$. Then $\alpha \upharpoonright [0, a]$ gives a directed topological path from $\overline{V_1}$ to $\overline{V_2}$. By a similar argument as above there is a point in $[0, a]$ mapped to an edge between $\overline{V_1}$ and $\overline{V_2}$, contradicting the choice of (a, b) .

(ii) First note that no inner point of a limit edge is a limit point of a set of vertices. Hence D has at least one end that is contained in both the closure of V_1 and V_2 . By Lemma 9.2.1 we find a necklace $N \subseteq D$ attached to $\{V_1, V_2\}$. Let ω be the end that is represented by N . Apply Corollary 9.2.2 to the solid digraph $D[V_1]$ and the infinite set $U_1 := V_1 \cap V(N)$ in order to obtain a necklace N_1 attached to U_1 . Let U_2 consist of all the vertices in V_2 that are contained in those beads of N that intersect a bead of N_1 . Apply Corollary 9.2.2 to the solid digraph $D[V_2]$ and the infinite set U_2 in order to obtain a necklace N_2 attached to U_2 . Note that both necklaces N_1 and N_2 represent ω . A ray in N_1 together with a reverse ray in N_2 defines a directed topological path that is as desired. \square

Now, let us turn to our applications. A finite digraph is called *Eulerian* if there is a closed directed walk that contains every edge exactly once. A *cut* of a digraph D is an ordered pair (V_1, V_2) of non-empty sets $V_1, V_2 \subseteq V(D)$ such that $V_1 \cup V_2 = V(D)$ and $V_1 \cap V_2 = \emptyset$. The sets V_1 and V_2 are the *sides* of the cut, and its *size* is the cardinality of the set of edges from V_1 to V_2 . We call a cut (V_1, V_2) *balanced* if its size equals that of (V_2, V_1) . An *unbalanced cut* is a cut that is not balanced. It is well known that a finite digraph (with a connected underlying graph) is Eulerian if and only if all of its cuts are balanced.

A closed directed topological path α that traverses every edge of $|D|$ exactly once is called *Euler tour*, i.e., for every edge e of $|D|$ there is exactly one subinterval

of $[0, 1]$ that is mapped homeomorphic to e via α . If $|D|$ has an Euler tour we call $|D|$ *Eulerian*. There are two obstructions for digraph D to be Eulerian: one is a vertex of infinite degree and the other one is an unbalanced cut. A digraph is *locally finite* if all its vertices have finite in- and out-degree. Theorem 9.3 states that there are no further obstructions. We need one more lemma for its proof:

Lemma 9.4.2. *Let D be a digraph with a connected underlying graph. If D is locally finite and every finite cut of D is balanced, then D is solid.*

Proof. Suppose for a contradiction that D is not solid and fix a finite vertex set $X \subseteq V(D)$ such that $D - X$ has infinitely many strong components. Our goal is to find a finite unbalanced cut of D . We may view the strong components of $D - X$ partially ordered by $C_1 \leq C_2$ if there is a path in $D - X$ from C_1 to C_2 . We first note that any strong component C of $D - X$ receives and sends out only finitely many edges in $D - X$. Indeed, if C sends out infinitely many edges, then (V_1, V_2) is a finite unbalanced cut, where V_1 is the union of all the strong components strictly greater than C and where $V_2 := V(D) \setminus V_1$. A similar argument shows that C receives only finitely many edges. Now, the (multi-)digraph D' obtained from D by contracting all the strong components of $D - X$ is locally finite. Note that also every finite cut of D' is balanced.

Now, D' is also strongly connected. Indeed, if there is a vertex $v \in V(D')$ that cannot reach all the other vertices, then (V_1, V_2) is a finite unbalanced cut of D' , where V_1 is the set of vertices in $V(D')$ that can be reached from v and where $V_2 := V(D') \setminus V_1$ (here we use that the graph underlying D is connected). Hence we may apply the directed star-comb lemma in D' to $V(D')$. As D' is locally finite, the return is a comb and a reverse comb sharing their attachment sets; we may assume that both avoid X . Let R be the spine of the comb and R' the spine of the reverse comb. Let V_2 be the set of all the vertices in $D' - X$ that can be reached from R' in $D' - X$ and $V_1 := V(D') \setminus V_2$.

As $D' - X$ is acyclic we have that the vertex set of R is included in V_1 . But then (V_1, V_2) is a finite unbalanced cut, which in turn gives rise to a finite unbalanced cut of D . \square

Theorem 9.3. *For a digraph D with a connected underlying graph the following assertions are equivalent:*

- (i) $|D|$ is Eulerian;
- (ii) D is locally finite and every finite cut of D is balanced.

Proof. For the forward implication (i) \rightarrow (ii) suppose that D has an Euler tour α . Using the directed jumping arc lemma it is straightforward to show that D has only balanced cuts. Let us show that D needs to be locally finite for α to be continuous. Suppose for a contradiction there is a $v \in V(D)$ with infinitely many edges e_0, e_1, \dots with head v ; the case where v is the tail of infinitely many edges is analogue. Let $(a_i, b_i) \subseteq [0, 1]$ the subinterval that is mapped homeomorphically by α to e_i . As the unit interval is compact, the sequence of the a_i has a convergent subsequence $(a_{i_n})_{n \in \mathbb{N}}$ and we write x for the limit point of this subsequence. Now, the subsequence $(b_{i_n})_{n \in \mathbb{N}}$ of the b_i forms a convergent subsequence, too, with limit point x . As $\alpha(b_{i_n}) = v$ for all the $n \in \mathbb{N}$ we have $\alpha(x) = v$, by the continuity of α ; but $\alpha(a_{i_n})$ is a sequence of neighbours of v which does not converge in $|D|$ to v , a contradiction.

For the backward implication (ii) \rightarrow (i) let us first show that $|D|$ contains no limit edges. As D is locally finite there is no limit edge between a vertex and an end. So suppose for a contradiction there is a limit edge $\omega\eta$ between two ends of D . Fix a finite vertex set X that separates ω and η . We may view the strong components of $D - X$ partially ordered by $C_1 \leq C_2$ if there is a path in $D - X$ from C_1 to C_2 . Let V_1 consist of all the vertices in strong components of $D - X$ that are strictly smaller than $C(X, \eta)$ and let $V_2 := V(D) \setminus V_1$. Then (V_1, V_2) is an unbalanced cut: On the one hand there are infinitely many edges from V_1 to V_2 because there are infinitely many from $C(X, \omega)$ to $C(X, \eta)$. On the other hand, there are only finitely many edges from V_2 to V_1 by our assumption that D is locally finite.

Let us now find an Euler tour for $|D|$. By Lemma 9.4.2 the digraph D is solid. As it is locally finite and its underlying graph is connected $V(D)$ is countable. Choose an enumeration of $V(D)$ and let X_n denote the set of the first n vertices. Then D/P_{X_n} contains no quotient edge and every cut of D/P_{X_n} is balanced. As the statement of Theorem 9.3 holds for finite digraphs, we have that D/P_{X_n} is Eulerian. Moreover, as D/P_{X_n} is a finite digraph there are only finitely many

(combinatorial) Euler tours of D/P_{X_n} . By König's infinity lemma there is a consistent choice of one Euler tour for every D/P_{X_n} . Now, take a parametrisation $\alpha_n: [0, 1] \rightarrow D/P_{X_n}$ of the Euler tour chosen for D/P_{X_n} such that the α_n are compatible. Using Theorem 9.2 it is straightforward to check that the universal property of the inverse limit gives an Euler tour for $|D|$. \square

It is well known that a finite digraph is strongly connected if and only if it has a directed closed walk that contains all its vertices. Clearly, the statement does not generalise verbatim to infinite digraphs nor does a spanning directed (double)-ray ensure the digraph to be strongly connected. Moreover, the statement does not hold if one adds the ends of the underlying undirected graph:




Figure 9.4.1.: A solid digraph (every undirected edge in the figure stands for two directed edges in opposite directions) that is not strongly connected. Adding the one end of the underlying undirected graph makes it possible to find a closed directed topological path that contains all the vertices.

Adding the ends and limit edges of the digraph turns out to be the right setting for the statement to generalise:

Theorem 9.4. *For a countable solid digraph D the following assertions are equivalent:*

- (i) D is strongly connected;
- (ii) there is a closed topological path in $|D|$ that contains all the vertices of D .

Proof. For the forward implication (i) \rightarrow (ii) fix an enumeration v_1, v_2, \dots of $V(D)$ and denote by X_n the set of the first n vertices. We will recursively define a sequence of walks W_1, W_2, \dots such that W_n is a directed closed walk of D/P_{X_n} that contains all the vertices of D/P_{X_n} and such that the projection of W_n to $D/P_{X_{n-1}}$ is exactly W_{n-1} .

Once the W_n are defined, it is not hard to find parametrisations α_n of each W_n such that $f_{X_n X_{n-1}} \circ \alpha_n = \alpha_{n-1}$. Then the universal property of the inverse limit together with Corollary 9.3.4 and Theorem 9.2 gives the desired closed directed topological path in $|D|$.

To begin, let W_1 be an arbitrary closed walk in $D - X_1$ that contains all its vertices. Now suppose that $n > 1$ and that W_{n-1} has already been defined. Let C be the strong component of $D - X_{n-1}$ that contains v_n . Note that the strong components of $D - X_n$ are exactly the strong components of $D - X_{n-1}$ that are distinct from C together with all the strong components of $C - v_n$. As C is strongly connected the digraph C/P_{v_n} is strongly connected, as well. We now extend W_{n-1} to W_n by plugging in a directed walk that contains all the vertices of C/P_{v_n} each time W_{n-1} meets C . Formally, we fix for every edge e_i of W_{n-1} with one of its endvertices in C an edge f_i in D/P_{X_n} that is mapped to e_i by $f_{X_n X_{n-1}}$. For every occurrence of C in W_{n-1} there are consecutive edges e_i and e_{i+1} in W_{n-1} such that C is the head of e_i and the tail of e_{i+1} . Now, fix a directed walk Q_i in C/P_{v_n} from the head of f_i to the tail of f_{i+1} that contains all the vertices of C/P_{v_n} . We define W_n by replacing any such consecutive edges e_i and e_{i+1} in W_{n-1} by $f_i Q_i f_{i+1}$.

We prove the implication (ii)→(i) via contraposition. Suppose that D is not strongly connected. Then there are vertices $v, w \in V(D)$ so that there is no path from v to w . Let V_1 consist of all the vertices that can be reached from v and let $V_2 := V(D) \setminus V_1$. As $w \notin V_1$, we have $V_2 \neq \emptyset$. Moreover, the edges between V_1 and V_2 form a cut with no edge from V_1 to V_2 (in particular no limit edge). Hence the intersection $\overline{V_1} \cap \overline{V_2}$ is empty. By the directed jumping arc lemma there is no directed topological path in $|D|$ from v to w . We conclude that there is no closed directed topological path in $|D|$ that contains all the vertices of D . \square

10. Normal arborescences

Depth-first search trees are a standard tool in finite graph and digraph theory. These trees arise from an algorithm on a graph or digraph called *depth-first search*. Starting from a fixed vertex, the ‘root’, the algorithm moves along the edges, going to a vertex not visited yet whenever this is possible, and going back otherwise. Depth-first search stops when all vertices have been visited, and the trees defined by the traversed edges are called *depth-first search* trees.

For connected finite graphs, the depth-first search trees are precisely the normal spanning trees. Here, a rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T . (A T -path in G is a non-trivial path that meets T exactly in its endvertices.) Normal spanning trees generalise depth-first search trees, since they are also defined for infinite graphs; they are perhaps the single most important structural tool in infinite graph theory [20].

In this chapter we introduce and study normal spanning arborescences. These are generalisations of depth-first search trees to infinite digraphs that promise to be as powerful for a structural analysis of digraphs as normal spanning trees are for graphs, both from a combinatorial and a topological point of view.

An *arborescence* is a rooted oriented tree T that contains for every vertex $v \in V(T)$ a directed path from the root to v . The vertices of any arborescence are partially ordered as $v \leq_T w$ if T contains a directed path from v to w . We write $[v]_T$ for the up-closure of v in T .

Consider any finite digraph D together with a spanning depth-first search tree $T \subseteq D$. If vw is an edge of D between \leq_T -incomparable vertices of T , then w is visited at an earlier stage of the depth-first search than v .¹ Together with all

¹Indeed, if v was visited before w , the algorithm would have traversed the edge vw rather than backtracking from v , which it must have done since v and w are incomparable. Note that all the visits to v happen while the algorithm searches $[v]_T$, and likewise for w , so visiting ‘before’ and ‘after’ are well-defined for incomparable vertices.




Figure 10.0.1.: A depth-first search arborescence visiting vertices from right to left.

such edges, T forms an acyclic subdigraph of D [16].²

Let us use this property of depth-first search trees in finite digraphs as the definition of our infinite analogue, i.e., as the defining property for ‘normal’ arborescences in infinite digraphs. More precisely, consider a (possibly infinite) digraph D and an arborescence $T \subseteq D$, not necessarily spanning. A T -path in D is a non-trivial directed path that meets T exactly in its endvertices. The *normal assistant* of T in D is the auxiliary digraph H that is obtained from T by adding an edge vw for every two \leq_T -incomparable vertices $v, w \in V(T)$ for which there is a T -path from $[v]_T$ to $[w]_T$ in D , regardless of whether D contains such an edge. The arborescence T is *normal* in D if the normal assistant of T in D is acyclic. It is straightforward to check that this indeed generalises depth-first search trees in that for finite D a spanning arborescence T of D is normal in D if and only if T defines a depth-first search tree; see Corollary 10.2.3.

One aspect of why normal spanning trees of infinite undirected graphs are so useful is that they are end-faithful. A spanning tree T of a graph G is *end-faithful* if the map that assigns to every end of T the end of G that contains it as a subset (of rays) is bijective, see [20]. Equivalently T is end-faithful if every end of G is represented by a unique ray in T that starts from a fixed root. Our first theorem will be an analogue of this for normal arborescences, so let us recall the definition

²Indeed, any cycle would, but cannot, lie in the up-closure of its first-visited vertex.

of ends of digraphs from Chapter 8.

A *directed ray* is an infinite directed path that has a first vertex (but no last vertex). The directed subrays of a directed ray are its *tails*. For the sake of readability we shall omit the word ‘directed’ in ‘directed path’ and ‘directed ray’ if there is no danger of confusion. We call a ray in a digraph *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. We call two solid rays in a digraph D *equivalent* if for every finite vertex set $X \subseteq V(D)$ they have a tail in the same strong component of $D - X$. The equivalence classes of this equivalence relation are the *ends* of D . For a finite vertex set $X \subseteq V(D)$ and an end ω of D we write $C(X, \omega)$ for the unique strong component of $D - X$ that contains a tail of every ray that represents ω ; the end ω is then said to *live* in that strong component. The set of ends of D is denoted by $\Omega(D)$.

Let $T \subseteq D$ be a spanning arborescence of a digraph D . We say that T is *end-faithful* if every end of D is represented by a unique ray in T starting from the root of T . (Note that, conversely, rays in T will only represent ends of D if they are solid in D .) Here is our first main result:

Theorem 10.1. *Every normal spanning arborescence of a digraph is end-faithful.*

In fact we will prove a localised version of this for normal arborescences in D that are not necessarily spanning.

The end space of any normal spanning tree T of an undirected graph G coincides with the end space of G , not only combinatorially but also topologically. Indeed, the map that assigns to every end of T the end of G that contains it as a subset is a homeomorphism between the end space of T and that of G , see [20]. Hence, in order to understand the end space of G one just needs to understand the simple structure of the tree T .

We also have an analogue of this for digraphs and their normal arborescences. To state this, let us recall the notion of limit edges of a digraph D .

For two distinct ends ω and η of D , we call the pair (ω, η) a *limit edge* from ω to η if D has an edge from $C(X, \omega)$ to $C(X, \eta)$ for every finite vertex set X for which ω and η live in distinct strong components of $D - X$. Similarly, for a vertex $v \in V(D)$ and an end ω of D we call the pair (v, ω) a *limit edge from v to ω* if D has an edge from v to $C(X, \omega)$ for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. And we call the pair (ω, v) a *limit edge from ω to v* if D has an

edge from $C(X, \omega)$ to v for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. The digraph D , its ends, and its limit edges together form a topological space $|D|$, in which the edges are copies of the real interval $[0, 1]$; see Chapter 9.

The *horizon* of a digraph D is the subspace of $|D|$ formed by the ends of D and all the limit edges between them. Arborescences do not themselves have ends or limit edges, but there is a natural way to endow an arborescence T in a digraph D with a meaningful horizon. The *solidification* of an arborescence $T \subseteq D$, or of its normal assistant H in D , is obtained from T or H , respectively, by adding all the edges wv with $vw \in E(T)$. Note that all the rays of T are solid in its solidification and thus represent ends there. Let us define the *horizon* of T as the horizon of the solidification of its normal assistant.

Recall that the digraphs D that are compactified by $|D|$ are precisely the *solid* ones, those such that $D - X$ has only finitely many strong components for every finite vertex set $X \subseteq V(D)$, see Chapter 9. Let T be a normal spanning arborescence of D , with root r , say. By Theorem 10.1, there exists a well-defined map ψ that sends every end ω of D to the end of the solidification \bar{T} of T represented by the unique ray $R \subseteq T$ starting from r that represents ω in D . This map ψ is clearly injective. If D is solid, every ray in T represents an end of D , so ψ is also surjective. Let ζ denote the map from the set of ends of \bar{T} to that of the solidification \bar{H} of the normal assistant H of T in D that assigns to every end of \bar{T} the end of \bar{H} that contains it as a subset (of rays). This is always bijective, see Lemma 10.4.1. Note that \bar{H} , unlike \bar{T} , can have limit edges. We say that T *reflects the horizon* of D if the map $\zeta \circ \psi: \Omega(D) \rightarrow \Omega(\bar{H})$ extends to a homeomorphism from the horizon of D to that of \bar{H} .

As our second main result we prove that normal spanning arborescences of solid digraphs reflect the horizon of the digraph they span:

Theorem 10.2. *Every normal spanning arborescence of a solid digraph reflects its horizon.*

Not every connected graph has a normal spanning tree; for example, uncountable complete graphs have none. Thus it is not surprising that there are also strongly connected digraphs without normal spanning arborescences—such as any digraph obtained from an uncountable complete graph by replacing every edge by its two orientations as separate directed edges.

Jung [43] characterised the connected graphs with a normal spanning tree in terms of dispersed sets. A set $U \subseteq V(G)$ of vertices of a graph G is *dispersed* if there is no comb in G with all its teeth in U . Recall that a *comb* is the union of a ray R with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R . The last vertices of those paths are the *teeth* of this comb, see [20]. Jung proved that a connected graph has a normal spanning tree if and only if its vertex set is a countable union of dispersed sets.

Translating this to digraphs, a *directed comb* is the union of a directed ray with infinitely many finite paths (possibly trivial) that have precisely their first vertex on R . Hence the underlying graph of a directed comb is an undirected comb. The *teeth* of a directed comb are the teeth of the underlying comb. We call a set $U \subseteq V(D)$ of vertices of a digraph D *dispersed* if there is no directed comb in D with all its teeth in U . For two vertices $v, w \in V(D)$, we say that v can reach w if D contains a path from v to w .

Theorem 10.3. *Let D be any digraph and suppose that $r \in V(D)$ can reach all the vertices of D . If $V(D)$ is a countable union of dispersed sets, then D has a normal spanning arborescence rooted in r .*

In fact we will prove a slightly stronger version of this where we show how to find a normal arborescence in D that contains a given set of vertices of D .

In an undirected graph, the levels of any normal spanning tree are dispersed, so the forward implication in Jung's characterisation is easy. Theorem 10.3 implies the harder backward implication when applied to the digraph obtained from the graph by replacing every edge by its two orientations as separate directed edges.

The easy forward implication in Jung's theorem does not have a directed analogue, since the converse implication in Theorem 10.3 may fail. For example consider the digraph obtained from a ray by adding a new vertex r and an edge from r to every vertex of the ray. In this digraph, the first level of the normal arborescence that consists of all edges at r is not dispersed. However, the converse of Theorem 10.3 does hold if the digraph D is solid.

This chapter is organised as follows. We provide further tools and terminology that we use throughout this chapter in Section 10.1. Then in Section 10.2 we introduce normal arborescences and provide some basic lemmas that we need for the proofs of our main results. In Section 10.3, we show that normal spanning

arborescences are end-faithful, Theorem 10.1. In Section 10.4, we prove that normal spanning arborescences reflect the horizon, Theorem 10.2. Finally, we prove our existence criterion for normal arborescences in digraphs, Theorem 10.3, in Section 10.5.

10.1. Tools and terminology

In this section we provide further tools and terminology that we use throughout this chapter. We also collect together some tools and terminology from Chapter 8 that become particularly important in this chapter.

Let D be any digraph. We write edges as ordered pairs (v, w) with $v, w \in V(D)$, and we usually write (v, w) simply as vw . The vertex v is the *tail* of vw and the vertex w its *head*. The *reverse* of an edge vw is the edge wv . More generally, the *reverse* of a digraph D is the digraph on $V(D)$ where we replace every edge of D by its reverse, i.e., the reverse of D has the edge set $\{vw \mid wv \in E(D)\}$. We write \overleftarrow{D} for the reverse of a digraph D .

The directed subrays of a ray are its *tails*. Call a ray *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. Two solid rays in D are *equivalent*, if they have a tail in the same strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. We call the equivalence classes of this relation the *ends* of D and we write $\Omega(D)$ for the set of ends of D .

Similarly, the reverse subrays of a reverse ray are its *tails*. We call a reverse ray *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. With a slight abuse of notation, we say that a reverse ray R *represents* an end ω if there is a solid ray R' in D that represents ω such that R and R' have a tail in the same strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$.

Given sets $A, B \subseteq V(D)$ of vertices a *path from A to B* , or A - B path is a path that meets A precisely in its first vertex and B precisely in its last vertex. We say that a vertex v can *reach* a vertex w in D and w can be *reached* from v in D if there is a v - w path in D . A non-trivial path P is an A -path for a set of vertices A if P has both its endvertices but none of its inner vertices in A . A set W of vertices is *strongly connected* in D if every vertex of W can reach every other vertex of W in $D[W]$.

We say that a digraph is acyclic if it contains no directed cycle as a subdigraph. The vertices of any acyclic digraph D are partially ordered by $v \leq_D w$ if D contains a path from v to w .

An *arborescence* is a rooted oriented tree that contains for every vertex $v \in V(T)$ a directed path from its root to v . Note that arborescences T are acyclic and that \leq_T coincides with the tree-order of the undirected tree underlying T . For vertices $v \in V(T)$, we write $\lfloor v \rfloor_T$ for the up-closure and $\lceil v \rceil_T$ for the down-closure of v with regard to \leq_T . The n th level of T is the n th level of the undirected tree underlying T .

A *directed comb* is the union of a ray with infinitely many finite paths (possibly trivial) that have precisely their first vertex on R . Hence the undirected graph underlying a directed comb is an undirected comb. The *teeth* of a directed comb are the teeth of the underlying undirected comb. The ray from the definition of a directed comb is the *spine* of the directed comb.

We write $\Lambda(D)$ for the set of all the limit edges of D . As we do for ‘ordinary’ edges of a digraph, we will suppress the brackets and the comma in our notation of limit edges. For example we write $\omega\eta$ instead of (ω, η) for a limit edge between ends ω and η . For limit edges we need the following proposition from Chapter 8.

Proposition 8.4.2. *For a digraph D , a vertex v and an end ω of D the following assertions are equivalent:*

- (i) D has a limit edge from v to ω ;
- (ii) there is a necklace $N \subseteq D$ that represents ω such that v sends an edge to every bead of N .

For vertex sets $A, B \subseteq V(D)$ let $E(A, B)$ be the set of edges from A to B , i.e., $E(A, B) = (A \times B) \cap E(D)$. Now, consider two ends $\omega, \eta \in \Omega(D)$ and a finite vertex set $X \subseteq V(D)$. If X separates ω and η we write $E(X, \omega\eta)$ as short for $E(C(X, \omega), C(X, \eta))$ and if additionally $\omega\eta$ is a limit edge, then we say that it *lives* in $E(X, \omega\eta)$.

10.2. Normal arborescences

In this section we introduce normal arborescences and we provide some basic lemmas that we need for the proofs of our main results.

Consider a digraph D and an arborescence $T \subseteq D$, not necessarily spanning. The *normal assistant* of T in D is the auxiliary digraph H that is obtained from T by adding an edge vw for every two \leq_T -incomparable vertices $v, w \in V(T)$ for which there is a T -path from $[v]_T$ to $[w]_T$ in D , regardless of whether D contains such an edge. The arborescence T is *normal* in D if the normal assistant of T in D is acyclic; in this case, we write $\trianglelefteq_T := \leq_H$ and we call \trianglelefteq_T the *normal order* of T . Similarly, a reverse arborescence T is *normal* in D if \overleftarrow{T} is normal in \overleftarrow{D} .

Lemma 10.2.1. *Let D be any digraph and let $T \subseteq D$ be an arborescence. If the normal assistant of T in D contains a cycle, then it also contains a cycle so that consecutive vertices on the cycle are \leq_T -incomparable.*

Proof. Let H be the normal assistant of T in D and let C be a cycle in H of minimal length. Suppose for a contradiction that C contains consecutive vertices that are \leq_T -comparable. As T is acyclic the cycle C cannot be contained entirely in T ; in particular C has length at least three. Thus we find a subpath $uvw \subseteq C$ such that u is the \leq_T -predecessor of v , and such that v and w are \leq_T -incomparable. But then also $uw \in E(H)$ and replacing the path uvw in C by the edge uw gives a shorter cycle. \square

An extension \preceq of \leq_T on an arborescence T is *branch sensitive* if for any two \leq_T -incomparable vertices $v \preceq w$ of T there is no $v' \in [v]_T$ with $w \preceq v'$. An extension \preceq of \leq_T on T is *path sensitive* if for no two \leq_T -incomparable vertices $v \preceq w$ the digraph D contains a T -path from w to v . Note that the normal order of any normal arborescence $T \subseteq D$ is both branch sensitive and path sensitive. A *sensitive order* on T is a linear extension of \leq_T on T that is both branch sensitive and path sensitive.

Lemma 10.2.2. *Let D be any digraph and let $T \subseteq D$ be an arborescence in D . Then T is normal in D if and only if there is a sensitive order on T .*

Proof. For the forward implication assume that T is normal in D . Let us write L_n for the n th level of T and let us write T_n for the arborescence that T induces on $\bigcup\{L_m \mid m \leq n\}$. We recursively construct an ascending sequence of orders $(\preceq_n)_{n \in \mathbb{N}}$ such that \preceq_n is a sensitive order on the arborescence T_n as follows. In the base case, we let $\preceq_0 := \leq_{T_0}$. In the recursive step, suppose that we have defined \preceq_n . Let us write for every $v \in L_n$ the set of up-neighbours (children) of v in T as N_v .

For every $v \in L_n$ let \preceq_v be a linear extension of the restriction of \preceq_T to N_v . And for every two distinct vertices $v, w \in L_n$ with $v \preceq_n w$ we define $v' \preceq_{vw} w'$ whenever $v' \in N_v$ and $w' \in [N_w]_T \setminus [v]_T$. Now, let \preceq_{n+1} be the transitive closure of

$$\preceq_n \cup \bigcup \{ \preceq_v \mid v \in L_n \} \cup \bigcup \{ \preceq_{vw} \mid v \neq w \text{ in } L_n \}.$$

It is straightforward to check that the order \preceq_{n+1} is a sensitive order on T_{n+1} . Hence $\bigcup \{ \preceq_n \mid n \in \mathbb{N} \}$ is a sensitive order on T as an ascending union of sensitive orders on subarborescences of T .

For the backward implication assume that T has a sensitive order \preceq on T . Suppose for a contradiction that T is not normal in D . Let H be the normal assistant of T in D . Then H contains a cycle C and by Lemma 10.2.1 we may assume that consecutive vertices on C are \preceq_T -incomparable. Let c be the \preceq -largest vertex on C and let c' be its successor on C . Note that $c' \preceq c$ by the choice of c . The edge cc' of $C \subseteq H$ is witnessed by a T -path P from $[c]_T$ to $[c']_T$. Let w be the first vertex and v the last vertex of P . As \preceq is branch sensitive, we have $v \preceq w$. But then the two vertices v and w together with P show that \preceq is not path sensitive contradicting that \preceq is a sensitive order on T . \square

Corollary 10.2.3. *A spanning arborescence of a finite digraph is normal if and only if it defines a depth-first search tree.*

Proof. Let T be a spanning arborescence of a finite digraph D . For the forward implication assume that T is normal in D . By Lemma 10.2.2, we find a sensitive order \preceq on T . Then T is defined by the traversed edges of the depth-first search that starts in the root of T and always chooses the \preceq -largest up-neighbour (child) in T in each step.

For the backward implication assume that T is a depth-first search tree and suppose for a contradiction that T is not normal in D . Then the normal assistant of T contains a cycle C and by Lemma 10.2.1 we may choose C so that consecutive vertices on C are \preceq_T -incomparable. Let x be the vertex on C that is visited first in the depth-first search and let y be its successor on C . The edge xy of the normal assistant of T is witnessed by a T -path from $[x]_T$ to $[y]_T$. As T is spanning this path is just an edge e . Note that all vertices in $[x]_T$ are visited earlier than those in $[y]_T$ in the depth-first search. Hence the edge e should have been visited by the depth-first search; this is a contraction because e is not an edge of T . \square

We think of (countable) normal spanning arborescence $T \subseteq D$ as being drawn in the plane with all the edges between \leq_T -incomparable vertices running from left to right; see Figure 10.0.1.

Let us see that, similar to their undirected counterparts, normal arborescences capture the separation properties of D , while they carry the simple structure of an arborescence:

Lemma 10.2.4. *Let D be any digraph and let $T \subseteq D$ be a normal arborescence in D . If $v, w \in V(T)$ are \leq_T -incomparable vertices of T with $w \not\leq_T v$, then every w - v path in D meets $X := [v]_T \cap [w]_T$. In particular, X separates v and w in D .*

Proof. Suppose for a contradiction that P is a w - v path in D that avoids X , for \leq_T -incomparable vertices $v, w \in V(T)$ with $w \not\leq_T v$. Let N_X consist of all neighbours of X in the digraph T that are contained in $V(T - X)$, let N_X^1 consist of all vertices $y \in N_X$ with $y \leq_T v$ and let $N_X^2 := N_X \setminus N_X^1$. Moreover, let Z_i be the union of the up-closures $[s]_T$ with $s \in N_X^i$, for $i = 1, 2$. Note that Z_1 and Z_2 partition $V(T - X)$. As \leq_T is branch sensitive, we observe that any two vertices $z_1, z_2 \in V(T) \setminus X$ with $z_1 \in Z_1$ and $z_2 \in Z_2$ are either incomparable with regard to \leq_T , or satisfy $z_1 \leq_T z_2$. Let z_1 be the first vertex of P in Z_1 and let z_2 be the last vertex of P in T that precedes z_1 in the path-order of P . Note that z_2 is contained in Z_2 by our assumption that P avoids X . Hence the T -path $z_2 P z_1$ witnesses that $z_2 \leq_T z_1$ contradicting our aforementioned observation. \square

The *dichromatic number* [52] of a digraph D is the smallest cardinal κ so that D admits a vertex partition into κ many partition classes that are acyclic in D . From the path sensitivity of normal arborescences we obtain the following:

Proposition 10.2.5. *Every digraph that has a normal spanning arborescence does have a countable dichromatic number.*

Proof. We denote by L_n the n th level of the arborescence T and claim that L_n is acyclic for every $n \in \mathbb{N}$. The vertices in L_n are pairwise \leq_T -incomparable. As T is path sensitive there is no w - v path in $D[L_n]$ between vertices $v \leq_T w$ in L_n . This would be violated by the \leq_T -largest vertex w and its successor v in C of any directed cycle $C \subseteq D[L_n]$. Hence the non-empty L_n define a partition of $V(D)$ into acyclic vertex sets, witnessing that D has a countable dichromatic number. \square

10.3. Arborescences are end-faithful

In this section we prove that normal spanning arborescences capture the end space combinatorially. Let $T \subseteq D$ be a fixed arborescence of a digraph D and let Ψ be a set of ends of D . We say that T is *end-faithful for Ψ* if every end in Ψ is represented by a unique ray of T that starts from the root. We call the rays in a normal arborescence T that start from the root *normal rays* of T . We say that an end ω of D is contained in the *closure* of a vertex set $U \subseteq V(D)$ if $C(X, \omega)$ meets U for every finite vertex set $X \subseteq V(D)$. Note that an end ω is contained in the closure of the vertex set of a ray R if and only if R represents ω .

Theorem 10.1. *Let D be any digraph and let $U \subseteq V(D)$ be any vertex set. If T is a normal arborescence containing U , then T is end-faithful for the set of ends in the closure of U .*

We will employ the following star-comb lemma [20, Lemma 8.2.2] in order to prove Theorem 10.1:

Lemma 10.3.1 (Star-comb lemma). *Let W be an infinite set of vertices in a connected undirected graph G . Then G contains a comb with all its teeth in W or a subdivided infinite star with all its leaves in W .*

Proof of Theorem 10.1. First, let R_1 and R_2 be distinct normal rays of T that represent ends of D in the closure of U , say ω_1 and ω_2 , respectively. Our goal is to show that ω_1 and ω_2 are distinct ends of D . By Lemma 10.2.4, the rays R_1 and R_2 have tails in distinct strong components of $D - X$ for $X = V(R_1) \cap V(R_2)$. Hence X witnesses that R_1 and R_2 are not equivalent; in particular $\omega_1 \neq \omega_2$.

It remains to show that every end ω in the closure of U is represented by a normal ray of T . We claim that there is a necklace N attached to U in D that represents ω . For this consider the auxiliary digraph D' obtained from D by adding a new vertex v^* and adding new edges v^*u , one for every $u \in U$. Since ω is contained in the closure of U , we have that $v^*\omega$ is a limit edge of D' . Note, that adding v^* does not change the set of ends, in the sense that every end of D' contains a unique end of D as a subset (of rays), and we may identify the ends of D' with the ends of D . Now, Proposition 8.4.2 yields a necklace $N \subseteq D$ that represents ω such that v^* sends an edge to every bead of N . By the definition of D' , we conclude that N is attached to U .

Having N at hand, fix a vertex from U of every bead of N and let W be the set of these fixed vertices. Now, apply the star-comb lemma in the undirected tree underlying T to W . We claim that the return is a comb. Indeed, suppose for a contradiction that we get a star and let c be its centre. By Lemma 10.2.4, the finite set $[c]_T$ separates any two leaves of the star, which is impossible because they are all contained in the necklace $N \subseteq D$.

So the return of the star-comb lemma is indeed a comb and we may assume that its spine R , considered as a ray in T , is a normal ray. Our aim is to prove that R represents ω and we may equivalently show that ω is contained in the closure of $V(R)$. So given a finite vertex set $X \subseteq V(D)$, fix teeth u and u' of the comb that are contained in $C(X, \omega)$. These exist because the teeth of the comb are contained in W and the choice of W . By Lemma 10.2.4, the strong component $C(X, \omega)$ contains a vertex of $[u]_T \cap [u']_T$. As this intersection is included in R we have verified that $C(X, \omega)$ contains a vertex of R . This completes the proof that ω is contained in the closure of R and with it the proof of this theorem. \square

Corollary 10.3.2. *Let D be any digraph and let $U \subseteq V(D)$ be any vertex set. If T is a reverse normal spanning arborescence containing U , then T is end-faithful for the set of ends in the closure of U .*

Proof. Applying Theorem 10.1 to the digraph \tilde{D} and the normal arborescence $\tilde{T} \subseteq \tilde{D}$ shows that \tilde{T} is end-faithful for the ends of \tilde{D} in the closure of U . Hence the statement is a consequence of the fact that the ends of D in the closure of U correspond bijectively to the ends of \tilde{D} in the closure of U , via the map that sends an end ω of D to the end of \tilde{D} that is represented by some (equivalently every) reverse ray of D that represents ω . \square

10.4. Arborescences reflect the horizon

One of the most useful facts about normal spanning trees is that the end space of any normal spanning tree T coincides with the end space of the graph G it spans—even topologically, i.e., the map that assigns to every end of T the end of G that contains it as a subset is a homeomorphism, see [17]. Hence, in order to understand the end space of G one just needs to understand the simple structure of the tree T .

In [11] we defined a topological space $|D|$ formed by a digraph D together with its ends and limit edges. The *horizon* of a digraph D is the subspace of $|D|$ formed by the ends of D and all the limit edges between them. In order to understand the results of this section it is not necessary to know the topology on $|D|$, as the subspace topology on the horizon of D is particularly simple. Let us give a brief description of the subspace topology for the horizon of D .

The ground set of the horizon of a digraph D is defined as follows. Take the set of ends $\Omega(D)$ of D together with a copy $[0, 1]_\lambda$ of the unit interval for every limit edge λ between two ends of D . Now, identify every end ω with the copy of 0 in $[0, 1]_\lambda$ for which ω is the tail of λ and with the copy of 1 in $[0, 1]_{\lambda'}$ for which ω is the head of λ' , for all the limit edges λ and λ' between ends of D . For inner points $z_\lambda \in [0, 1]_\lambda$ and $z_{\lambda'} \in [0, 1]_{\lambda'}$ of limit edges λ and λ' between ends of D we say that z_λ *corresponds* to $z_{\lambda'}$ if both correspond to the same point of the unit interval.

We describe the topology of the horizon of D by specifying the basic open sets. Neighbourhoods $\Omega_\varepsilon(X, \omega)$ of an end ω are of the following form: Given $X \in \mathcal{X}(D)$ let $\Omega_\varepsilon(X, \omega)$ be the union of

- the set of all the ends that live in $C(X, \omega)$ and the points of limit edges between ends that live in $C(X, \omega)$ and
- half-open partial edges $(\varepsilon, y]_\lambda$ respectively $[y, \varepsilon)_\lambda$ for every limit edge λ between ends for which y lives in $C(X, \omega)$.

Neighbourhoods $\Lambda_{\varepsilon, z}(X, \lambda)$ of inner points z of a limit edge λ between ends are of the following form: Given $X \in \mathcal{X}(D)$ that separates the endpoints of λ let $\Lambda_{\varepsilon, z}(X, \lambda)$ be the union of all the open balls of radius ε around points $z_{\lambda'}$ with λ' a limit edge between ends that lives in the bundle $E(X, \lambda)$ and with $z_{\lambda'}$ corresponding to z . Here we make the convention that for limit edges λ between ends the ε of open balls $B_\varepsilon(z)$ of radius ε around points $z \in \lambda$ is implicitly chosen small enough to guarantee $B_\varepsilon(z) \subseteq \lambda$.

Arborescences do not themselves have ends or limit edges, but there is a natural way to endow an arborescence T in a digraph D with a meaningful horizon. The *solidification* of an arborescence $T \subseteq D$, or of its normal assistant H in D , is obtained from T or H , respectively, by adding all the edges wv with $vw \in E(T)$.

Note that all the rays of T are solid in its solidification and thus represent ends there. Let us define the *horizon* of T as the horizon of the solidification of its normal assistant.

Now suppose that we have fixed a root r of T and suppose that every ray of D is solid in D . By Theorem 10.1, there exists a well-defined map ψ that sends every end ω of D to the end of the solidification \overline{T} of T represented by the unique ray $R \subseteq T$ starting from r that represents ω in D . This map ψ is clearly injective. Note that the map ψ is also surjective, by our assumption that every ray of D is solid in D . Let ζ denote the map from the set of ends of \overline{T} to that of the solidification \overline{H} of the normal assistant H of T in D that assigns to every end of \overline{T} the end of \overline{H} that contains it as a subset (of rays). This is always bijective:

Lemma 10.4.1. *Let D be a digraph, T a normal spanning arborescence of D and H the normal assistant of T in D . The map $\zeta: \Omega(\overline{T}) \rightarrow \Omega(\overline{H})$ that assigns to every end of \overline{T} the end of \overline{H} that contains it as a subset is bijective.*

Proof. To see that ζ is injective, let ω_1 and ω_2 be distinct ends of \overline{T} and let R_i be the ray in T starting from the root of T that represents ω_i for $i = 1, 2$. By Lemma 10.2.4 the two rays R_1 and R_2 have a tail in distinct strong components of $\overline{H} - X$ for $X := [R_1]_T \cap [R_2]_T$; hence ζ maps the ends ω_1 and ω_2 to distinct ends of \overline{H} .

To see that ζ is onto, let ω be an end of \overline{H} and let R be any solid ray in \overline{H} that represents ω . Our goal is to find a solid ray R' in \overline{T} that is equivalent to R in \overline{H} : then the end of \overline{T} that is represented by R' is included in ω as a subset of rays. For this apply the star-comb lemma in the undirected tree underlying T to the vertex set of R . If the return is a comb, then the comb's spine defines the desired ray R' . Indeed, the paths between the comb's spine and its teeth, define (in \overline{H}) a family of disjoint directed paths from R' to R and from R to R' , hence R' and R are equivalent in \overline{H} . It now suffices to show that the return of the star-comb lemma is always a comb; so suppose for a contradiction that it is a star with centre c say. Then, by Lemma 10.2.4, the down-closure of c in T separates infinitely many vertices of $V(R)$ in \overline{H} , contradicting that R has a tail in a strong component of $\overline{H} - X$ for every finite vertex set X . \square

Note that \overline{H} , unlike \overline{T} , can have limit edges. We say that T *reflects the horizon* of D if the map $\zeta \circ \psi: \Omega(D) \rightarrow \Omega(\overline{H})$ extends to a homeomorphism from the

horizon of D to that of \overline{H} .

The horizon of a normal spanning arborescences might differ from the horizon of the digraph it spans. This is due to the nature of connectivity in digraphs: a digraph might have a normal spanning arborescence and many strong components at the same time. For example consider the digraph D depicted in Figure 10.4.1. On the one hand, every ray in D is solid in D . On the other hand, consider the unique normal spanning arborescence T that is rooted in the leftmost vertex of the bottom ray. Note that T coincides with its normal assistant. Hence T is normal in D and the end ω in the horizon of T is a limit point of the ends ω_i in the horizon of T . In contrast to that, all points in the horizon of D are isolated as every end lives in exactly one strong component of D .




Figure 10.4.1.: A digraph D with a normal spanning arborescence T where the horizon of T differs from that of D . Every undirected edge in the figure represents a pair of inversely directed edges. Every line that ends with an arrow stands for a symmetric ray.

However, it turns out that the horizon of a digraph D coincides with the horizon of any normal spanning arborescence of D if D belongs to an important class of digraphs, namely to the class of solid digraphs:

Theorem 10.2. *Every normal spanning arborescence of a solid digraph reflects its horizon.*

The proof of this will be a consequence of the following two lemmas:

Lemma 10.4.2. *Let D be a solid digraph, let T be a normal spanning arborescence of D and H the normal assistant of T in D . For ends ω of D and ω' of \overline{H} , with $\omega' = \zeta(\psi(\omega))$ the following statements hold:*

- (i) *For every finite vertex set $X \subseteq V(D)$ there is a finite vertex set $X' \subseteq V(\overline{H})$ such that the vertex set of $C(X', \omega')$ is contained in that of $C(X, \omega)$.*
- (ii) *For every finite vertex set $X' \subseteq V(\overline{H})$ there is a finite vertex set $X \subseteq V(D)$ such that the vertex set of $C(X, \omega)$ is contained in that of $C(X', \omega')$.*

Proof. (i) Let X be any finite vertex set of D . We may assume that X is down-closed with respect to \leq_T . We write $R_{\omega'}$ for the unique ray of T that starts from the root of T and represents ω' , Theorem 10.1. Note that since D is solid every ray in T is solid in D . We will find a vertex $v \in R_{\omega'}$ such that the up-closure of v in T is contained in $C(X, \omega)$; then $X' := [v]_T \setminus \{v\}$ is as desired by the separation properties of normal arborescences, Lemma 10.2.4.

For this, we call a vertex $v \in R_{\omega'}$ *bad* if $[v]_T$ meets $V(T) \setminus C(X, \omega)$. Let us show that $R_{\omega'}$ has only finitely many bad vertices. As T is normal in D and X is down-closed we have that every strong component of $D - X$ other than $C(X, \omega)$ receives at most one edge of T from $C(X, \omega)$. Now, using that $D - X$ has only finitely many strong components, it follows that only finitely many edges of T leave $C(X, \omega)$. Let B be the finite set of all tails of edges of T that leave $C(X, \omega)$. Then also $[B]_T$ is finite and no vertex of $R_{\omega'} - [B]_T$ is bad. This shows that there are indeed only finitely many bad vertices on $R_{\omega'}$. Now, choosing a vertex v on $R_{\omega'}$ higher than any bad vertex and high enough so that the subray of $R_{\omega'}$ that starts at v is included in $C(X, \omega)$ gives $[v]_T \subseteq C(X, \omega)$.

(ii) Let X' be any finite vertex set of \overline{H} . By the separation properties of normal arborescences, Lemma 10.2.4, the finite vertex set $X := [X']_T$ is as desired. \square

Lemma 10.4.3. *Let D be a solid digraph, T a normal spanning arborescence of D and H the normal assistant of T in D . Then $\omega\eta$ is a limit edge of D if and only if $\omega'\eta'$ is a limit edge of \overline{H} , where ω' and η' is the image under the map $\zeta \circ \psi$ of ω and η , respectively.*

Proof. We write ω' and η' for the image under the map $\zeta \circ \psi$ of ω and η , respectively. Let us first show that $\omega'\eta'$ is a limit edge of \overline{H} if $\omega\eta$ is a limit edge of D . For this let any finite vertex set X' that separates ω' and η' in \overline{H} be given.

Our goal is to find an edge in \overline{H} from $C(X', \omega')$ to $C(X', \eta')$. By Theorem 10.1 there are rays R_ω and R_η in T that represent ω and η in D , respectively. As $\omega\eta$ is a limit edge of D , there is an edge in D from $[v_\omega]_T$ to $[v_\eta]_T$ for any two \leq_T -incomparable vertices $v_\omega \in R_\omega$ and $v_\eta \in R_\eta$. Now, choose such vertices v_ω and v_η so that both $[v_\omega]_T$ and $[v_\eta]_T$ avoid X' . In \overline{H} both $[v_\omega]_T$ and $[v_\eta]_T$ are strongly connected, by the definition of the solidification. And as R_ω and R_η have a tail in $C(X', \omega')$ and $C(X', \eta')$, respectively, we have that $[v_\omega]_T \subseteq C(X', \omega')$ and $[v_\eta]_T \subseteq C(X', \eta')$. In particular, $[v_\omega]_T \cap [v_\eta]_T = \emptyset$. Consequently, any edge in D from $[v_\omega]_T$ to $[v_\eta]_T$ has \leq_T -incomparable endvertices and therefore is an edge in \overline{H} from $C(X', \omega')$ to $C(X', \eta')$.

Now, let $\omega'\eta'$ be a limit edge of \overline{H} . We write ω and η for the unique preimage under $\zeta \circ \psi$ of ω' and η' , respectively. We show that $\omega\eta$ is a limit edge in D . For this let any finite vertex set X that separates ω and η in D be given. Our goal is to find an edge in D from $C(X, \omega)$ to $C(X, \eta)$. As in the proof of Lemma 10.4.2, there are vertices v_ω and v_η such that $[v_\omega]_T \subseteq C(X, \omega)$ and $[v_\eta]_T \subseteq C(X, \eta)$. Let $X' = ([v_\omega]_T \cup [v_\eta]_T) \setminus \{v_\omega, v_\eta\}$ and consider $C(X', \omega')$ and $C(X', \eta')$ in \overline{H} . Using that T is normal in D it is easy to show that $C(X', \omega') = [v_\omega]_T$ and $C(X', \eta') = [v_\eta]_T$. As $\omega'\eta'$ is a limit edge of \overline{H} there is an edge e in \overline{H} from $C(X', \omega')$ to $C(X', \eta')$. Furthermore, the endpoints of e are \leq_T -incomparable. Now, e was added to T in the definition of H because there is an edge f of D from $[v_\omega]_T$ to $[v_\eta]_T$ and this edge f is as desired. \square

Proof of Theorem 10.2. By Lemma 10.4.1 and its preceding text, the map $\zeta \circ \psi$ is a bijection. We extend this map to a bijection Θ between the horizon of D and that of \overline{H} as follows. Let y be an inner point of a limit edge $\omega\eta$ between ends of D . We write ω' and η' for the image under $\zeta \circ \psi$ of ω and η , respectively. By Lemma 10.4.3 we have that $\omega'\eta'$ is a limit edge of \overline{H} . Then we declare $\Theta(y) := y'$ for y' the point that corresponds to y on $\omega'\eta'$. Again, by Lemma 10.4.3, the map Θ is bijective; we claim that Θ is even a homeomorphism. Indeed, using Lemma 10.4.2 (ii) it is straightforward to check that Θ is continuous and using Lemma 10.4.2 (i) it is straightforward to check that the inverse of Θ is continuous. \square

10.5. Existence of arborescences

Not every digraph with a vertex that can reach all the other vertices has a normal spanning arborescence, for example any digraph D obtained from an uncountable complete graph by replacing every edge by its two orientations as separate directed edges has none. Indeed, if T is a normal arborescence of D , then any two of its vertices must be contained in the same ray starting from the root of T . Hence T cannot be spanning. In this section we give a Jung-type existence criterion for normal spanning arborescence.

For a digraph D we call a set $U \subseteq V(D)$ of vertices *dispersed* in D if there is no comb in D with all its teeth in U . Our main result of this section reads as follows:

Theorem 10.3. *Let D be any digraph, $U \subseteq V(D)$ and suppose that $r \in V(D)$ can reach all the vertices in U . If U is a countable union of dispersed sets, then D has a normal arborescence that contains U and is rooted in r .*

The converse of this is false in general. To see this consider the digraph $D = (\omega_1, E)$ with $E = \{(\alpha, \beta) \mid \alpha < \beta\}$ and $U = V(D)$. Here ω_1 denotes the first uncountable ordinal. On the one hand, no infinite subset of ω_1 is dispersed, so ω_1 cannot be written as a countable union of dispersed sets. On the other hand, the spanning arborescence that consists of all the edges with tail 0 is normal in D .

However, the converse of Theorem 10.3 holds in an important case, namely if the digraph D is solid. Indeed, if D is solid then any arborescence $T \subseteq D$ that is normal in D is locally finite by the separation properties of normal arborescences, Lemma 10.2.4. Hence the levels of T are finite; in particular, dispersed.

An analogue of Theorem 10.3 holds for reverse normal arborescences:

Corollary 10.5.1. *Let D be any digraph and suppose that $U \subseteq V(D)$ is a countable union of dispersed sets in \check{D} . If $r \in V(D)$ can be reached by all the vertices in U , then D has a reverse normal arborescence that contains U and is rooted in r .*

Proof. Apply Theorem 10.3 to the reverse of D . □

Proof of Theorem 10.3. Suppose that the vertex set U can be written as a countable union $\bigcup\{U_n \mid n \in \mathbb{N}\}$ of sets that are dispersed in D . Then we can write U as a collection $\{u_\alpha \mid \alpha < \kappa\}$ for a finite or limit ordinal κ such that every proper

initial segment of the collection is dispersed in D as follows: We may assume that the U_n are pairwise disjoint. Choose a well-ordering \leq_n of every U_n . Then write $u \leq u'$ for vertices $u \in U_m$ and $u' \in U_n$ with $m < n$, or with $m = n$ and $u \leq_m u'$. It is straightforward to show that \leq defines a well-ordering of U that is as desired.

We may assume that for every limit ordinal $\alpha < \kappa$ the vertex u_α coincides with some u_ξ with $\xi < \alpha$; indeed, just increment the subscripts of the u_α by one for α an infinite ordinal, and recursively redefine u_α to be some u_ξ with $\xi < \alpha$ for α a limit ordinal.

Now, we recursively define ascending sequences $(T_\alpha)_{\alpha < \kappa}$ and $(\preceq_\alpha)_{\alpha < \kappa}$ such that T_α is an arborescence and \preceq_α is a sensitive order of T_α that satisfies the following conditions:

- (i) T_α contains $\{u_\xi \mid \xi \leq \alpha\}$ cofinally³ with regard to \leq_{T_α} ;
- (ii) if $v, w \in T_\alpha$ with $v \preceq_\alpha w$ are distinct and have a common \leq_{T_α} -predecessor, then $w \in T_\xi$ and $v \notin T_\xi$ for some $\xi < \alpha$;
- (iii) there is no infinite strictly ascending sequence of vertices in T_α with regard to \preceq_α .

Once the T_α are defined the arborescence $T := \bigcup\{T_\alpha \mid \alpha < \kappa\}$ is as desired; indeed, $\bigcup\{\preceq_\alpha \mid \alpha < \kappa\}$ is a sensitive order on T and thus T is normal in D by Lemma 10.2.2. Finally, $V(T)$ contains U by condition (i).

Conditions (ii) and (iii) become relevant in the construction of the T_α , which now follows. If $\alpha = 0$, then let T_0 be any r - u_0 path in D and let $\preceq_0 := \leq_{T_0}$. Otherwise $\beta > 0$. If β is a limit ordinal, then let $T_\beta := \bigcup\{T_\alpha \mid \alpha < \beta\}$ and $\preceq_\beta = \bigcup\{\preceq_\alpha \mid \alpha < \beta\}$. Then \preceq_β is a sensitive order on T_β as each \preceq_α with $\alpha < \beta$ is a sensitive order on T_α . Condition (i) for β follows from (i) for $\alpha < \beta$ and our assumption that u_β coincides with u_α for some $\alpha < \beta$. Similarly, condition (ii) for β follows from (ii) for $\alpha < \beta$. Condition (iii) can be seen as follows. Suppose for a contraction that there is an infinite strictly ascending sequence $(w_n)_{n \in \mathbb{N}}$ in T_β with regard to \preceq_β . Apply the star-comb lemma to the set $\{w_n \mid n \in \mathbb{N}\}$ in the undirected tree underlying T_β . The return is an infinite subdivided undirected star since an undirected comb would give rise to a directed comb in T_β with all its teeth in U ; here we use that by (i) every tooth has a vertex of U in its up-closure

³A subset B of a poset $A = (A, \leq)$ is *cofinal* in A if for every $a \in A$ there is a $b \in B$ with $a \leq b$.

and that every proper initial segment of $\{u_\alpha \mid \alpha < \kappa\}$ is dispersed. Let Z be the set of \leq_{T_β} -up-neighbours of the centre of the subdivided star that contain a tooth in their \leq_{T_β} -up-closure. Since \preceq_β is branch sensitive we may write Z as a strictly ascending collection $Z = \{z_n \mid n \in \mathbb{N}\}$ with regard to \preceq_β . Choose $z^* \in Z \cap V(T_\alpha)$ so that α is minimal with $Z \cap V(T_\alpha) \neq \emptyset$. By (ii) we have that $z_n \preceq_\beta z^*$ for every $z_n \neq z^*$ contradicting that the z_n form a strictly ascending sequence with regard to \preceq_β .

Now, suppose that $\beta = \alpha + 1$ is a successor ordinal. If T_α already contains u_β , we let $T_\beta := T_\alpha$. Otherwise u_β is not contained in T_α . As r can reach u_β there is a T_α - u_β path P . By (iii) for α , we may choose P such that its first vertex v_P is \preceq_β -maximal among all the starting vertices of T_α - u_β paths. We let $T_\beta := T_\alpha \cup P$. Note that this ensures condition (i) for T_β .

In order to define \preceq_β we only need to describe how the vertices from $\dot{v}_P P$ relate to the vertices in T_α . We define vertices of $P - v_P$ to be smaller than all the vertices larger than v_P and larger than all others (with regard to the normal order of T_α). Note that this ensures (ii). Condition (iii) holds because there is no infinite strictly ascending sequence of vertices in T_α with regard to \preceq_α and T_β extends T_α finitely.

It remains to show that \preceq_β is a sensitive order on T_β . That \preceq_β is branch sensitive is immediate from the construction so let us prove that it is path sensitive. Suppose for a contradiction that Q is a T_β -path from w to v with \leq_{T_β} -incomparable vertices $v \preceq_\beta w$. Since T_α is normal either v or w are contained in $P - v_P$. If $w \in P - v_P$, then $v_P P w Q v$ is a path violating that \preceq_α is path sensitive; unless v_P and v are \leq_{T_β} comparable, but then we would have $w \preceq_\beta v$ by the definition of \preceq_β . In the other case, the path $w Q v P u_\beta$ would have been a better choice for P . □

Part III.

Stars and combs

It is well known, and easy to see, that every finite connected graph contains either a long path or a vertex of high degree. Similarly,

Every infinite connected graph contains either a ray or a vertex of infinite degree (*)

[20, Proposition 8.2.1]. Here, a *ray* is a one-way infinite path. Call two properties of infinite graphs *dual*, or *complementary*, in a class of infinite graphs if they partition that class. Despite (*), the two properties of ‘containing a ray’ and ‘containing a vertex of infinite degree’ are not complementary in the class of all infinite graphs: an infinite complete graph, for example, contains both. Hence it is natural to ask for structures, more specific than vertices of infinite degree and rays, whose existence is complementary to that of rays and vertices of infinite degree, respectively. Such structures do indeed exist.

For example, the property of having a vertex of infinite degree is trivially complementary, for connected infinite graphs, to the property that all distance classes from any fixed vertex are finite. This duality is employed to prove (*): if all the distance classes from some vertex are finite, then applying the infinity lemma [20, Lemma 8.1.2] to these classes yields a ray.

Similarly, having a rank in the sense of Schmidt [64] is complementary for infinite graphs to containing a ray, see [20, Lemma 8.5.2]. This duality allows for an alternative proof of (*) that avoids the use of compactness, as follows. If G is rayless, connected and infinite, then it has some rank $\alpha > 0$. Hence there is a finite vertex set $X \subseteq V(G)$ such that every component of $G - X$ has rank $< \alpha$. Then $G - X$ must have infinitely many components, and so by pigeonhole principle some vertex in X has infinite degree in G .

A stronger and localised version of (*) is the star-comb lemma [20, Lemma 8.2.2], a standard tool in infinite graph theory. Recall that a *comb* is the union of a ray R (the comb’s *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R . The last vertices of those paths are the *teeth* of this comb. Given a vertex set U , a *comb attached to U* is a comb with all its teeth in U , and a *star attached to U* is a subdivided infinite star with all its leaves in U . Then the set of teeth is the *attachment set* of the comb, and the set of leaves is the *attachment set* of the star.

Star-comb lemma. *Let U be an infinite set of vertices in a connected graph G . Then G contains either a comb attached to U or a star attached to U .*

Although the star-comb lemma trivially implies assertion (*), with $U := V(G)$, it is not primarily about the existence of one subgraph or another. Rather, it tells us something about the nature of connectedness in infinite graphs: that the way in which they link up their infinite sets of vertices can take two fundamentally different forms, a star and a comb. These two possibilities apply separately to all their infinite sets U of vertices, and clearly, the smaller U the stronger the assertion.

Just like the existence of rays or vertices of infinite degree, the existence of stars or combs attached to a given set U is not complementary (in the class of all infinite connected graphs containing U). In the first chapter of this part, we determine structures that are complementary to stars, and structures that are complementary to combs (always with respect to a fixed set U).

As stars and combs can interact with each other, this is not the end of the story. For example, a given set U might be connected in G by both a star and a comb, even with infinitely intersecting sets of leaves and teeth. To formalise this, let us say that a subdivided star S *dominates* a comb C if infinitely many of the leaves of S are also teeth of C . A *dominating star* in a graph G then is a subdivided star $S \subseteq G$ that dominates some comb $C \subseteq G$; and a *dominated comb* in G is a comb $C \subseteq G$ that is dominated by some subdivided star $S \subseteq G$. In three further chapters we shall find complementary structures to the existence of these substructures (again, with respect to some fixed set U). Here, then is an overview of the first four chapters in this part, each naming the substructure for which duality theorems are proved in its title:

10: ARBITRARY STARS AND COMBS

11: DOMINATING STARS AND DOMINATED COMBS

12: UNDOMINATED COMBS

13: UNDOMINATING STARS

Just like the original star-comb lemma, our results can be applied as structural tools in other contexts. Examples of such applications can be found in Chapter 12, 13, 14 and in this part's last chapter, which can be seen as a spin-off of its previous

chapters: There we provide further applications of our duality theorems and of the techniques that we use in order to prove them.

11. Arbitrary stars and comb

In this chapter we prove five duality theorems for combs, and two for stars. The complementary structures they offer are quite different, and not obviously inter-derivable.

Our first result is obtained by techniques of Jung [43]. Recall that a rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T , cf. [20].

Theorem 11.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *comb* attached to U ;
- (ii) there is a *rayless normal tree* $T \subseteq G$ that contains U .

To see that (ii) implies that G —in fact, the normal tree T —contains a star attached to U when U is infinite, pick from among the nodes of T that lie below infinitely many vertices of T in U one that is maximal in the tree-order of T . Then its up-closure in T contains the desired star.

Even though the normal tree from (ii) is in general not spanning, its separation properties still tell us a lot about the ambient graph G . Our next result captures this overall structure of G more explicitly (refer to [20] for the definition of tree-decompositions and adhesion sets):

Theorem 11.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *comb* attached to U ;
- (ii) G has a *rayless tree-decomposition* into parts each containing at most finitely many vertices from U and whose parts at non-leaves of the decomposition tree are all finite.

Moreover, the tree-decomposition in (ii) can be chosen with connected adhesion sets.

For $U = V(G)$, this theorem implies the following characterisation of rayless graphs by Halin [39]: *G is rayless if and only if G has a rayless tree-decomposition into finite parts.*

While Theorems 11.1 and 11.2 tell us about the structure of the graph around U , they further imply a more localised duality theorem for combs. Call a finite vertex set $X \subseteq V(G)$ *critical* if the collection $\check{\mathcal{C}}_X$ of the components of $G - X$ having their neighbourhood precisely equal to X is infinite.

Theorem 11.3. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *comb* attached to U ;
- (ii) for every infinite $U' \subseteq U$ there is a *critical vertex set* $X \subseteq V(G)$ such that infinitely many of the components in $\check{\mathcal{C}}_X$ meet U' .

Critical vertex sets were introduced in [48]. As tangle-distinguishing separators, they have a surprising background involving the Stone-Ćech compactification of G , Robertson and Seymour's tangles from their graph-minor series, and Diestel's tangle compactification, cf. [21, 47, 63]. Moreover, it turns out that Theorem 11.3 implies another characterisation of rayless graphs by Halin [38].

Schmidt's ranking of rayless graphs was employed by Bruhn, Diestel, Georgakopoulos and Sprüssel [3] to prove the unfriendly partition conjecture for the class of rayless graphs by an involved transfinite induction on their rank. We will show how the notion of a rank can be adapted to take into account a given set U , so as to give a recursive definition of those graphs that do not contain a comb attached to U . This yields our fourth duality theorem for combs:

Theorem 11.4. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *comb* attached to U ;
- (ii) G has a *U -rank*.

With these four complementary structures for combs at hand, the question arises whether there is another complementary structure combining them all. Our fifth duality theorem for combs shows that this is indeed possible:

Theorem 11.5. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *comb* attached to U ;
- (ii) G has a *tree-decomposition* that has the list (\dagger) of properties.

For the precise statement of this theorem, see Section 11.2.5. Essentially, the list (\dagger) consists of the following four properties:

- its decomposition tree stems from a normal tree as in Theorem 11.1;
- it has the properties of the tree-decomposition in Theorem 11.2;
- the infinite-degree nodes of its decomposition tree correspond bijectively to the critical vertex sets of G that are relevant in Theorem 11.3;
- the rank of its decomposition tree is equal to the U -rank of G from Theorem 11.4.

Now that we have stated all the duality theorems for combs, let us turn to our two duality theorems for stars. Recall that a vertex v of G *dominates* a ray $R \subseteq G$ if there is an infinite v – $(R - v)$ fan in G . Rays not dominated by any vertex are *undominated*, cf. [20]. Our first duality theorem for stars reads as follows:

Theorem 11.6. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *star* attached to U ;
- (ii) there is a *locally finite normal tree* $T \subseteq G$ that contains U and all whose rays are undominated in G .

To see that (ii) implies that G —in fact, the normal tree—contains a comb attached to U when U is infinite, pick a ray in the locally finite down-closure of U in the tree and extend it to a comb attached to U .

We have seen normal trees before in our first duality theorem for combs, Theorem 11.1. Theorem 11.6 above compares with Theorem 11.1 as follows. The only additional property required of the normal trees that are complementary to combs is that they are rayless. Similarly, the normal trees that are complementary to

stars have the additional property that they are locally finite. However, they have the further property that all their rays are undominated in G .

This further property is necessary to ensure that the normal trees and stars in Theorem 11.6 exclude each other. To see this, let G be obtained from a ray R by completely joining its first vertex r to all the other vertices of R , and suppose that $U = V(G)$. Then $R \subseteq G$ with root r is a locally finite normal tree containing U . But the edges of G at r form a star attached to U , so the further property is indeed necessary.

By contrast, we do not need to require in Theorem 11.1 that all the stars in the normal trees that are complementary to combs are undominating in G : this is already ensured by the nature of normal trees (see Lemma 11.2.4 for details).

Our second duality theorem for stars is phrased in terms of tree-decompositions, similar to Theorem 11.2:

Theorem 11.7. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *star* attached to U ;
- (ii) G has a *locally finite tree-decomposition* with finite and pairwise disjoint adhesion sets such that each part contains at most finitely many vertices from U .

Moreover, the tree-decomposition in (ii) can be chosen with connected adhesion sets.

This chapter is organised as follows. Section 11.1 provides the tools and terminology that we use throughout this part. Section 11.2 and 11.3 are dedicated to the duality theorems for combs and stars respectively.

Throughout this part, $G = (V, E)$ is an arbitrary graph.

11.1. Tools and terminology

An independent set M of edges in a graph G is called a *matching* of A and B for vertex sets $A, B \subseteq V(G)$ if every edge in M has one endvertex in A and the other in B .

11.1.1. Star-Comb Lemma

The predecessors of the star-comb lemma are the following facts:

Lemma 11.1.1 ([20, Proposition 9.4.1]). *For every $m \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that each connected finite graph with at least n vertices either contains a path of length m or a star with m leaves as a subgraph.*

Lemma 11.1.2 ([20, Proposition 8.2.1]). *A connected infinite graph contains either a ray or a vertex of infinite degree.*

The latter is a direct consequence of the infinity lemma, [20, Lemma 8.1.2]. Lemma 11.1.1 has been generalised to higher connectivity, [34, 42, 53], and so has Lemma 11.1.2 in [36, 40, 53]. For an overview we recommend the introduction of [36].

For locally finite trees, Lemma 11.1.2 already yields a comb:

Lemma 11.1.3. *If U is an infinite set of vertices in a locally finite rooted tree T , then T contains a comb attached to U whose spine starts at the root.*

Proof. The down-closure of U in the tree-order of T induces a locally finite subtree which, by Lemma 11.1.2 above, contains a ray starting at the root, say. This ray can be extended recursively to the desired comb. \square

For rayless trees, the situation is simpler:

Lemma 11.1.4. *If U is an infinite set of vertices in a rayless rooted tree T , then T contains a star attached to U which is contained in the up-closure of its central vertex in the tree-order of T .*

Proof. Among all the nodes of T that lie below some infinitely nodes from U , pick one node t , say, that is maximal in the tree-order of T . Then t has infinite degree and we find the desired star with centre t in the up-closure of t . \square

We already stated the star-comb lemma in its basic form in the introduction, but a stronger version is known:

Lemma 11.1.5 (Star-comb lemma). *Let G be any connected graph and let $U \subseteq V(G)$ be infinite. If $\kappa \leq |U|$ is a regular¹ cardinal, then U has a subset U' of size κ such that at least one of the following assertions holds:*

¹A cardinal κ is *regular* if there is no family $(\kappa_\alpha \mid \alpha < \lambda)$ with $\lambda < \kappa$ and all $\kappa_\alpha < \kappa$ such that $\bigcup_{\alpha < \lambda} \kappa_\alpha = \kappa$. For example, \aleph_0 and \aleph_1 are regular while $\aleph_\omega = \bigcup_{n < \omega} \aleph_n$ is not.

- (i) G contains a comb attached to U whose attachment set is U' ;
- (ii) G contains a star attached to U whose attachment set is U' .

In particular, if κ is uncountable, then (i) fails and (ii) holds for every such U' .

For singular cardinals κ this version of the star-comb lemma is not true in general, as the following example demonstrates. Consider the singular cardinal $\kappa = \aleph_\omega$. Let G be the rayless tree that is obtained from a $K_{1,\omega}$ with ω as set of leaves by adding pairwise disjoint copies of K_{1,\aleph_n} , one for each non-zero $n < \omega$, such that K_{1,\aleph_n} meets $K_{1,\omega}$ precisely in n and n happens to be the central vertex of K_{1,\aleph_n} . Then the rayless tree G cannot contain a comb, and it cannot contain subdivision of a star $K_{1,\kappa}$ since every vertex of G has degree $< \kappa$, but the vertex set of G has size κ .

Recently, Gollin and Heuer [36] introduced a way more complex version of the star-comb lemma above for the more difficult singular case, the *Frayed-Star-Comb Lemma*, [36, Corollary 8.1].

The version for regular cardinals has been proved in, e.g., [27] and [36]. We repeat the short proof here for the sake of convenience:

Proof of Lemma 11.1.5. Using Zorn's lemma we find a maximal tree $T \subseteq G$ all whose edges lie on a U -path in T . Then T contains U .

If T has a vertex v of degree κ , then its incident edges extend to v - U paths whose union is the desired star with U' its attachment set.

Otherwise every vertex of T has degree $< \kappa$. After fixing an arbitrary vertex, an inductive argument—utilising the regularity of κ —shows that every distance class of T has size $< \kappa$. As $V(T)$ is the countable union of these distance classes, we deduce from the regularity of κ that $\kappa = \aleph_0$ is the only possibility. This, then, means that T is locally finite, and hence contains a ray by Lemma 11.1.2. As every edge of T lies on a U -path in T , an inductive construction turns this ray into a comb attached to U , and we may let U' consist of its $\aleph_0 = \kappa$ many teeth. \square

We remark that this version of the star-comb lemma can be proved alternatively by means of [46, Lemma III.6.14].

11.1.2. Separations

For a vertex set $X \subseteq V(G)$ we denote the collection of the components of $G - X$ by \mathcal{C}_X . If any $X \subseteq V(G)$ and $\mathcal{C} \subseteq \mathcal{C}_X$ are given, then these give rise to a separation of G which we denote by

$$\{X, \mathcal{C}\} := \{V \setminus V[\mathcal{C}], X \cup V[\mathcal{C}]\}$$

where $V[\mathcal{C}] = \bigcup \{V(C) \mid C \in \mathcal{C}\}$. Note that every separation $\{A, B\}$ of G can be written in this way. For the orientations of $\{X, \mathcal{C}\}$ we write

$$(X, \mathcal{C}) := (V \setminus V[\mathcal{C}], X \cup V[\mathcal{C}]) \quad \text{and} \quad (\mathcal{C}, X) := (V[\mathcal{C}] \cup X, V \setminus V[\mathcal{C}]).$$

We write $\{X, C\}$ and (X, C) and (C, X) instead of $\{X, \{C\}\}$ and $(X, \{C\})$ and $(\{C\}, X)$ respectively. The set of all finite-order separations of a graph G is denoted by $S_{\aleph_0} = S_{\aleph_0}(G)$.

11.1.3. Ends of graphs

We write $\mathcal{X} = \mathcal{X}(G)$ for the collection of all finite subsets of the vertex set V of G , partially ordered by inclusion. An *end* of G , as defined by Halin [37], is an equivalence class of rays of G , where a ray is a one-way infinite path. Here, two rays are said to be *equivalent* if for every $X \in \mathcal{X}$ both have a subray (also called *tail*) in the same component of $G - X$. So in particular every end ω of G chooses, for every $X \in \mathcal{X}$, a unique component $C(X, \omega) = C_G(X, \omega)$ of $G - X$ in which every ray of ω has a tail. In this situation, the end ω is said to *live* in $C(X, \omega)$. The set of ends of a graph G is denoted by $\Omega(G)$. We use the convention that Ω always denotes the set of ends $\Omega(G)$ of the graph named G .

A vertex v of G *dominates* a ray $R \subseteq G$ if there is an infinite v - $(R - v)$ fan in G . Rays not dominated by any vertex are *undominated*. An end of G is *dominated* and *undominated* if one (equivalently: each) of its rays is dominated and undominated, respectively. If v does not dominate ω , then there is an $X \in \mathcal{X}$ which *strictly separates* v from ω in that $v \notin X \cup C(X, \omega)$. More generally, if no vertex of $Y \in \mathcal{X}$ dominates ω , then there is an $X \in \mathcal{X}$ *strictly separating* Y from ω in that Y avoids the union $X \cup C(X, \omega)$. Let us say that an oriented finite-order separation (A, B) *strictly separates* a set $X \subseteq V(G)$ of vertices from a set $\Psi \subseteq \Omega$ of ends if $X \subseteq A \setminus B$ and every end in Ψ lives in a component of $G[B \setminus A]$.

Let us say that an end ω of G is contained *in the closure* of M , where M is either a subgraph of G or a set of vertices of G , if for every $X \in \mathcal{X}$ the component $C(X, \omega)$ meets M . Equivalently, ω lies in the closure of M if and only if G contains a comb attached to M with its spine in ω . We write $\partial_\Omega M$ for the subset of Ω that consists of the ends of G lying in the closure of M . Note that $\partial_\Omega H$ usually differs from $\Omega(H)$ for subgraphs $H \subseteq G$: For example, if G is a ladder and H is its outer double ray, then $\partial_\Omega H$ consists of the single end of G while $\Omega(H)$ consists of the two ends of the double ray in H . Readers familiar with $|G|$ as in [20] will note that $\partial_\Omega M$ is the intersection of Ω with the closure of M in $|G|$, which in turn coincides with the topological frontier of $M \setminus \overset{\circ}{E}$ in the space $|G| \setminus \overset{\circ}{E}$. If an end ω of G does not lie in the closure of M , and if $X \in \mathcal{X}$ witnesses this (in that $C(X, \omega)$ avoids M), then X is said to *separate* ω from M (and M from ω). Carmesin [15] observed that

Lemma 11.1.6. *Let G be any graph. If $H \subseteq G$ is a connected subgraph and ω is an undominated end of G lying in the closure of H , then H contains a ray from ω .*

Proof. Since ω lies in the closure of H we find a comb in G attached to H with spine in ω . And as ω is undominated in G , the star-comb lemma in H must return a comb in H attached to the attachment set of the first comb. Then the two combs' spines are equivalent in G . \square

Another way of viewing the ends of a graph goes via its *directions*: choice maps f assigning to every $X \in \mathcal{X}$ a component of $G - X$ such that $f(X') \subseteq f(X)$ whenever $X' \supseteq X$. Every end ω defines a unique direction f_ω by mapping every $X \in \mathcal{X}$ to $C(X, \omega)$. Conversely, Diestel and Kühn proved in [27] (Theorem 11.1.7 below) that every direction in fact comes from a unique end in this way, thus giving a one-to-one correspondence between the ends and the directions of a graph.

The advantage of this point of view stems from an inverse limit² description of the directions: note that \mathcal{X} is directed³ by inclusion; for every $X \in \mathcal{X}$ let \mathcal{C}_X consist of the components of $G - X$; endow each \mathcal{C}_X with the discrete topology; and let $\mathbf{c}_{X', X}: \mathcal{C}_{X'} \rightarrow \mathcal{C}_X$ for $X' \supseteq X$ send each component of $G - X'$ to the component of $G - X$ containing it; then $\{\mathcal{C}_X, \mathbf{c}_{X', X}, \mathcal{X}\}$ is an inverse system whose inverse limit, by construction, consists of the directions.

²For details on inverse limits, see e.g. [32] or [62].

³A poset (P, \leq) is said to be *directed* if for all $p, q \in P$ there is an $r \in P$ with $r \geq p$ and $r \geq q$.

Theorem 11.1.7 ([27, Theorem 2.2]). *Let G be any graph. Then the map $\omega \mapsto f_\omega$ is a bijection between the ends of G and its directions, i.e. $\Omega = \varprojlim \mathcal{C}_X$.*

From now on we do not distinguish between Ω and the inverse limit space $\varprojlim \mathcal{C}_X$ with the inverse limit topology, and we call Ω the end space.

If a graph G is locally finite, then the star-comb lemma always yields a comb. This fact has been generalised in Lemma 11.1.8 below, where the proof relies on the combination of Halin’s combinatorial definition of an end with the topological inverse limit point of view on ends as directions:

Lemma 11.1.8. *Let G be any graph and let $U \subseteq V(G)$ be infinite. If for every $X \in \mathcal{X}$ only finitely many components of $G - X$ meet U , then $\partial_\Omega U$ is a non-empty and compact subspace of Ω .*

Proof. For every $X \in \mathcal{X}$ let $\mathcal{K}_X \subseteq \mathcal{C}_X$ consist of the finitely many components of $G - X$ that meet U . Then the closed subspace $\partial_\Omega U$ of the inverse limit $\Omega = \varprojlim \mathcal{C}_X$ is non-empty and compact as inverse limit of its non-empty compact Hausdorff projections \mathcal{K}_X , cf. [32, Corollary 2.5.7]. \square

The combination of topology and infinite graph theory is known as *topological infinite graph theory*.⁴ And in fact, Lemma 11.1.8 can be employed⁵ to deduce a well-known result of Diestel from this field, [24, Theorem 4.1], which states that a graph is compactified by its ends if and only if it is *tough* in that deleting any finite set of vertices always leaves only finitely many components.

Since Lemma 11.1.8 yields combs even when there are both combs and stars (for example if G is an infinite complete graph), this plus of control makes it a useful addition to the star-comb lemma.

11.1.4. Critical vertex sets

We have indicated above that adding the ends generally does not suffice to compactify a graph with the usual topologies.

⁴An overview on this young field is presented in [19, 20].

⁵If G is tough and a covering of $G \sqcup \Omega$ with basic open sets is given, first apply Lemma 11.1.8 to V to obtain a finite subcover \mathcal{O} of Ω , then apply Lemma 11.1.8 to $U = V \setminus \bigcup \mathcal{O}$ to deduce that U is finite and, therefore, $G \setminus \bigcup \mathcal{O}$ is compact.

However, every graph is naturally compactified by its ends plus critical vertex sets, where a finite set X of vertices of an infinite graph G is *critical* if the collection

$$\mathcal{C}_X^\checkmark := \{ C \in \mathcal{C}_X \mid N(C) = X \}$$

is infinite (cf. [21, 47, 48]). When G is connected, all its critical vertex sets are non-empty, and so it follows that G having a critical vertex set is stronger than G containing an infinite star: On the one hand, given a critical vertex set X , each $x \in X$ sends an edge to each of the infinitely many components $C \in \mathcal{C}_X^\checkmark$ and therefore is the centre of an infinite star. On the other hand, if G is obtained from a ray R by completely joining its first vertex r to all the other vertices of R , then G contains an infinite star but no critical vertex set.

Let us say that a critical vertex set X of G lies *in the closure* of M where M is either a subgraph of G or a set of vertices of G , if infinitely many components in \mathcal{C}_X^\checkmark meet M . The collection of all critical vertex sets of G is denoted by $\text{crit}(G)$. The *combinatorial remainder* of a graph G is the disjoint union $\Gamma(G) := \Omega(G) \sqcup \text{crit}(G)$. As usual, $\Gamma = \Gamma(G)$, and $\partial_\Gamma M$ consists of those $\gamma \in \Gamma$ lying in the closure of M . We obtain a slight strengthening of the star-comb lemma:

Lemma 11.1.9. *Let G be any graph and let $U \subseteq V(G)$ be infinite. Then at least one of the following assertions holds:*

- (i) G has an end lying in the closure of U ;
- (ii) G has a critical vertex set lying in the closure of U .

Proof. If there is a vertex set $X' \in \mathcal{X}$ such that infinitely many components of $G - X'$ meet U , then X' includes a critical vertex set X such that infinitely many components in \mathcal{C}_X^\checkmark meet U , giving (ii). Otherwise Lemma 11.1.8 gives (i). \square

11.1.5. Normal trees

A rooted tree $T \subseteq G$, not necessarily spanning, is said to be *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T , [20, p. 220]. We say that a vertex set $W \subseteq V(G)$ is *normally spanned* in G if there is a normal tree in G that contains W . A graph G is *normally spanned* if $V(G)$ is normally spanned, i.e., if G has a normal spanning tree.

The *generalised up-closure* $\llbracket x \rrbracket$ of a vertex $x \in T$ is the union of $\lfloor x \rfloor$ with the vertex set of $\bigcup \mathcal{C}(x)$, where the set $\mathcal{C}(x)$ consists of those components of $G - T$ whose neighbourhoods meet $\lfloor x \rfloor$. Every graph G reflects the separation properties of each normal tree $T \subseteq G$ (we generalise [20, Lemma 1.5.5] to possibly non-spanning normal trees):

Lemma 11.1.10. *Let G be any graph and let $T \subseteq G$ be any normal tree.*

- (i) *Any two vertices $x, y \in T$ are separated in G by the vertex set $\lfloor x \rfloor \cap \lfloor y \rfloor$.*
- (ii) *Let $W \subseteq V(T)$ be down-closed. Then the components of $G - W$ come in two types: the components that avoid T ; and the components that meet T , which are spanned by the sets $\llbracket x \rrbracket$ with x minimal in $T - W$.*

Proof. (i) The proof is that of [20, Lemma 1.5.5 (i)].

(ii) In a first step, we prove that if a component C of $G - W$ meets T and x is minimal in $C \cap T$, then $C = G[\llbracket x \rrbracket]$. The backward inclusion holds because $\llbracket x \rrbracket$ is connected, avoids W and contains x . The forward inclusion can be seen as follows. On the one hand, $C \cap T \subseteq \lfloor x \rfloor$. Indeed, by (i), any x - y path in C with $y \in C \cap T$ contains a vertex below both x and y and every such vertex must be the minimal vertex x itself. On the other hand, $C - T \subseteq \bigcup \mathcal{C}(x)$. Indeed, every component C' of $C - T$ is a component of $G - T$ since $W \subseteq T$, and by $C \cap T \subseteq \lfloor x \rfloor$ each neighbour of C' inside C must be contained in $\lfloor x \rfloor$.

Now let us deduce (ii). Without loss of generality W is not empty. To begin, we prove that each component C of $G - W$ meeting T is spanned by $\llbracket x \rrbracket$ for some minimal x in $T - W$. By the first step, it suffices to show that a minimal vertex x of $C \cap T$ is also minimal in $T - W$, a fact that we verify as follows. The vertices below x form a chain $\lceil t \rceil$ in T . As t is a neighbour of x , the maximality of C as a component of $G - W$ implies that $t \in W$, giving $\lceil t \rceil \subseteq W$ since W is down-closed. Hence x is also minimal in $T - W$.

Conversely, if x is any minimal element of $T - W$, it is clearly also minimal in $C \cap T$ for the component C of $G - W$ to which it belongs. Together with the first step we conclude that C is a component of $G - W$ meeting T and spanned by $\llbracket x \rrbracket$. □

As a consequence, the *normal rays* of a normal spanning tree $T \subseteq G$, those that start at the root, reflect the end structure of G in that every end of G contains exactly one normal ray of T , [20, Lemma 8.2.3]. More generally,

Lemma 11.1.11. *If G is any graph and $T \subseteq G$ is any normal tree, then every end of G in the closure of T contains exactly one normal ray of T . Moreover, sending these ends to the normal rays they contain defines a bijection between $\partial_\Omega T$ and the normal rays of T .*

Proof. Let ω be any end of G in the closure of T . By Lemma 11.1.10 (i) at most one normal ray of T is contained in ω , and so it remains to find a normal ray of T that lies in ω . For this, we pick a comb in G attached to T with its spine in ω . We construct a normal ray of T in ω , as follows.

Starting with the root v_0 of T , recursively choose nodes v_0, v_1, v_2, \dots of T such that v_{n+1} is the minimal vertex of $T - [v_n]$ for which $\llbracket v_{n+1} \rrbracket$ spans the component of $G - [v_n]$ that contains all but finitely many vertices of the comb. Such a vertex v_{n+1} exists by Lemma 11.1.10 (ii). And it is an upward neighbour of v_n , which can be seen by applying Lemma 11.1.10 (i) to v_n and v_{n+1} . In conclusion $v_0 v_1 v_2 \dots$ is a normal ray of T that is equivalent in G to the spine of the comb.

The ‘moreover’ part holds as every normal ray of T has its end in G contained in the closure of T . □

Consequently, if G contains a comb attached to T , then T contains exactly one normal ray that is equivalent in G to that comb’s spine.

Lemma 11.1.12. *Let G be any graph and let $T \subseteq G$ be any normal tree. Then every critical vertex set of G in the closure of T is contained in T as a chain.*

Proof. Let X be any critical vertex set of G that lies in the closure of T . For every component $C \in \mathcal{C}_X$ that meets T , pick a C - X edge from T . By the pigeonhole principle, some infinitely many of these edges have the same endpoint $x \in X$, giving rise to an infinite star in T . Then, by Lemma 11.1.10, $[x]$ pairwise separates all the leaves of the star above x at once; let us write L for the set of these leaves. Since $[x]$ is finite, all but finitely many of the infinitely many components in \mathcal{C}_X that meet L are also components of $G - [x]$. And every vertex from X defines at least one path of length two between distinct such components, by the definition of critical vertex sets. Therefore, no vertex in X can be contained in a component of $G - [x]$; in other words, X is contained in the chain $[x]$. □

11.1.6. Containing vertex sets cofinally

We say that a rooted tree $T \subseteq G$ contains a set W *cofinally* if $W \subseteq V(T)$ and W is cofinal⁶ in the tree-order of T . Interestingly, our next lemma does not require T to be normal.

Lemma 11.1.13. *Let G be any graph. If $T \subseteq G$ is a rooted tree that contains a vertex set W cofinally, then $\partial_\Gamma T = \partial_\Gamma W$.*

Proof. We first prove that $\partial_\Omega T = \partial_\Omega W$. The backward inclusion $\partial_\Omega T \supseteq \partial_\Omega W$ holds as T contains W . For the forward inclusion we prove equivalently that every end of G that is not contained in the closure of W also does not lie in the closure of T . So consider any end $\omega \in \Omega \setminus \partial_\Omega W$, and pick a finite vertex set $X \subseteq V(G)$ separating W from ω . We claim that the finite set X' consisting of the vertices in X and all vertices in the down-closure of $X \cap V(T)$ in T , i.e. $X' := X \cup \downarrow[X \cap V(T)]_T$, separates T from ω . Indeed, suppose for a contradiction that the component $C := C(X', \omega)$ of $G - X'$ meets T . Consider a vertex $v \in C \cap T$. As $X' \cap V(T)$ is down-closed in T , the up-closure $\uparrow v$ in T is included in C . Hence—as T contains W cofinally—the component C also contains a vertex from W , contradicting the assumption that $X \subseteq X'$ separates W from ω .

It remains to show that $\partial_\Gamma T$ and $\partial_\Gamma W$ coincide on $\text{crit}(G)$. From $W \subseteq T$ we infer $\partial_\Gamma W \subseteq \partial_\Gamma T$, so it suffices to show that every critical vertex set that lies in the closure of T does also lie in the closure of W . For this, let any critical vertex set $X \in \partial_\Gamma T$ be given. We pick, for every component $C \in \mathcal{C}_X^\checkmark$ meeting T , a vertex $u(C)$ of T in C . Then applying the star-comb lemma in T to this infinite vertex set yields either a star or a comb attached to it. Since the finite vertex set X pairwise separates every two vertices in the attachment set at once, we in fact get a star. Consider the centre of the star. This is a vertex of T that has infinitely many pairwise incomparable vertices $u(C)$ above it. Using that T contains W cofinally, we find a vertex $w(C)$ in $T \cap W$ above every $u(C)$. As X is finite, we may assume without loss of generality that every vertex $w(C)$ is contained C . Then X lies in the closure of the vertex set formed by the vertices $w(C)$, and hence $X \in \partial_\Gamma W$ follows. \square

⁶A subset X of a poset $P = (P, \leq)$ is *cofinal* in P , and \leq , if for every $x \in X$ there is a $p \in P$ with $p \geq x$.

11.1.7. Tree-decompositions and S -trees

In this part we assume familiarity with [20, Section 12.3] up to but not including Lemma 12.3.2, and with the concepts of oriented separations and S -trees for S a set of separations of a given graph as presented in [20, Section 12.5]. Whenever we introduce a tree-decomposition as (T, \mathcal{V}) we tacitly assume that $\mathcal{V} = (V_t)_{t \in T}$. Usually we refer to the adhesion sets of a tree-decompositions as separators. We call a tree-decomposition *rayless* and *locally finite* if the decomposition tree T is rayless and locally finite, respectively. A *star-decomposition* is a tree-decomposition whose decomposition-tree is a star $K_{1,\kappa}$ for some cardinal κ . A *rooted* tree-decomposition is a tree-decomposition (T, \mathcal{V}) where T is rooted. We say that a rooted tree-decomposition (T, \mathcal{V}) of G *covers* a vertex set $U \subseteq V(G)$ *cofinally* if the set of nodes of T whose parts meet U is cofinal in the tree-order of T .

We will need the following standard facts about tree-decompositions:

Lemma 11.1.14 ([20, Lemma 12.3.1]). *Let G be any graph with a tree-decomposition (T, \mathcal{V}) and let $t_1 t_2$ be any edge of T and let T_1, T_2 be the components of $T - t_1 t_2$, with $t_1 \in T_1$ and $t_2 \in T_2$. Then $V_{t_1} \cap V_{t_2}$ separates $A_1 := \bigcup_{t \in T_1} V_t$ from $A_2 := \bigcup_{t \in T_2} V_t$ in G .*

Corollary 11.1.15. *Let (T, \mathcal{V}) be any tree-decomposition of any graph G . If a connected subgraph $H \subseteq G$ avoids a part V_t , then there is a unique component T' of $T - t$ with $H \subseteq \bigcup_{t' \in T'} G[V_{t'}]$ and H avoids every part that is not at a node of the component T' . \square*

A tree-decomposition (T, \mathcal{V}) makes T into an S -tree for the set S of separations it induces, cf. [20]. The converse is true, for example if T is rayless, but false in general (it is no longer clear that every vertex of G lives in some part if T contains a ray). By a simple distance argument, however, the converse holds in a special case for which we need the following definition. Suppose that (T, α) is an S -tree with T rooted in $r \in T$. We say that the separators of (T, α) are *upwards disjoint* if for every two edges $\vec{e} < \vec{f}$ pointing away from the root r the separators of $\alpha(\vec{e})$ and $\alpha(\vec{f})$ are disjoint. Then every S -tree with upwards disjoint separators induces a tree-decomposition.

We use the following non-standard notation for S -trees (T, α) : for an edge $xy = e$ of the decomposition tree T we abbreviate $\alpha(e, x, y) = \alpha(x, y)$.

11.1.8. Tree-decompositions and S -trees displaying sets of ends

In this section we give a brief summary of how the ends of G relate to the decomposition trees of tree-decompositions and S -trees. For the sake of readability, we introduce all needed concepts for S -trees and let the tree-decompositions inherit these concepts from their corresponding S -trees.

Let (T, α) be any S_{\aleph_0} -tree. If ω is an end of G , then ω orients every finite-order separation $\{A, B\} \in S_{\aleph_0}$ of G towards the side $K \in \{A, B\}$ for which every ray in ω has a tail in $G[K]$. In this way, ω induces a consistent orientation of \vec{S}_{\aleph_0} and, via α , also induces a consistent orientation O of $\vec{E}(T)$. Then ω either *lives* at a unique node $t \in T$ in that the star $\vec{F}_t = \{(e, s, t) \in \vec{E}(T) \mid e = st \in T\}$ at t is included in O , or *corresponds* naturally to a unique end η of T in that for some (equivalently: every) ray $t_1 t_2 \dots$ in η all oriented edges $(t_n t_{n+1}, t_n, t_{n+1})$ are contained in O . When (T, α) corresponds to a tree-decomposition (T, \mathcal{V}) and ω lives at t , then we also say that ω *lives* in the part V_t at t . Moreover, we remark that ω lives in V_t if and only if some (equivalently: every) ray in ω has infinitely many vertices in V_t . Likewise, ω corresponds to η if and only if some (equivalently: every) ray $R \in \omega$ follows the course of some (equivalently: every) ray $W \in \eta$ (in that for every tail $W' \subseteq W$ the ray R has infinitely many vertices in $\bigcup_{t \in W'} V_t$). In both cases ‘having infinitely many vertices in’ cannot be replaced with ‘having a tail in’, e.g. consider decomposition trees that are infinite stars or combs whose teeth avoid their spines.

Consider the map $\tau: \Omega(G) \rightarrow \Omega(T) \sqcup V(T)$ that takes each end of G to the end or node of T which it corresponds to or lives at respectively. This map essentially captures how the ends of G relate to the ends of T . We say that (T, α) *displays* a set of ends $\Psi \subseteq \Omega(G)$ if τ restricts to a bijection $\tau \upharpoonright \Psi: \Psi \rightarrow \Omega(T)$ between Ψ and the end space of T and maps every end that is not contained in Ψ to some node of T .

It is a natural and largely open question for which subsets $\Psi \subseteq \Omega(G)$ a graph G has a tree-decomposition (T, \mathcal{V}) that displays Ψ . Only recently, Carmesin achieved a major breakthrough by providing a positive answer for Ψ the set of undominated ends of G . In order to state his result in its full strength, we introduce two more definitions and motivate them in a lemma.

Suppose that T is rooted in $r \in T$. Let us say that the separators of (T, α) are *upwards disjoint* if for every two edges $\vec{e} < \vec{f}$ pointing away from the root r the separators of $\alpha(\vec{e})$ and $\alpha(\vec{f})$ are disjoint. Here, $\vec{e} = (e, s, t)$ *points away from* r if $r \leq_T s <_T t$, i.e., if $s \in rTt$. If the finite separators of (T, α) are upwards disjoint, then by the star-comb lemma and a simple distance argument, every end of T has some ends of G corresponding to it (i.e. $\tau^{-1}(\eta) \neq \emptyset$ for every end η of T). And if additionally (T, α) is *upwards connected* in that for every edge \vec{e} pointing away from the root r the induced subgraph $G[B]$ stemming from $(A, B) = \alpha(\vec{e})$ is connected, then T already displays the set of those ends of G that correspond naturally to ends of T (i.e. $|\tau^{-1}(\eta)| = 1$ for every end η of T):

Lemma 11.1.16. *Let G be any graph. Every upwards connected rooted S_{\aleph_0} -tree (T, α) with upwards disjoint separators displays the ends of G that correspond to the ends of T .*

Proof. By our preliminary remarks it remains to show that for every end η of T there is at most one end of G corresponding to η . Suppose for a contradiction that η is an end of T such that two distinct ends $\omega \neq \omega'$ of G correspond to it, and write R for the rooted ray of T that represents η . Pick $X \in \mathcal{X}$ such that ω and ω' live in distinct components of $G - X$. As the separators of (T, α) are upwards disjoint, by a distance argument we find an edge $e \in R$ with orientation \vec{e} away from the root such that the separation $(A, B) = \alpha(\vec{e})$ satisfies $B \cap X = \emptyset$. Now both of the two ends ω and ω' have rays in $G[B]$ because both of them correspond to η . And in $G[B]$ we find paths connecting these rays, since (T, α) is upwards connected. But then these rays and paths avoid X , contradicting the choice of X . \square

Now we are ready to state the following result of Carmesin [15] that solved a conjecture of Diestel [17] from 1992 (in amended form) and, as a corollary, also solved a conjecture of Halin [37] from 1964 (again in amended form):

Theorem 11.1.17 (Carmesin 2014). *Every connected graph G has an upwards connected rooted tree-decomposition with upwards disjoint finite separators that displays the undominated ends of G .*

The theorem above accumulates Carmesin's Theorem 1, Remark 6.6 and the second paragraph of his 'Proof that Theorem 1 implies Corollary 2.6'.

Our [6, Lemma 3.7] will allow us to strengthen Carmesin’s theorem so that it states that *every connected graph G has a tree-decomposition with pairwise disjoint finite connected separators that displays the undominated ends of G .*

11.2. Combs

Jung [43] noted that, given any connected graph G and any vertex set $U \subseteq V(G)$, the absence of a comb attached to U is equivalent to U being *dispersed* in G , meaning that for every ray $R \subseteq G$ there is a finite vertex set $X \subseteq V(G)$ separating R from U . This equivalence then gives another equivalence as U being dispersed rephrases to ‘no end of G lies in the closure of U ’. For readers familiar with the topological space $|G| = G \sqcup \Omega$ as in [20], this is to say that U is closed in $|G|$. These assertions—while equivalent to the absence of a comb—are abstract and do not immediately provide concrete structures that are complementary to combs. Providing concrete complementary structures is the aim of this section.

11.2.1. Normal trees

In this section we prove

Theorem 11.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) *G contains a comb attached to U ;*
- (ii) *there is a rayless normal tree $T \subseteq G$ that contains U .*

Moreover, the normal tree T in (ii) can be chosen such that it contains U cofinally.

For this, we need the following key results of Jung’s proof of his 1967 characterisation, Theorem 11.2.5, of the connected graphs that have normal spanning trees.

Proposition 11.2.1 (Jung). *Let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. If U is a countable union $\bigcup_{n \in \mathbb{N}} U_n$ of dispersed sets $U_n \subseteq V(G)$ and v is any vertex of G , then G contains an ascending sequence $T_0 \subseteq T_1 \subseteq \dots$ of rayless normal trees $T_n \subseteq G$ such that each T_n contains $U_0 \cup \dots \cup U_n$ cofinally and is rooted in v . In particular, the overall union $T := \bigcup_{n \in \mathbb{N}} T_n$ is a normal tree in G that contains U cofinally and is rooted in v .*

Proof. It suffices to show that, given a rayless normal tree T_n containing $U_{\leq n} := U_0 \cup \dots \cup U_n$ cofinally, we find a rayless normal tree T_{n+1} extending T_n and containing $U_{\leq n+1} = U_{\leq n} \cup U_{n+1}$ cofinally. For this, let any T_n be given. Consider the collection of all normal trees $T \supseteq T_n$ with $T \cap U_{\leq n+1}$ cofinal in the tree-order of T , partially ordered by letting $T \leq T'$ whenever T is extended by T' as normal tree. Since U_{n+1} is dispersed and T_n is rayless, all of these trees must be rayless. Let T_{n+1} be a maximal tree that Zorn's lemma provides for this poset. In the following we show that T_{n+1} is as desired.

Assume for a contradiction that some vertex $u \in U_{\leq n+1}$ is not contained in T_{n+1} . Since T_{n+1} is normal, the neighbourhood of the component C of $G - T_{n+1}$ that contains u forms a chain in the tree-order of T_{n+1} . As T_{n+1} is rayless, this chain has a maximal node $x \in T_{n+1}$. Let T' be the union of T_{n+1} and an x - u path P with $\overset{\circ}{P} \subseteq C$. Then the neighbourhood in T' of any new component $C' \subseteq C$ of $G - T'$ is a chain in T' , so T' is again normal. But then T' contradicts the maximality of T_{n+1} , completing the proof that T_{n+1} is as desired. \square

Corollary 11.2.2 (Jung). *Let G be any graph and let $U \subseteq V(G)$ be any vertex set. If U is dispersed itself and v is any vertex of G , then G contains a rayless normal tree that contains U cofinally and is rooted in v .* \square

Corollary 11.2.3 (Jung). *Let G be any graph and let $U \subseteq V(G)$ be any vertex set. If U is countable and v is any vertex of G , then G contains a normal tree that contains U cofinally and is rooted in v .* \square

Lemma 11.2.4. *Let G be any graph. The vertex set of any rayless normal tree $T \subseteq G$ is dispersed. In particular, the levels of any normal tree $T \subseteq G$ are dispersed.*

Proof. Lemma 11.1.11. \square

Jung's abstract characterisation of the normally spanned graphs goes as follows:

Theorem 11.2.5 (Jung, [43, Satz 6]). *Let G be any graph. A vertex set $W \subseteq V(G)$ is normally spanned in G if and only if it is a countable union of dispersed sets. In particular, G is normally spanned if and only if $V(G)$ is a countable union of dispersed sets.*

For an excluded-minor characterisation of the connected graphs with normal spanning trees see Diestel and Leader's [29].

Proof of Theorem 11.2.5. The backward implication holds by Proposition 11.2.1. The forward implication holds as the levels of any normal tree are dispersed, Lemma 11.2.4. \square

We are now ready to prove Theorem 11.1:

Proof of Theorem 11.1. First, to show that at most one of (i) and (ii) holds, we show (ii) \rightarrow \neg (i). If $T \subseteq G$ is a rayless normal tree containing U , then $V(T)$ is dispersed by Lemma 11.2.4, and hence so is $U \subseteq V(T)$.

It remains to show that at least one of (i) and (ii) holds; we show \neg (i) \rightarrow (ii). Since the absence of a comb with all its teeth in U means that U is dispersed, Corollary 11.2.2 yields a rayless normal tree in G that contains U cofinally. \square

11.2.2. Tree-decompositions

In this section, we show how the rayless normal tree from Theorem 11.1 gives rise to a tree-decomposition that is complementary to combs.

Theorem 11.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a comb attached to U ;
- (ii) G has a rayless tree-decomposition into parts each containing at most finitely many vertices from U and whose parts at non-leaves of the decomposition tree are all finite.

Moreover, the rayless tree-decomposition in (ii) displays $\partial_{\Omega}U$ and may be chosen with connected separators.

We start with a lemma which shows that at most one of (i) and (ii) holds.

Lemma 11.2.6. *Let G be any graph and let $U \subseteq V(G)$ be any vertex set. Suppose that G has a rayless tree-decomposition into parts each containing at most finitely many vertices from U and whose parts at non-leaves of the decomposition tree are all finite. Then for every infinite $U' \subseteq U$ there is a critical vertex set of G that lies in the closure of U' .*

Proof. Let such a tree-decomposition (T, \mathcal{V}) of G be given for U , and let U' be an arbitrary infinite subset of U . For every $u \in U'$ we choose a node $t_u \in T$ with $u \in V_{t_u}$. Since each part of the tree-decomposition contains at most finitely many vertices from U , we may assume without loss of generality (moving to an infinite subset of U') that the nodes t_u are pairwise distinct. Hence applying Lemma 11.1.4 in the rayless tree T yields a star S attached to $\{t_u \mid u \in U'\}$. Without loss of generality (as before) we may assume that the nodes t_u form precisely the attachment set of S and that no vertex u from U' is contained in the finite part V_c at the central node c of $S \subseteq T$. For every $u \in U'$ let C_u be the component of $G - V_c$ containing u . Then distinct vertices from U' are contained in distinct components of $G - V_c$ by Lemma 11.1.14. Since the finite part V_c contains the neighbourhood of each component C_u , by pigeon-hole principle we find a subset $X \subseteq V_c$ which is precisely equal to the neighbourhood of C_u for some infinitely many $u \in U'$. \square

Proof of Theorem 11.2. By Lemma 11.2.6 at most one of (i) and (ii) holds. It remains to show that at least one of (i) and (ii) holds.

We show $\neg(\text{i}) \rightarrow (\text{ii})$. Let $T_{\text{NT}} \subseteq G$ be a rayless normal tree containing U as provided by Theorem 11.1. We construct the desired tree-decomposition from T_{NT} . As T_{NT} is rayless and normal, the neighbourhood of any component C of $G - T_{\text{NT}}$ is a finite chain in the tree-order of T_{NT} , and hence has a maximal element $t_C \in T_{\text{NT}}$. Now, let the tree T be obtained from T_{NT} by adding each component C of $G - T_{\text{NT}}$ as a new vertex and joining it precisely to t_C . The tree T will be our decomposition tree; it remains to name the parts. For nodes $t \in T_{\text{NT}} \subseteq T$ we let V_t consist of the down-closure $[t]_{T_{\text{NT}}}$ of t in the normal tree T_{NT} . And for newly added nodes $C \in T - T_{\text{NT}}$ we let V_C be the union of V_{t_C} and the vertex set of the component C , i.e., we put $V_C = [t_C]_{T_{\text{NT}}} \cup V(C)$. It is straightforward to check that T with these parts forms a tree-decomposition of G that meets the requirements of (ii) and satisfies the theorem's 'moreover' part. \square

Our next example shows that Theorem 11.2 (ii) cannot be strengthened so as to get a star as decomposition tree or to have pairwise disjoint separators:

Example 11.2.7. Suppose that G consists of the first three levels of T_{\aleph_0} , the tree all whose vertices have countably infinite degree, and let $U = V(G)$. Then G is rayless so there is no comb attached to U .

First, G has no star-decomposition into parts each containing at most finitely many vertices from U : Indeed, assume for a contradiction that G has such a star-decomposition (S, \mathcal{V}) , and let c be the centre of the infinite star S . As the part V_c contains at most finitely many vertices from $U = V(G)$ it must be finite. Then each component of $G - V_c$ is contained in some $G[V_\ell]$ with ℓ a leaf of S by Corollary 11.1.15. As each part of (S, \mathcal{V}) contains at most finitely many vertices from U , this means that every component of $G - V_c$ contains at most finitely many vertices from $U = V(G)$ and hence is finite. But as V_c is finite, $G - V_c$ must have an infinite component, a contradiction.

Second, G also has no rayless tree-decomposition with finite and pairwise disjoint separators such that each part contains at most finitely many vertices from U : Indeed, suppose for a contradiction that G has such a tree-decomposition (T, \mathcal{V}) . Without loss of generality we may assume that all its parts are non-empty. The rayless decomposition tree T has a vertex t of infinite degree, so V_t contains infinitely many of the finite and pairwise disjoint separators. As G is connected, all of these are non-empty by Lemma 11.1.14, so V_t is infinite, and hence so is $V_t \cap U = V_t$. But this contradicts our assumptions.

11.2.3. Critical vertex sets

The absence of a comb attached to U is equivalent to U being dispersed, which is to say that no end of G lies in the closure of U . With the combinatorial remainder $\Gamma(G) = \Omega(G) \sqcup \text{crit}(G)$ compactifying G in mind, this means that only critical vertex sets of G lie in the closure of U , i.e. $\partial_\Gamma U \subseteq \text{crit}(G)$. Phrasing this combinatorially gives

Theorem 11.3. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a comb attached to U ;
- (ii) for every infinite $U' \subseteq U$ there is a critical vertex set $X \subseteq V(G)$ such that infinitely many of the components in \mathcal{C}_X meet U' .

Quantifying over all U' in Theorem 11.3 is necessary for (ii) \rightarrow \neg (i), e.g., if G is an infinite star of rays with $U = V(G)$. We remark that Theorem 11.3 implies Halin's [38, Satz 1] from 1965 which reads as follows: *A graph G is rayless if and*

only if every infinite $M \subseteq V(G)$ has an infinite subset M' for which there is a finite $H \subseteq G$ such that every component of $G - H$ contains only finitely many vertices of M' .

Since, by now, the right tools are at hand, we can prove Theorem 11.3 in two efficient ways:

Combinatorial proof of Theorem 11.3 using Theorem 11.1 or 11.2. Evidently, at most one of (i) and (ii) can hold. And if G contains no comb attached to U , then (ii) holds by Theorem 11.1 with Lemma 11.1.4 or by Theorem 11.2 with Lemma 11.2.6. □

Inverse limit proof of Theorem 11.3. Lemma 11.1.9 gives $\neg(\text{i}) \rightarrow (\text{ii})$. □

Note that condition (ii) yields a star attached to U .

11.2.4. Rank

In 1983, Schmidt [64] introduced a notion that is now known as the *rank* of a graph, cf. Chapter 8.5 of [20]. His rank provides a recursive definition of the class of rayless graphs which enables us to prove assertions about rayless graphs by transfinite induction. An outstanding application of this technique is the proof of the unfriendly partition conjecture for rayless graphs, cf. [3, 20]. Since the absence of a comb attached to U is equivalent to the existence of a *rayless* normal tree containing U , Theorem 11.1, one may wonder whether there somehow is a link to Schmidt's rank. In this section we show that this is indeed the case.

Schmidt defines the rank of a graph as follows. He assigns *rank* 0 to all the finite graphs. And given an ordinal $\alpha > 0$, he assigns *rank* α to every (not necessarily connected) graph G that does not already have a rank $\beta < \alpha$ and which has a finite set X of vertices such that every component of $G - X$ has some rank $< \alpha$.

Lemma 11.2.8 ([20, Lemma 8.5.2]). *Let G be any graph. Then the following assertions are complementary:*

- (i) G contains a ray;
- (ii) G has a rank.

Now we introduce the notion of a U -rank, based on Schmidt's rank, which additionally takes into account a fixed set U . For this, suppose that U is any set. Even though, formally, U is an arbitrary set, we think of U as a set of vertices. Let us assign U -rank 0 to all the graphs that contain at most finitely many vertices from U . Given an ordinal $\alpha > 0$, we assign U -rank α to every graph G that does not already have a U -rank $\beta < \alpha$ and which has a finite set X of vertices such that every component of $G - X$ has some U -rank $< \alpha$. Note that the rank of G is equal to the V -rank of G .

The U -rank behaves quite similar to Schmidt's rank, [20, p. 243]: When disjoint graphs G_i have U -ranks $\alpha_i < \alpha$, their union clearly has a U -rank of at most α ; if the union is finite, it has U -rank $\max_i \alpha_i$. Induction on α shows that subgraphs of graphs of U -rank α also have a U -rank of at most α . Conversely, joining finitely many new vertices to a graph, no matter how, will not change its U -rank.

Not every graph has a U -rank. Indeed, a comb attached to U cannot have a U -rank, since deleting finitely many of its vertices always leaves a component that is a comb attached to U . As subgraphs of graphs with a U -rank also have a U -rank, this means that only graphs without such combs can have a U -rank. But all these do:

Theorem 11.4. *Let G be any graph and let U be any set. Then the following assertions are complementary:*

- (i) G contains a comb attached to U ;
- (ii) G has a U -rank.

Phrased differently, the U -rank provides a recursive definition of the class of the graphs in which U is dispersed.

Proof of Theorem 11.4. We show the equivalence (i) \leftrightarrow \neg (ii). The forward implication has already been pointed out above. For the backward implication suppose that G has no U -rank; we show that G must contain a comb attached to U . As G has no U -rank, one of its components, C_0 say, has no U -rank as well. Pick $u_0 \in U \cap C_0$ arbitrarily. Since C_0 has no U -rank, it follows that $C_0 - u_0$ has a component C_1 that has no U -rank; let $u_1 \in U \cap C_1$ and pick a u_0 - u_1 path P_1 in C_0 with $\overset{\circ}{P}_1 \subseteq C_1$. Next, delete P_1 from C_1 and let $C_2 \subseteq C_1 - P_1$ be a component that has no U -rank. Let $u_2 \in U \cap C_2$, pick any P_1 - u_2 path P_2 in C_1 with $\overset{\circ}{P}_2 \subseteq C_2$

and note that P_2 meets P_1 in $\dot{u}_1 P_1$. Therefore, if we continue inductively to find paths P_1, P_2, \dots in G , then their union $\bigcup_n P_n$ is a comb with attachment set $\{u_n \mid n \in \mathbb{N}\} \subseteq U$. \square

There is a way to see immediately that for a connected graph G having a U -rank is stronger than G containing a star attached to U when U is infinite. For this, suppose that G has U -rank α . Then $\alpha > 0$ as $U \subseteq V(G)$ is infinite. Hence G has a finite set X of vertices such that every component of $G - X$ has some U -rank $< \alpha$. In particular, $G - X$ must have some infinitely many components that meet U . Each of these components gives some $U-X$ path avoiding all other components, so the pigeon-hole principle yields a star attached to U as desired.

The U -rank of a graph has many properties. In the remainder of this section, we prove three such properties that we will put to use in the next section.

Lemma 11.2.9. *Let G be any graph, let U be any set and suppose that G has U -rank α . Then the following assertions hold:*

- (i) *for every subset $U' \subseteq U$ the graph G has U' -rank $\leq \alpha$;*
- (ii) *for every subgraph $H \subseteq G$ the graph H has U -rank $\leq \alpha$.*

Proof. Induction on α . \square

Lemma 11.2.10. *Let U be any set. If T is a rooted rayless tree containing $U \cap V(T)$ cofinally, then the U -rank of T is equal to the rank of T .*

Proof. Let α be the U -rank of T and let β be the rank of T . Since the $V(T)$ -rank of T is the same as the rank of T , Lemma 11.2.9 (i) gives the inequality $\alpha \leq \beta$. An induction on α shows the converse inequality (in the induction step consider a set $X \subseteq V(T)$ witnessing that T has U -rank α and employ the induction hypothesis to see that every component of $T - X$ has rank $< \alpha$; it is convenient to assume X to be down-closed, which is possible by Lemma 11.2.9 (ii)). \square

Lemma 11.2.11. *If G is any graph and $T \subseteq G$ is a rayless normal tree containing $U \cap G$ cofinally, then the following three ordinals are all equal:*

- (i) *the rank of T ;*
- (ii) *the U -rank of T ;*
- (iii) *the U -rank of G .*

Proof. The equality (i) = (ii) is the subject of Lemma 11.2.10. Lemma 11.2.9 gives the inequality (ii) \leq (iii). We show the remaining inequality (iii) \leq (ii) by induction on the U -rank of T , as follows.

If the U -rank of T is 0, then $U \cap T = U \cap G$ is finite, and thus the U -rank of G is 0 as well. For the induction step, suppose that T has U -rank $\alpha > 0$, and let $X \subseteq V(T)$ be any finite vertex set such that every component of $T - X$ has U -rank $< \alpha$. By Lemma 11.2.9 (ii) we may assume that X is down-closed in T . It suffices to show that every component of $G - X$ has a U -rank $< \alpha$.

If C is a component of $G - X$, then either C avoids $T \supseteq U \cap C$ and has U -rank $0 < \alpha$, or C meets T . In the case that C meets T , by Lemma 11.1.10 we know that C is spanned by $\llbracket x \rrbracket$ with x minimal in $T - X$, so $T \cap C \subseteq C$ is a normal tree containing $U \cap C$ cofinally. Finally, by the induction hypothesis,

$$(U\text{-rank of } C) \leq (U\text{-rank of } T \cap C) < \alpha.$$

□

11.2.5. Combining the duality theorems

So far we have seen duality theorems for combs in terms of normal trees, tree-decompositions, critical vertex sets and rank. With these four complementary structures for combs at hand, the question arises whether it is possible to combine them all. In this section we will answer the question in the affirmative. That is, we will present a fifth complementary structure for combs that combines all of the four above.

This fifth structure will be a tree-decomposition that is more specific than the one listed above. It will stem from a normal tree in a way that we call ‘squeezed expansion’. Just like the tree-decomposition listed above, all its parts will meet U finitely, and all its parts at non-leaves will be finite. Moreover, it will display not only the ends in the closure of U , but also the critical vertex sets in the closure of U . In order to realise this, we will extend the definition of ‘display’ in a reasonable way. Finally, the decomposition tree will have a rank that is equal to the U -rank of the whole graph. The combined duality theorem reads as follows:

Theorem 11.5. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a comb attached to U ;
- (ii) G has a rooted tame tree-decomposition (T, \mathcal{V}) that covers U cofinally and satisfies the following four assertions:
 - (T, \mathcal{V}) is the squeezed expansion of a normal tree in G that contains the vertex set U cofinally;
 - every part of (T, \mathcal{V}) meets U finitely and parts at non-leaves are finite;
 - (T, \mathcal{V}) displays $\partial_T U \subseteq \text{crit}(G)$;
 - the rank of T is equal to the U -rank of G .

Corollary 11.2.12. *If a connected graph G is rayless (equivalently: if G has a rank), then G has a tame tree-decomposition into finite parts that displays the combinatorial remainder of G and has a decomposition tree whose rank is equal to the rank of G . □*

Here we remark that, in this chapter, we consider Schmidt’s ranking of rayless graphs as discussed in Section 11.2.4. In particular, when we consider the rank of a (possibly rooted) tree, we do not mean the rank for rooted trees that defines recursive prunability (cf. [20, p. 242 & 243]).

The proof of the theorem above is organised as follows. First, we will state Proposition 11.2.13, which lists some useful properties of squeezed expansions. Then, we will employ this proposition in a high level proof of Theorem 11.5. In order to follow the line of argumentation up to here, it is not necessary to know the definitions of ‘display’ and ‘squeezed’ ‘expansion’, which is why we will introduce them subsequently to our high level proof. Finally, we will prove Proposition 11.2.13.

Proposition 11.2.13. *Let G be any graph and suppose that $T_{\text{NT}} \subseteq G$ is a normal tree such that every component of $G - T_{\text{NT}}$ has finite neighbourhood, that (T, \mathcal{V}) is the expansion of T_{NT} and that (T', \mathcal{W}) is a squeezed (T, \mathcal{V}) . Then the following assertions hold:*

- (i) (T, \mathcal{V}) is upwards connected;
- (ii) both (T, \mathcal{V}) and (T', \mathcal{W}) display $\partial_{\Gamma} T_{\text{NT}}$;
- (iii) all the parts of (T, \mathcal{V}) and (T', \mathcal{W}) meet T_{NT} finitely;
- (iv) parts of (T', \mathcal{W}) at non-leaves of T' are finite;
- (v) T' is rayless if and only if T is rayless if and only if T_{NT} is rayless;
- (vi) if one of T' , T and T_{NT} is rayless, then the ranks of T' , T and T_{NT} all exist and are all equal.

The proposition has a corollary that is immediate because every normal spanning tree will have an expansion, and expansions will be rooted and tame:

Corollary 11.2.14. *Every normally spanned graph has a rooted tame tree-decomposition displaying its combinatorial remainder. \square*

Now we prove Theorem 11.5 using Proposition 11.2.13 above:

Proof of Theorem 11.5. (i) and (ii) exclude each other for various reasons we have already discussed.

For the implication $\neg(\text{i}) \rightarrow (\text{ii})$ suppose that G contains no comb attached to U . By Theorem 11.1 there is a rayless normal tree $T_{\text{NT}} \subseteq G$ that contains U cofinally. We show that the squeezed expansion (T, \mathcal{V}) of T_{NT} is as desired. By Proposition 11.2.13 every part of (T, \mathcal{V}) meets $T_{\text{NT}} \supseteq U$ finitely and parts at non-leaves of T are finite. As we have $\partial_{\Gamma} T_{\text{NT}} = \partial_{\Gamma} U$ by Lemma 11.1.13, Proposition 11.2.13 also ensures that the squeezed expansion (T, \mathcal{V}) of T_{NT} displays $\partial_{\Gamma} U$. Finally, the U -rank of G exists by Theorem 11.4 and is equal to the rank of T_{NT} by Lemma 11.2.11, which in turn is equal to the rank of T by Proposition 11.2.13. \square

Next, we provide all the definitions needed: First, we extend the definition of ‘display’ to include critical vertex sets (Definition 11.2.16). Second, we define

the ‘expansion’ of a normal tree (Definition 11.2.17), which is a certain tree-decomposition. Finally we define what it means to ‘squeeze’ a tree-decomposition (Definition 11.2.18).

Recall that the definition of ‘display’, as discussed in Section 11.1, highly relies on the fact that the ends of a graph orient all its finite-order separations. Now, critical vertex sets are closely related to ends, as they together with the ends turn graphs into compact topological spaces. This is why we may hope that every critical vertex set X orients the finite-order separations so as to lead immediately to a notion of ‘displaying a collection of critical vertex sets’. Probably the most natural way how a critical vertex set X could orient a finite-order separation $\{A, B\}$ towards a side $K \in \{A, B\}$ is that X together with all but finitely many of the components in $\check{\mathcal{C}}_X$ are contained in K .

However, this is too much to ask: For example consider an infinite star. The centre c of the star forms a critical vertex set $X = \{c\}$, and any separation with separator X that has infinitely many leaves on both sides will not be oriented by X in this way.

But focusing on a suitable class of separations, those that are *tame*, leads to a natural extension of ‘display’ to include critical vertex sets: A finite-order separation $\{X, \mathcal{C}\}$ of G is *tame* if for no $Y \subseteq X$ both \mathcal{C} and $\check{\mathcal{C}}_X \setminus \mathcal{C}$ contain infinitely many components whose neighbourhoods are precisely equal to Y . The tame separations of G are precisely those finite-order separations of G that respect the critical vertex sets:

Lemma 11.2.15. *A finite-order separation $\{A, B\}$ of a graph G is tame if and only if every critical vertex set X of G together with all but finitely many components from $\check{\mathcal{C}}_X$ is contained in one side of $\{A, B\}$.*

Proof. For the forward implication, note that every distinct two vertices of a critical vertex set are linked in $G[X \cup \check{\mathcal{C}}_X]$ by infinitely many independent paths, so every critical vertex set of G meets at most one component of $G - (A \cap B)$. \square

We say that an S_{\aleph_0} -tree (T, α) is *tame* if all the separations in the image of α are tame. And we say that a tree-decomposition is *tame* if it corresponds to a tame S_{\aleph_0} -tree.

If X is a critical vertex set of G and (T, α) is a tame S_{\aleph_0} -tree, then X induces a consistent orientation of the image of α by orienting every tame finite-order

separation $\{A, B\}$ towards the side that contains X and all but finitely many of the components from $\check{\mathcal{C}}_X$ (cf. Lemma 11.2.15 above). This consistent orientation also induces a consistent orientation of $\vec{E}(T)$ via α . Then, just like for ends, the critical vertex set X either *lives* at a unique node $t \in T$ or *corresponds* to a unique end of T . In this way, we obtain an extension $\sigma: \Gamma(G) \rightarrow \Omega(T) \sqcup V(T)$ of the map $\tau: \Omega(G) \rightarrow \Omega(T) \sqcup V(T)$ from Section 11.1.8.

Since σ extends τ from the end space $\Omega(G)$ of G to the full combinatorial remainder $\Gamma(G)$ of G , it is reasonable to wonder why the target set of σ is that of τ , namely $\Omega(T) \sqcup V(T)$, rather than analogously taking the target set $\Gamma(T) \sqcup V(T)$. At a closer look, the critical vertex sets of T are already contained in the target set $\Omega(T) \sqcup V(T)$, for they are precisely the infinite degree nodes of T . This, and the fact that every critical vertex set X of G naturally comes with an oriented tame separation $(X, \check{\mathcal{C}}_X)$ of G , motivate the following definition.

Definition 11.2.16. [Display $\Psi \subseteq \Gamma(G)$] Let G be any graph. A rooted tame S_{\aleph_0} -tree (T, α) *displays* a subset Ψ of the combinatorial remainder $\Gamma(G) = \Omega(G) \sqcup \text{crit}(G)$ of G if σ satisfies the following three conditions:

- σ restricts to a bijection between $\Psi \cap \Omega(G)$ and $\Omega(T)$;
- σ restricts to a bijection between $\Psi \cap \text{crit}(G)$ and the infinite-degree nodes of T so that: whenever σ sends a critical vertex set $X \in \Psi$ to $t \in T$, then t has a predecessor $s \in T$ with $\alpha(s, t) = (X, \mathcal{C})$ such that $\mathcal{C} \subseteq \check{\mathcal{C}}_X$ is cofinite and α restricts to a bijection between \vec{F}_t and the star in \vec{S}_{\aleph_0} that consists of the separation (X, \mathcal{C}) and all the separations (C, X) with $C \in \mathcal{C}$;
- σ sends all the elements of $\Gamma(G) \setminus \Psi$ to finite-degree nodes of T .

Note that this definition of displays is not exactly an extension of the original definition given in Section 11.1.8. Indeed, if (T, α) displays Ψ and $\omega \in \Psi$ is an end, then with the original definition ω may correspond to an infinite degree vertex of T , but not with the new definition. However, the new definition is stronger than the original one: if (T, α) displays $\Psi \subseteq \Gamma(G)$ in the new sense, then (T, α) displays $\Psi \cap \Omega(G)$ in the original sense.

We solve this ambiguity as follows. Whenever we say that a tree-decomposition or S_{\aleph_0} -tree displays some set Ψ of ends of G and it is clearly understood that we view Ψ as a subset of $\Omega(G)$, e.g. when we let Ψ consist of the undominated ends

of G or consider $\Psi = \partial_\Omega U$, then by ‘displays’ we refer to the original definition from Section 11.1.8. But whenever we explicitly introduce Ψ as a subset of the combinatorial remainder $\Gamma(G)$ of G , e.g. when we let Ψ consist of critical vertex sets or consider $\Psi = \partial_\Gamma U$, then by ‘displays’ we refer to the new definition introduced above.

We wish to make a few remarks on our new definition. If (T, α) is a rooted tame S_{\aleph_0} -tree displaying some $\Psi \subseteq \Gamma(G)$ and the tree-decomposition (T, \mathcal{V}) corresponding to (T, α) exists, then $V_{\sigma(X)} = X$ whenever X is a critical vertex set in Ψ . We do not require $\mathcal{C} = \check{\mathcal{C}}_X$ in the definition of displays because there are simply structured normally spanned graphs for which otherwise none of their tree-decompositions would display their combinatorial remainder. See [31, Examples 3.6 & 3.7] for details.

Now, let us turn to the expansion of a normal tree. Given vertex sets $Y \subseteq X \subseteq V(G)$ we write $\mathcal{C}_X(Y)$ for the collection of all components $C \in \mathcal{C}_X$ with $N(C) = Y$.

Definition 11.2.17 (Expansion of a normal tree). In order to define the expansion, suppose that G is any connected graph and $T_{\text{NT}} \subseteq G$ is any normal tree such that every component of $G - T_{\text{NT}}$ has finite neighbourhood. From the normal tree T_{NT} we obtain the *expansion* (T, \mathcal{V}) of T_{NT} in G in two steps, as follows.

For the first step, let us suppose without loss of generality that for all nodes $t \in T_{\text{NT}}$ every up-neighbour t' of t in T_{NT} is named as the component $\llbracket t' \rrbracket$ of $G - [t]$ containing t' . We define a map $\beta: \vec{E}(T_{\text{NT}}) \rightarrow \vec{S}_{\aleph_0}$ by letting $\beta(t, C) := (N(C), C)$ and $\beta(C, t) := \beta(t, C)^*$ whenever C is an up-neighbour of a node t in T_{NT} . Then (T_{NT}, β) is a rooted tame S_{\aleph_0} -tree that displays $\partial_\Omega T_{\text{NT}} \subseteq \Omega(G)$.

In the second step, we obtain from (T_{NT}, β) a rooted tame S_{\aleph_0} -tree (T, α) displaying $\partial_\Gamma T_{\text{NT}} \subseteq \Gamma(G)$. Informally speaking we sort the separations of the form $\beta(t, C)$ with $t \in T_{\text{NT}}$ an infinite degree-node and C an up-neighbour of t in T_{NT} by the critical vertex sets $X \subseteq [t]$ in the closure of T_{NT} with $C \in \check{\mathcal{C}}_X$. Formally this is done as follows (cf. Figure 11.2.1).

For every infinite-degree node $t \in T_{\text{NT}}$ and every critical vertex set $X \in \partial_\Gamma T_{\text{NT}}$ satisfying $t \in X \subseteq [t]$ we do the following:

- (i) we add a new vertex named X to T_{NT} and join it to t ;
- (ii) for every component $C \in \mathcal{C}_{[t]}(X) \subseteq \check{\mathcal{C}}_X$ we delete the edge tC (this is




Figure 11.2.1.: The second step in the construction of the expansion of normal trees. The critical vertex sets X and X' are in the closure of T_{NT} , while X'' is not. The three sets X , X' and X'' are all the critical vertex sets of G that contain t and are contained in $[t]$.

redundant when T_{NT} avoids C) and add the new edge XC (note that in particular the vertex C gets added as well, even if T_{NT} avoids C);

- (iii) we let $\alpha(t, X) := (X, \mathcal{C}_{[t]}(X))$, and for every component $C \in \mathcal{C}_{[t]}(X)$ we let $\alpha(X, C) := (X, C)$.

Then we take T to be the resulting tree, and we extend α to all of $\vec{E}(T)$ by letting $\alpha(\vec{e}) := \beta(\vec{e})$ whenever the edge e of T is also an edge of the normal tree T_{NT} . The rooted tame tree-decomposition (T, \mathcal{V}) corresponding to (T, α) is the *expansion* of T_{NT} in G . \diamond

And here is the definition of squeezing:

Definition 11.2.18 (Squeezing a tree-decomposition). Suppose that (T, \mathcal{V}) and (T', \mathcal{W}) are tree-decompositions of G . We say that (T', \mathcal{W}) is a *squeezed* (T, \mathcal{V}) if (T', \mathcal{W}) is obtained from (T, \mathcal{V}) as follows. The tree T' is obtained from T by adding, for every node $t \in T$ that has finite degree > 1 and whose part V_t is infinite, a new node t' to T and joining it to t . For all these nodes t the part W_t is the union of the separators of (T, \mathcal{V}) associated with the edges of T at t , and the part $W_{t'}$ is taken to be the part V_t . For all other nodes t the part W_t is V_t .

Note that if (T', \mathcal{W}) is the squeezed (T, \mathcal{V}) and all separators of (T, \mathcal{V}) are finite, then all the infinite parts V_t with t an internal finite-degree node of T become finite parts W_t . Thus, all parts W_t with t an internal finite-degree node of T' are finite. Achieving this property is the purpose of squeezing.

Squeezing preserves tameness:

Lemma 11.2.19. *Let G be any graph, let (T, \mathcal{V}) be any tree-decomposition of G with finite separators and let (T', \mathcal{W}) be the squeezed (T, \mathcal{V}) . If (T, \mathcal{V}) is tame, then (T', \mathcal{W}) is tame as well.*

Proof. Suppose that (T, \mathcal{V}) is a tame tree-decomposition of G and that (T', \mathcal{W}) is the squeezed (T, \mathcal{V}) . Separations of G that are induced by (T', \mathcal{W}) are tame when they are induced by edges of T' that are also edges of $T \subseteq T'$. Hence it suffices to show that for every leaf $\ell \in T' - T$ with neighbour $t \in T \subseteq T'$ the separation induced by $\ell t \in T'$ is tame. For this, let any edge $\ell t \in T'$ be given and write s_0, \dots, s_n for the finitely many neighbours of t in T . Let (T', α') be the S_{\aleph_0} -tree corresponding to (T', \mathcal{W}) , let $(A, B) := \alpha'(\ell, t)$ and define $(A_i, B_i) := \alpha'(t, s_i)$ for all $i \leq n$. Then, by the definition of (T', \mathcal{W}) , we have $A = \bigcap_i A_i$ and $B = \bigcup_i B_i$. Our aim is to show that the separation $\{A, B\}$ is tame. By Lemma 11.2.15 it suffices to show that for every critical vertex set X of G there is a cofinite subset $\mathcal{C} \subseteq \check{\mathcal{C}}_X$ such that either $G[X \cup \bigcup \mathcal{C}] \subseteq G[A]$ or $G[X \cup \bigcup \mathcal{C}] \subseteq G[B]$. For this, let any critical vertex set X of G be given.

The critical vertex set X lives at or correspond to the unique node or end $\sigma(X)$ of T with regard to (T, \mathcal{V}) because (T, \mathcal{V}) is tame. If $\sigma(X)$ is distinct from t , then there is a cofinite subset $\mathcal{C} \subseteq \check{\mathcal{C}}_X$ such that $G[X \cup \bigcup \mathcal{C}] \subseteq G[B_i]$ for some $i \leq n$, and $G[X \cup \bigcup \mathcal{C}] \subseteq G[B]$ follows as desired. Hence we may assume that $\sigma(X) = t$. Thus, for every $i \leq n$ there is a cofinite subset $\mathcal{C}(i) \subseteq \check{\mathcal{C}}_X$ such that $G[X \cup \bigcup \mathcal{C}(i)] \subseteq G[A_i]$. Then $G[X \cup \bigcup \mathcal{C}] \subseteq G[A]$ as desired for the cofinite subset $\mathcal{C} := \bigcap_i \mathcal{C}(i) \subseteq \check{\mathcal{C}}_X$. \square

Now that we have formally introduced all the definitions involved, we are ready to prove Proposition 11.2.13:

Proof of Proposition 11.2.13. (i) The expansion is upwards connected by its definition.

(ii) Using Lemma 11.1.11 and the fact that every component of $G - T_{\text{NT}}$ has finite neighbourhood, it is straightforward to check that the tree-decomposition (T, \mathcal{V}) displays $\partial_{\Omega} T_{\text{NT}} \subseteq \Omega(G)$. We verify that (T, \mathcal{V}) even displays $\partial_{\Gamma} T_{\text{NT}} \subseteq \Gamma(G)$. On the one hand, by Lemma 11.1.12 every critical vertex set $X \in \partial_{\Gamma} T_{\text{NT}}$ is contained in T_{NT} as a chain, and hence appears precisely once as a node of T by the definition of the expansion. On the other hand, every node of infinite degree of T stems from such a critical vertex set. Together we obtain that (T, \mathcal{V}) displays $\partial_{\Gamma} T_{\text{NT}}$. The

tree-decomposition (T', \mathcal{W}) is tame because (T, \mathcal{V}) is, cf. Lemma 11.2.19. From here, it is straightforward to show that (T', \mathcal{W}) displays $\partial_\Gamma T_{\text{NT}}$ as well.

(iii) and (iv) are straightforward.

(v) follows from (ii) and Lemma 11.1.11.

(vi) It is straightforward to check by induction on the rank that the rank is preserved under taking contraction minors with finite branch sets. Similarly, one can show that two infinite trees have the same rank if one is obtained from the other by adding new leaves to some of its nodes of infinite degree. Now we deduce (vi) as follows. For every node $t \in T_{\text{NT}}$ let us write S_t for the finite star with centre t and leaves the critical vertex sets $X \in \partial_\Gamma T_{\text{NT}}$ with $t \in X \subseteq [t]$. The decomposition tree T of the expansion of T_{NT} is obtained from an $IT_{\text{NT}} \subseteq T$ with finite branch sets (the non-trivial branch sets are precisely the vertex sets of the stars S_t for the nodes $t \in T_{\text{NT}}$ of infinite degree) by adding leaves to nodes of infinite degree (each leaf is a component $C \in \mathcal{C}_{[t]}(X)$ avoiding T_{NT} for some $X \in S_t$ and gets joined to $X \in IT_{\text{NT}} \subseteq T$). Therefore, the ranks of T and T_{NT} coincide. The decomposition tree T' is obtained from T by adding at most one new leaf to each node of T , and new leaves are only added to finite-degree nodes of T . An induction on the rank shows that the rank is preserved under this operation, and so the ranks of T' and T coincide as well. \square

Carmesin [15] showed that every connected graph G has a tree-decomposition with finite separators that displays Ψ for Ψ the set undominated ends of G , cf. Theorem 11.1.17. He then asked for a characterisation of those pairs of a graph G and a subset $\Psi \subseteq \Omega(G)$ for which G has such a tree-decomposition displaying Ψ . In the same spirit, our findings motivate the following problem:

Problem 11.2.20. *Characterise, for all connected graphs G , the subsets $\Psi \subseteq \Gamma(G)$ for which G admits a rooted tame tree-decomposition displaying Ψ .*

11.3. Stars

11.3.1. Normal trees

In this section we prove a duality theorem for stars in terms of normal trees.

Theorem 11.6. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a star attached to U ;
- (ii) there is a locally finite normal tree $T \subseteq G$ that contains U and all whose rays are undominated in G .

Moreover, the normal tree T in (ii) can be chosen such that it contains U cofinally and every component of $G - T$ has finite neighbourhood.

Proof of Theorem 11.6 without the ‘moreover’ part. First, we show that at most one of (i) and (ii) holds. Assume for a contradiction that both hold. Let $T \subseteq G$ be a normal tree as in (ii) and let $U' \subseteq U$ form the attachment set of some star attached to U . By Lemma 11.1.3 the locally finite tree T contains a comb attached to U' . That comb’s spine, then, is dominated in G by the centre of the star, a contradiction.

It remains to show that at least one of (i) and (ii) holds; we show $\neg(\text{i}) \rightarrow (\text{ii})$. We have that U is countable, since otherwise the star-comb lemma yields a star attached to U . By Corollary 11.2.3 we find a normal tree $T \subseteq G$ that contains U cofinally. Clearly, T must be locally finite since G contains no star attached to U . For the same reason, every ray of T is undominated in G . \square

The remaining ‘moreover’ part is a consequence of Theorem 12.1, which is why its proof is placed in Chapter 12. To see immediately that a locally finite normal tree T as in (ii) is more specific than a comb when U is infinite, apply Lemma 11.1.3 to T .

11.3.2. Tree-decompositions

For combs we have provided a duality theorem in terms of normal trees, and that theorem then gave rise to another duality theorem in terms of tree-decompositions. Since we have shown a duality theorem for stars in terms of normal trees in the

previous section, a natural question to ask is whether this theorem gives rise to a duality theorem for stars in terms of tree-decompositions, just like for combs. It turns out that stars have a duality theorem in terms of tree-decompositions. But this theorem cannot be proved by imitating the proof of the respective theorem for combs, and so we will have to come up with a whole new strategy. Our theorem reads as follows:

Theorem 11.7. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a star attached to U ;
- (ii) G has a locally finite tree-decomposition with finite and pairwise disjoint separators such that each part contains at most finitely many vertices of U .

Moreover, the tree-decomposition in (ii) can be chosen with connected separators and such that it displays $\partial_\Gamma U$ which consists only of ends.

We remark that (ii) is equivalent to the assertion that ‘ G has a ray-decomposition with finite and pairwise disjoint separators such that each part contains at most finitely many vertices of U ’ since the distance classes of locally finite trees are finite.

To see that a tree-decomposition as in (ii) is more specific than a comb, start with a ray in the decomposition tree (cf. Lemma 11.1.3) and then inductively construct a comb in the connected parts along that ray. To see that a locally finite tree-decomposition (T, \mathcal{V}) as in (ii) is more specific than a comb attached to U , consider the nodes of T whose parts meet U and apply Lemma 11.1.4 in T to find a comb C attached to them. Then inductively construct a comb in G attached to U working inside the connected parts along $C \subseteq T$.

To prove the theorem, we start by showing that (i) and (ii) exclude each other:

Lemma 11.3.1. *In Theorem 11.7 the graph G cannot satisfy both (i) and (ii).*

Proof. Let (T, \mathcal{V}) be a tree-decomposition as in (ii) of Theorem 11.7. Assume for a contradiction that G contains a star S attached to U . As the separators of (T, \mathcal{V}) are pairwise disjoint, the centre c of S is contained in at most two parts of (T, \mathcal{V}) . Let $T' \subseteq T$ be the finite subtree induced by the nodes of these parts plus their neighbours in T . As the parts at the nodes of T' altogether contain at

most finitely many vertices from U , the star S must send infinitely many paths to vertices in parts at $T - T'$. But the centre c is separated from the parts at $T - T'$ by the finite union of the finite separators associated with the edges of T leaving T' , a contradiction. \square

Now, to prove Theorem 11.7 it remains to show $\neg(\text{i}) \rightarrow (\text{ii})$. This time, however, it is harder to see how the normal tree from Theorem 11.6 can be employed to yield a tree-decomposition. That is why we do not take the detour via normal trees and instead construct the tree-decomposition directly. Still, this requires some effort.

First of all, assuming the absence of a star as in (i), we need a strategy to construct a tree-decomposition as in (ii). Fortunately, we do not have to start from scratch. In the proof of [27, Theorem 2.2], Diestel and Kühn proved the following as a technical key result: *If ω is an undominated end of G , then there exists a sequence $(X_n)_{n \in \mathbb{N}}$ of non-empty finite vertex sets $X_n \subseteq V(G)$ such that, for all $n \in \mathbb{N}$, the component $C(X_n, \omega)$ contains $X_{n+1} \cup C(X_{n+1}, \omega)$.* Now if $\partial_\Omega U$ is a singleton $\{\omega\}$, then ω must be undominated as (i) fails, and we consider such a sequence $(X_n)_{n \in \mathbb{N}}$. By making all the X_{n+1} connected in $C(X_n, \omega)$ first, and then moving to a suitable subsequence, we obtain a ray-decomposition of G that meets the requirements of (ii). Our strategy is to generalise this fundamental observation using that $\partial_\Omega U$ is compact in our situation:

Lemma 11.3.2. *If G contains no star attached to U , then $\partial_\Omega U$ is non-empty, compact and contains only undominated ends.*

Proof. By the pigeonhole principle, for every $X \in \mathcal{X}$ only finitely many components of $G - X$ may meet U . Thus we have that $\partial_\Omega U$ is non-empty and compact by Lemma 11.1.8. \square

Our next lemma generalises the fact that a vertex can be strictly separated from every end which it does not dominate.

Lemma 11.3.3. *Suppose that X is a finite set of vertices in a (possibly disconnected) graph G such that $G - X$ is connected, and that $\Psi \subseteq \Omega(G)$ is a non-empty and compact subspace consisting only of undominated ends. Then there is a finite-order separation of G that strictly separates X from Ψ and whose separator is connected.*

Proof. No end in Ψ is dominated and X is finite, so for every end $\omega \in \Psi$ we find a finite vertex set $Y(\omega) \subseteq V(G)$ with $Y(\omega) \cup C(Y(\omega), \omega)$ disjoint from X . Since the components $C(Y(\omega), \omega)$ induce a covering of Ψ by open sets, the compactness of Ψ yields finitely many ends $\omega_1, \dots, \omega_n \in \Psi$ such that every end in Ψ lives in at least one of the components $C(Y(\omega_i), \omega_i)$. Let the vertex set Y be obtained from the finite union of the finite sets $Y(\omega_i)$ by adding some finitely many vertices from the connected subgraph $G - X$ so as to ensure that $G[Y]$ is connected. Note that Y avoids X , and write \mathcal{D} for the collection of the components of $G - Y$ in which ends of Ψ live. We claim that (Y, \mathcal{D}) strictly separates X from Ψ . For this, let ω be any end in Ψ . Pick an index k for which ω lives in the component $C(Y(\omega_k), \omega_k) =: C$. Then, by the choice of $Y(\omega_k)$, there is no X - C path in $G - Y(\omega_k)$. By $Y(\omega_k) \subseteq Y$ and $C(Y, \omega) \subseteq C$ then there certainly is no X - $C(Y, \omega)$ path in $G - Y$. Therefore, (Y, \mathcal{D}) strictly separates X from Ψ . \square

Proposition 11.3.4. *Let G be any connected graph and suppose that $\Psi \subseteq \Omega$ is a non-empty and compact subspace that consists only of undominated ends. Then there exists a locally finite S_{\aleph_0} -tree (T, α) with connected pairwise disjoint separators that displays Ψ .*

Proof. We inductively construct a sequence $((T_n, \alpha_n))_{n \in \mathbb{N}}$ of rooted S_{\aleph_0} -trees with root $r \in T_0 \subseteq T_1 \subseteq \dots$ and $\alpha_0 \subseteq \alpha_1 \subseteq \dots$, as follows.

To define (T_0, α_0) , let T_0 consist of one edge rt and put $\alpha_0(r, t) := (\{v\}, V)$ for an arbitrary vertex v of G . Now, to obtain (T_{n+1}, α_{n+1}) from (T_n, α_n) , we do the following for every edge $t\ell$ of T_n at a leaf $\ell \neq r$. Consider the separation $\alpha(t, \ell) = (X, \mathcal{C})$ with C_1, \dots, C_n the finitely many components in \mathcal{C} in which ends of Ψ live (these are finitely many as Ψ is compact). For each component C_i apply Lemma 11.3.3 in $G[X \cup C_i]$ to X and $\Psi \cap \partial_\Omega C_i$ to obtain a finite-order separation (A_i, B_i) of $G[X \cup C_i]$ that strictly separates X from $\Psi \cap \partial_\Omega C_i$ in $G[X \cup C_i]$ and has a connected separator $A_i \cap B_i$. Then (A'_i, B'_i) with $A'_i := A_i \cup (V \setminus C_i)$ and $B'_i := B_i$ is a finite-order separation of G that strictly separates X from $\Psi \cap \partial_\Omega C_i$ in G and has a connected separator $A'_i \cap B'_i = A_i \cap B_i$. We add each C_i as a new node to T_n , join it precisely to the leaf ℓ and let $\alpha_{n+1}(\ell, C_i) := (A'_i, B'_i)$. This completes the description of our construction.

We claim that the pair (T, α) given by $T := \bigcup_n T_n$ and $\alpha := \bigcup_n \alpha_n$ is as required. Our construction ensures that T is locally finite and that the separators of (T, α)

are connected and pairwise disjoint. Furthermore, our construction ensures that every end in Ψ corresponds to an end of T . It remains to show that (T, α) displays Ψ . By Lemma 11.1.16 it suffices to show that, for every end of T , there is an end in Ψ corresponding to it. And indeed, every ray in T avoiding the root is, literally, a descending sequence $C_1 \supseteq C_2 \supseteq \dots$ of components for which some end of the compact Ψ lives in all C_n by the finite intersection property of the collection $\{\Psi \cap \partial_\Omega C_n \mid n \in \mathbb{N}\}$. \square

Proof of Theorem 11.7. By Lemma 11.3.1 at most one of (i) and (ii) can hold. To establish that at least one of them holds, we show $\neg(\text{i}) \rightarrow (\text{ii})$. Suppose that G contains no star attached to U . By Lemma 11.3.2 we know that the subspace $\partial_\Omega U \subseteq \Omega$ consisting of the ends lying in the closure of U actually contains only undominated ones, and is both non-empty and compact. Proposition 11.3.4 then yields a locally finite S_{\aleph_0} -tree (T, α) with connected pairwise disjoint separators that displays $\partial_\Omega U$. Let (T, \mathcal{V}) be the tree-decomposition corresponding to (T, α) . As G contains no star attached to U , there is no critical vertex set in the closure of U , and hence (T, \mathcal{V}) even displays $\partial_\Gamma U$. It remains to show that each part of (T, \mathcal{V}) contains at most finitely many vertices from U . Suppose for a contradiction that some part V_t contains some infinitely many vertices from U , and write U' for that subset of U . As (i) fails, applying Lemma 11.3.2 in G to U' yields an end in $\partial_\Omega U'$. But then this end lies in Ψ but does not correspond to an end of T , a contradiction. \square

12. Dominating stars and dominated combs

In the first chapter of this part we found structures whose existence is complementary to the existence of a star or a comb attached to a given set U of vertices. A *comb* is the union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R . The last vertices of those paths are the *teeth* of this comb. Given a vertex set U , a *comb attached to U* is a comb with all its teeth in U , and a *star attached to U* is a subdivided infinite star with all its leaves in U . Then the set of teeth is the *attachment set* of the comb, and the set of leaves is the *attachment set* of the star.

As stars and combs can interact with each other, this is not the end of the story. For example, a given vertex set U might be connected in a graph G by both a star and a comb, even with infinitely intersecting sets of leaves and teeth. To formalise this, let us say that a subdivided star S *dominates* a comb C if infinitely many of the leaves of S are also teeth of C . A *dominating star* in a graph G then is a subdivided star $S \subseteq G$ that dominates some comb $C \subseteq G$; and a *dominated comb* in G is a comb $C \subseteq G$ that is dominated by some subdivided star $S \subseteq G$. In this chapter we determine structures whose existence is complementary to the existence of dominating stars and dominated combs. Note that duality theorems for dominated combs are by nature also duality theorems for dominating stars, because for a graph G and a vertex set $U \subseteq V(G)$ the existence of a dominated comb attached to U is equivalent to the existence of a dominating star attached to U . For the sake of readability, we will state our duality theorems only for dominated combs.

Our first duality theorem for dominated combs is phrased in terms of normal trees. A rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T . A vertex v of G *dominates* a ray $R \subseteq G$ if there is an infinite v - $(R - v)$ fan in G . For example, a comb is dominated

in G if and only if its spine is dominated in G . Rays not dominated by any vertex are *undominated*. An end of G is *dominated* and *undominated* if one (equivalently: each) of its rays is dominated and undominated, respectively. (See Diestel’s textbook [20].)

Theorem 12.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *dominated comb* attached to U ;
- (ii) there is a *normal tree* $T \subseteq G$ that contains U and all whose rays are *undominated* in G .

Moreover, the normal tree T in (ii) can be chosen such that it contains U cofinally and every component of $G - T$ has finite neighbourhood.

When a graph contains no star or no comb attached to U , then in particular it contains no dominated comb attached to U . Hence, by our theorem, the graph contains a certain normal tree. If there is no star, then this normal tree will be locally finite; and if there is no comb, then it will be rayless. Therefore, our duality theorem for dominated combs in terms of normal trees implies our duality theorems for arbitrary stars and combs in terms of normal trees from Chapter 11, Theorems 12.1.1 and 12.1.2. This is surprising given that infinite trees cannot be locally finite and rayless at the same time.

As an application, we will partially generalise Diestel’s structural characterisation [24] of the graphs for which the topological spaces obtained by adding their ends are metrisable. Depending on the topology chosen, Diestel characterised these graphs in terms of normal spanning trees, dominated combs, and infinite stars. Applying Theorem 12.1, we can now provide, for any given set U of vertices, existence criteria for metrisable (standard) subspaces containing U in the various topologies. Our criteria will be in terms of normal trees containing U , dominated combs attached to U , and stars attached to U . For one of the topologies we obtain a characterisation.

Theorem 12.1 is significantly strengthened by its ‘moreover’ part. It will be needed in the proof of our second duality theorem for dominated combs which is phrased in terms of tree-decompositions. For the definition of tree-decompositions see [20]. ‘Essentially disjoint’ and ‘displaying’ are defined in Section 12.2. An end

ω of a graph G is contained *in the closure* of a vertex set $U \subseteq V(G)$ in G if G contains a comb attached to U whose spine lies in ω .

Theorem 12.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a *dominated comb* attached to U ;
- (ii) G has a *tree-decomposition* (T, \mathcal{V}) such that:
 - each part contains at most finitely many vertices from U ;
 - all parts at non-leaves of T are finite;
 - (T, \mathcal{V}) has essentially disjoint connected adhesion sets;
 - (T, \mathcal{V}) displays the ends of G in the closure of U in G .

Similar to Theorem 12.1, our duality theorem for dominated combs in terms of tree-decompositions implies our duality theorems for arbitrary stars and combs in terms of tree-decompositions from Chapter 11, Theorems 12.2.1 and 12.2.2.

In our proof of Theorem 12.2 we employ a profound theorem of Carmesin [15], which states that every graph has a tree-decomposition displaying all its undominated ends. As it will be the case in this paper, Carmesin’s theorem might often be used for graphs with normal spanning trees. For this particular case we provide a substantially shorter proof.

This chapter is organised as follows. Section 12.1 establishes our duality theorem for dominated combs in terms of normal trees. In Section 12.2 we prove our duality theorems for dominated combs in terms of tree-decompositions. Our short proof of Carmesin’s theorem for graphs with a normal spanning tree can be found there as well.

Throughout this chapter, $G = (V, E)$ is an arbitrary infinite graph. We assume familiarity with the tools and terminology described in the first chapter of this part.

12.1. Normal trees

In this section we obtain the following duality theorem for dominated combs in terms of normal trees:

Theorem 12.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) *G contains a dominated comb attached to U ;*
- (ii) *there is a normal tree $T \subseteq G$ that contains U and all whose rays are undominated in G .*

Moreover, the normal tree T in (ii) can be chosen such that it contains U cofinally and every component of $G - T$ has finite neighbourhood.

The inconspicuous ‘moreover’ part will pave the way for our duality theorem for dominated combs in terms of tree-decompositions (Theorem 12.2).

Before we provide a proof of Theorem 12.1 above, we shall discuss some consequences and applications. As a first consequence, Theorem 12.1 above builds a bridge between the duality theorems for combs (Theorem 12.1.1) and stars (Theorem 12.1.2) in terms of normal trees (from the first chapter of this part), which we recall here.

Theorem 12.1.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) *G contains a comb attached to U ;*
- (ii) *there is a rayless normal tree $T \subseteq G$ that contains U .*

Moreover, the normal tree T in (ii) can be chosen so that it contains U cofinally.

Theorem 12.1.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) *G contains a star attached to U ;*
- (ii) *there is a locally finite normal tree $T \subseteq G$ that contains U and all whose rays are undominated in G .*

Moreover, the normal tree T in (ii) can be chosen such that it contains U cofinally and every component of $G - T$ has finite neighbourhood.

Our duality theorem for dominated combs in terms of normal trees implies the corresponding duality theorems for combs and stars above. This becomes apparent by a close look at Figure 12.1.1. The three columns of the diagram summarise the three duality theorems. Arrows depict implications between the statements; the dashed arrows indicate that further assumptions are needed to obtain their implications. On the left hand side, the extra assumption is that there is no comb attached to U ; on the right hand side, the extra assumption is that there is no star attached to U .

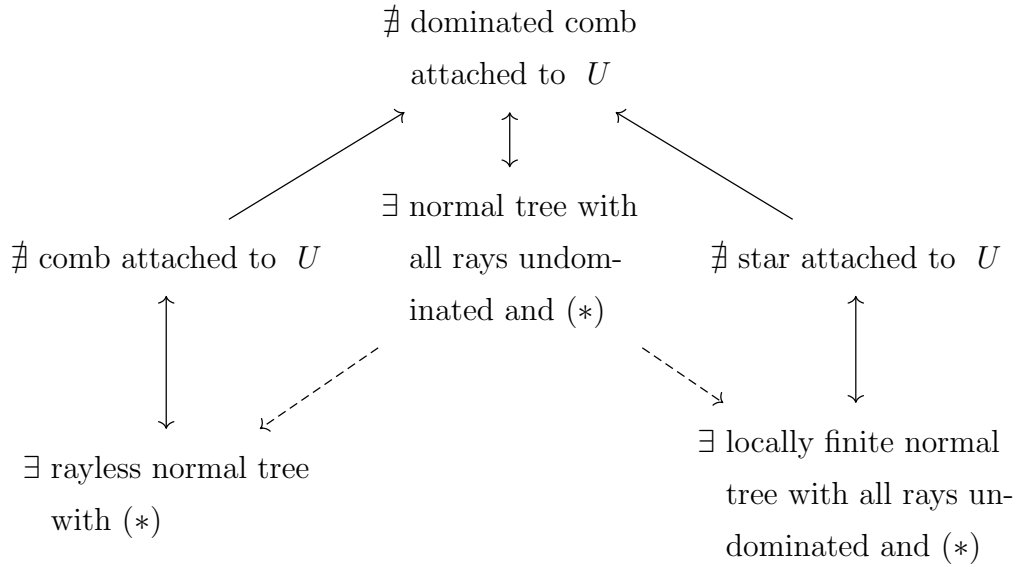


Figure 12.1.1.: The relations between the duality theorems for combs, stars and dominated combs in terms of normal trees.

Condition $(*)$ says that the normal tree contains U cofinally and every component of the graph minus the normal tree has finite neighbourhood.

As a consequence of the two dashed arrows, we obtain the implications $\neg(i) \rightarrow (ii)$ of Theorem 12.1.1 and of Theorem 12.1.2 from the corresponding implication of Theorem 12.1. Indeed, if G does not contain a comb attached to U , then in particular it does not contain a dominated comb attached to U . Hence Theorem 12.1 yields a normal tree, which additionally must be rayless. Similarly, if G does not contain a star attached to U , then in particular it does not contain a dominated comb attached to U . Hence Theorem 12.1 yields a normal tree, which additionally must be locally finite and satisfy that all its rays are undominated.

Since (i) and (ii) of Theorem 12.1.1 and of Theorem 12.1.2 exclude each other almost immediately we have, so far, derived these two duality theorems for combs and stars from our duality theorem for dominated combs—except for the ‘moreover’ part of Theorem 12.1.2.

We proved Theorem 12.1.2 without its ‘moreover’ part in the first chapter of this part. There, instead of proving the ‘moreover’ part as well, we announced that we would prove it in this chapter. And here we prove it, by deriving it from the identical ‘moreover’ part of Theorem 12.1:

Proof of Theorem 12.1.2. Employ Theorem 12.1 as above. □

Another consequence of Theorem 12.1 is a fact whose previous proof, [24, Lemma 2.3], relied on the theorem of Halin [40] which states that every connected graph without a subdivided K^{\aleph_0} has a normal spanning tree:

Corollary 12.1.3. *If G is a connected graph none of whose ends is dominated, then G is normally spanned.* □

For the proof of Theorem 12.1, we shall need the following four lemmas and a result by Jung (cf. [43, Satz 6] or Theorem 11.2.5). The first lemma is from the first chapter of this part and we remark that the original statement also takes critical vertex sets in the closure of T or W into account.

Lemma 12.1.4 (see Chapter 11). *Let G be any graph. If $T \subseteq G$ is a rooted tree that contains a vertex set W cofinally, then $\partial_\Omega T = \partial_\Omega W$.*

Recall that for a graph G and a normal tree $T \subseteq G$ the *generalised up-closure* $\llbracket x \rrbracket$ of a vertex $x \in T$ is the union of $[x]$ with the vertex set of $\bigcup \mathcal{C}(x)$, where the set $\mathcal{C}(x)$ consists of those components of $G - T$ whose neighbourhoods meet $[x]$.

Lemma 11.1.10. *Let G be any graph and $T \subseteq G$ any normal tree.*

- (i) *Any two vertices $x, y \in T$ are separated in G by the vertex set $[x] \cap [y]$.*
- (ii) *Let $W \subseteq V(T)$ be down-closed. Then the components of $G - W$ come in two types: the components that avoid T ; and the components that meet T , which are spanned by the sets $\llbracket x \rrbracket$ with x minimal in $T - W$.*

Lemma 11.1.11. *If G is any graph and $T \subseteq G$ is any normal tree, then every end of G in the closure of T contains exactly one normal ray of T . Moreover, sending these ends to the normal rays they contain defines a bijection between $\partial_\Omega T$ and the normal rays of T .*

Lemma 12.1.5. *Let G be a connected graph, let D_0, D_1, \dots be the distance classes of G with respect to an arbitrary vertex of G , and let $n \geq 1$. Then for every infinite $U \subseteq D_n$ the induced subgraph $G[D_0 \cup \dots \cup D_n]$ contains a star attached to U .*

Proof. By induction on n . For $n = 1$ there is a star in $G[D_0 \cup D_1]$ with centre in D_0 and attachment set U . Now suppose that $n > 1$, and let any infinite $U \subseteq D_n$ be given. For every $u \in U$ pick an edge e_u at u incident with some vertex w_u in D_{n-1} , and let $W \subseteq D_{n-1}$ consist of the vertices w_u . If some vertex $w \in W$ is incident with infinitely many edges of the form e_u , we have the desired star. Otherwise every vertex $w \in W$ is incident with only finitely many such edges. In that case, we find an infinite subset $W' \subseteq W$ together with a matching of W' and an infinite subset of U formed by edges e_u . Then we employ the induction hypothesis to W' to yield a star S in $G[D_0 \cup \dots \cup D_{n-1}]$ attached to W' , and we extend S to the desired star by adding edges of the matching. \square

Theorem 11.2.5 (Jung). *Let G be any graph. A vertex set $W \subseteq V(G)$ is normally spanned if and only if it is a countable union of dispersed sets. In particular, G is normally spanned if and only if $V(G)$ is a countable union of dispersed sets.*

Now we are ready to prove our first duality theorem for dominated combs:

Proof of Theorem 12.1. First, we show that at most one of (i) and (ii) holds. Assume for a contradiction that both hold, let R be the spine of a dominated comb attached to U and let T be a normal tree as in (ii). Then the end of R lies in the closure of $U \subseteq T$, so by Lemma 11.1.11 the normal tree T contains a normal ray from that end. But then the vertices dominating R in G also dominate that normal ray, a contradiction.

It remains to show that at least one of (i) and (ii) holds; we show $\neg(\text{i}) \rightarrow (\text{ii})$. For this, pick an arbitrary vertex v_0 of G and write D_n for the n th distance class of G with respect to v_0 . If for some distance class D_n there was a comb in G attached to $D_n \cap U$, then that comb would be dominated by Lemma 12.1.5 contrary to our

assumptions. Therefore, all the sets $D_n \cap U$ with $n \in \mathbb{N}$ are dispersed. Now, Jung's Theorem 11.2.5 yields a normal tree $T' \subseteq G$ that contains U , and by replacing T' with the down-closure of U we may assume that T' even contains U cofinally. The normal rays of T' cannot be dominated in G because a normal ray of T' that is dominated in G would give rise to a dominated comb attached to U .

For the 'moreover' part it remains to find a normal tree $T \subseteq G$ just like T' , but such that additionally every component of $G - T$ has finite neighbourhood. Our proof proceeds in three steps, as follows.

It will turn out that if a component C of $G - T'$ has infinite neighbourhood, then there are rays in C whose ends in G lie in the closure of U . In step one we define a superset $\hat{U} \supseteq U$ that extends $V(T')$ by carefully chosen vertex sets of such rays, and we verify $\partial_\Omega \hat{U} = \partial_\Omega U$. The choice of \hat{U} allows us in step two to apply Theorem 12.1 (without the 'moreover' part) to \hat{U} , yielding a normal tree $T'' \subseteq G$ (which contains $V(T')$ but in general does not extend T') for which we then verify that every component of $G - T''$ has finite neighbourhood. As T'' contains \hat{U} cofinally, it also contains U , but it need not do so cofinally. Hence in step three we fix this by taking T to be the down-closure of U in T'' , and we verify that T is as desired.

As our first step, we prepare the construction of T'' . Write $\mathcal{D}_{T'}$ for the collection of the components of $G - T'$ that have infinite neighbourhood. For each component $C \in \mathcal{D}_{T'}$ the down-closure $[N(C)]$ is a normal ray in T' which we denote by R_C .

Using Zorn's lemma we choose, for every component $C \in \mathcal{D}_{T'}$, an inclusionwise maximal collection \mathcal{R}_C of pairwise disjoint rays in the end of R_C in G such that all these rays are contained in C . We write U_C for the vertex set of $\bigcup \mathcal{R}_C$ and put

$$\hat{U} := V(T') \cup \bigcup \{ U_C \mid C \in \mathcal{D}_{T'} \}$$

while noting $U \subseteq V(T') \subseteq \hat{U}$.

We claim that $\partial_\Omega \hat{U} = \partial_\Omega U$ holds. The backward inclusion is immediate from the inclusion $\hat{U} \supseteq U$. For the forward inclusion, consider any end ω of G that is not contained in the closure of U ; we show $\omega \notin \partial_\Omega \hat{U}$. As T' contains U cofinally, it follows from Lemma 12.1.4 that the end ω does not lie in the closure of T' either. Let $X \subseteq V(G)$ be a finite set of vertices witnessing that ω does not lie in the closure of T' . The plan is to slightly expand X so that it witnesses that ω does not lie in the closure of \hat{U} as well. The component $C(X, \omega)$ avoids T' , and

in particular $C(X, \omega)$ avoids U . But $C(X, \omega)$ may meet some U_C with $C \in \mathcal{D}_{T'}$. However, the rays in the union of all sets \mathcal{R}_C over $C \in \mathcal{D}_{T'}$ are pairwise disjoint by the choice of the sets \mathcal{R}_C , and none of these rays' ends lives in $C(X, \omega) \subseteq G - T'$. So as X is finite this means that at most finitely many vertices of $C(X, \omega)$ belong to rays from the sets \mathcal{R}_C , and therefore adding these vertices to X results in the finite X separating ω from \hat{U} as well.

Now that we have $\partial_\Omega \hat{U} = \partial_\Omega U$ we apply Theorem 12.1 (without its 'moreover' part which we are currently proving) to \hat{U} in G and obtain a normal tree $T'' \subseteq G$ that contains \hat{U} cofinally and all whose rays are undominated in G . We claim that every component C of $G - T''$ has finite neighbourhood. For this, assume for a contradiction that some component C of $G - T''$ has infinite neighbourhood. Let R be the normal ray in T'' given by the down-closure of that neighbourhood in T'' , and write Z for the set of those vertices in C that send edges to T'' . Since T'' contains \hat{U} cofinally it follows from Lemma 12.1.4 that $\partial_\Omega T'' = \partial_\Omega \hat{U}$ and thus also $\partial_\Omega T'' = \partial_\Omega U$. As a consequence we know that the end ω of R in G lies in the closure of U .

If some $z \in Z$ would send infinitely many edges to T'' , then z would dominate R , contradicting the choice of T'' . Thus every vertex in Z may send only finitely many edges to R , and in particular Z must be infinite. Therefore, we find an infinite subset $Z' \subseteq Z$ for which G contains a matching of Z' and an infinite subset of $V(R)$. Applying the star-comb lemma in C to Z' then, as R was just noted to be undominated, must yield a comb in C attached to Z' . That comb's spine R' is equivalent in G to R . Now consider the component D of $G - T'$ that contains C . Having in mind that ω lies in the closure of U , we find that the normal tree T' that contains U cofinally does contain a normal ray equivalent to R , cf. Lemma 11.1.11. This normal ray in T' must be R_D , so in particular we have $D \in \mathcal{D}_{T'}$. But then the spine $R' \subseteq C$ is disjoint from all the rays in \mathcal{R}_D since C avoids $U_D \subseteq T''$, contradicting the maximality of \mathcal{R}_D . Thus, every component C of $G - T''$ must have finite neighbourhood.

Finally, let $T \subseteq G$ be the normal tree given by the down-closure of U in T'' . Then T contains U cofinally. We claim that every component of $G - T$ has a finite neighbourhood. Indeed, consider any component C of $G - T$. If C is also a component of $G - T''$, then—as we have already seen—it has finite neighbourhood. Otherwise, by Lemma 11.1.10, the component C is spanned by $\llbracket x \rrbracket$ with respect

to T'' for the minimal node x in $C \cap T''$. Now, as T'' is normal, C can only send edges to the finite set $[x] \setminus \{x\}$. Hence the component C has finite neighbourhood as claimed. \square

Let us discuss an application of our duality theorem for dominated combs in terms of normal trees. In [24], Diestel proves the following theorem that relates the metrisability of $|G|$ to the existence of normal spanning trees (we refer to [24, Section 2] for definitions concerning $|G|$, MTOP , VTOP and TOP):

Theorem 12.1.6 ([24, Theorem 3.1]). *Let G be any connected graph.*

- (i) *In MTOP , $|G|$ is metrisable if and only if G has a normal spanning tree.*
- (ii) *In VTOP , $|G|$ is metrisable if and only if no end of G is dominated.*
- (iii) *In TOP , $|G|$ is metrisable if and only if G is locally finite.*

Assertions (ii) and (iii) of this theorem can be reformulated so as to speak about normal spanning trees: By Theorem 12.1 with $U = V(G)$, the graph G having no dominated end is equivalent to G having a normal spanning tree all of whose normal rays are undominated. And by Theorem 12.1.2 with $U = V(G)$, the graph G being locally finite is equivalent to G having a locally finite normal spanning tree all of whose normal rays are undominated. That is why we may hope that these theorems allow us to localise Theorem 12.1.6 above to arbitrary vertex sets $U \subseteq V(G)$. We will show that this is largely possible.

Recall that a *standard subspace* of $|G|$ (with regard to MTOP , VTOP or TOP) is a subspace Y of $|G|$ that is the closure \overline{H} of a subgraph H of G (see Diestel's textbook [20, p. 246]).

Lemma 12.1.7. *Let G be any graph, let $T \subseteq G$ be any normal tree and consider the spaces $|T|$ and $|G|$, both in the same choice of one of the three topologies MTOP , VTOP or TOP . Then $|T|$ is homeomorphic to the standard subspace \overline{T} of $|G|$.*

Proof. By Lemma 11.1.11, the identity on T extends to a bijection $|T| \rightarrow \overline{T} \subseteq |G|$ that sends every end of T to the unique end of G including it. Using Lemma 11.1.10 it is straightforward to verify that the bijection is a homeomorphism, no matter which of the three topologies we chose. \square

Theorem 12.1.8. *Let G be any connected graph and $U \subseteq V(G)$ any vertex set.*

- (i) *In MTOP , $|G|$ has a metrisable standard subspace containing U if and only if there is a normal tree $T \subseteq G$ that contains U .*
- (ii) *In VTOP , $|G|$ has a metrisable standard subspace containing U whenever there is no dominated comb in G attached to U .*
- (iii) *In TOP , $|G|$ has a metrisable standard subspace containing U whenever there is no star in G attached to U .*

Proof. (i) First, suppose that there is a metrisable standard subspace containing U . We imitate Diestel's proof of the corresponding implication of Theorem 12.1.6 (i). Recall from [24] that a set of vertices of G is dispersed in G if and only if it is closed in $|G|$. So by Jung's Theorem 11.2.5, it suffices to show that U can be written as a countable union of closed vertex sets. For this, the sets U_n consisting of the vertices in U that have distance $\geq 1/n$ from every end can be taken: On the one hand, every U_n is the intersection of complements of open balls of radius $1/n$, and hence closed. On the other hand, every vertex $u \in U$ is contained in U_n for some $n \in \mathbb{N}$ because G is open in $|G|$.

Now, suppose that there is a normal tree $T \subseteq G$ containing U and consider the standard subspace \overline{T} . By Lemma 12.1.7 the spaces \overline{T} and $|T|$ are homeomorphic. Since T normally spans itself, $|T|$ is metrisable by Theorem 12.1.6 (i).

(ii) Suppose that G contains no dominated comb attached to U . By Theorem 12.1, there is a normal tree $T \subseteq G$ that contains U cofinally. Then $\overline{T} \cong |T|$ by Lemma 12.1.7, and $|T|$ is metrisable by Theorem 12.1.6 (ii).

(iii) If G contains no star attached to U , then by Theorem 12.1.2 there is a locally finite normal tree $T \subseteq G$ that contains U cofinally. By Lemma 12.1.7 we have that the standard subspace that arises from T is homeomorphic to $|T|$ with TOP . Since T is locally finite, TOP coincides with MTOP on $|T|$ which is metrisable by Theorem 12.1.6 (i). □

The statements (ii) and (iii) of Theorem 12.1.8 cannot be extended so as to give equivalent statements: Let R be a ray, $U = V(R)$ and consider the graph $G := R * v$ where $v \notin R$ is any vertex (that is, G is obtained from $R + v$ by adding all possible v - R edges). By Lemma 12.1.7 the standard subspace that arises from R is homeomorphic to $|R|$, which in turn is metrisable by Theorem 12.1.6. But $R \subseteq G$ is a dominated comb attached to U .

12.2. Tree-decompositions

In the previous section, we have presented a duality theorem for dominated combs in terms of normal trees. And we have deduced from this theorem the hard implications $\neg(\text{i}) \rightarrow (\text{ii})$ of Theorem 12.1.1 and of Theorem 12.1.2 (the duality theorems for combs and stars in terms of normal trees).

Therefore we may expect from a duality theorem for dominated combs in terms of tree-decompositions to reestablish the hard implications $\neg(\text{i}) \rightarrow (\text{ii})$ of the duality theorems for combs and stars in terms of tree-decompositions (Theorem 12.2.1 and Theorem 12.2.2 below)—by following arrows in Figure 12.2.1 like we did in Figure 12.1.1.

Theorem 12.2.1 (see Chapter 11). *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) *G contains a comb attached to U ;*
- (ii) *G has a rayless tree-decomposition into parts each containing at most finitely many vertices from U and whose parts at non-leaves of the decomposition tree are all finite.*

Moreover, the tree-decomposition in (ii) can be chosen with connected separators.

Recall from Chapter 11 that a tree-decomposition (T, \mathcal{V}) of a given graph G with finite separators *displays* a set Ψ of ends of G if τ restricts to a bijection $\tau \upharpoonright \Psi: \Psi \rightarrow \Omega(T)$ between Ψ and the end space of T and maps every end that is not contained in Ψ to some node of T , where $\tau: \Omega(G) \rightarrow \Omega(T) \sqcup V(T)$ maps every end of G to the end or node of T which it corresponds to or lives at, respectively.

Theorem 12.2.2 (see Chapter 11). *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) *G contains a star attached to U ;*
- (ii) *G has a locally finite tree-decomposition with finite and pairwise disjoint separators such that each part contains at most finitely many vertices of U .*

Moreover, the tree-decomposition in (ii) can be chosen with connected separators and so that it displays $\partial_\Omega U$.

In Section 12.2.1, we will prove a duality theorem for dominated combs in terms of tree-decompositions, making the left but not the right dashed arrow in Figure 12.2.1 true. In Section 12.2.2, the situation is reversed: we will prove a duality theorem making the right but not the left dashed arrow in Figure 12.2.1 true. Here we also provide a short proof of Carmesin’s result [15], which states that every graph has a tree-decomposition displaying all its undominated ends, for normally spanned graphs. Finally, in Section 12.2.3, we will prove a duality theorem that makes both the left and the right dashed arrow in Figure 12.2.1 true. This will be achieved by combining our proof technique from Section 12.2.1 and our duality theorem from Section 12.2.2.

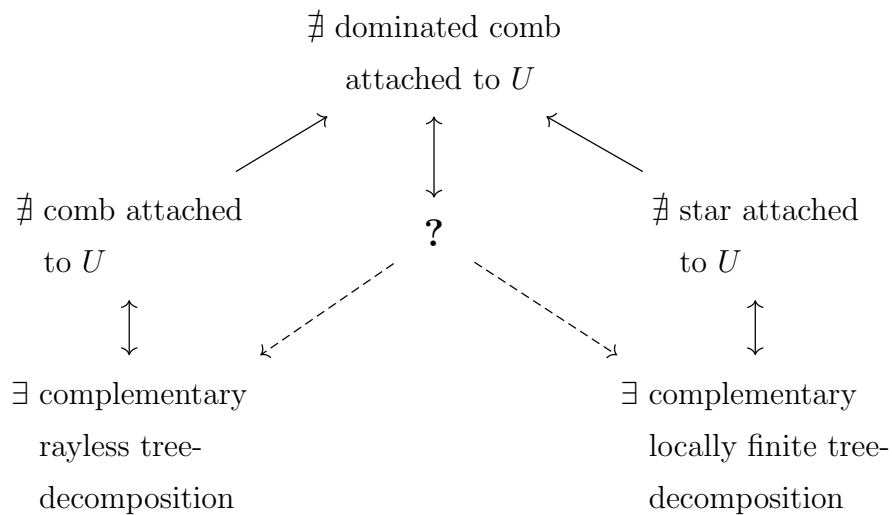


Figure 12.2.1.: The desired relation between stars, combs, dominated combs and complementary tree-decompositions.

The left and right dashed arrow describe an implication whenever there is no comb and no star attached to U , respectively.

12.2.1. A duality theorem related to combs

Here we present a duality theorem for dominated combs in terms of tree-decompositions making the left but not the right dashed arrow of Figure 12.2.1 true:

Theorem 12.2.3. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a dominated comb attached to U ;
- (ii) G has a tree-decomposition (T, \mathcal{V}) that satisfies:
 - (a) each part contains at most finitely many vertices from U ;
 - (b) all parts at non-leaves of T are finite;
 - (c) every dominated end of G lives in a part at a leaf of T .

Moreover, the tree-decomposition in (ii) can be chosen with connected separators and so that it displays $\partial_\Omega U$.

Before we provide a proof of this theorem, let us deduce the left dashed arrow of Figure 12.2.1 from it (also see Figure 12.2.2 which shows the first two columns of Figure 12.2.1 in greater detail and with Theorem 12.2.3 (ii) including the theorem's 'moreover' part inserted for '?'): If G does not contain a comb attached to U , then in particular it does not contain a dominated comb attached to U . Hence Theorem 12.2.3 returns a tree-decomposition (T, \mathcal{V}) of G which we may choose so that it satisfies the theorem's 'moreover' part; in particular (T, \mathcal{V}) displays $\partial_\Omega U$. Our assumption that there is no comb attached to U implies that $\partial_\Omega U$ is empty and hence T is rayless. Using the corresponding conditions from Theorem 12.2.3 (ii) including the theorem's 'moreover' part, we conclude that (T, \mathcal{V}) is as in Theorem 12.2.1 (ii) including the theorem's 'moreover' part.

Finally, we prove Theorem 12.2.3:

Proof of Theorem 12.2.3. First, we show that at most one of (i) and (ii) holds. Assume for a contradiction that G contains a dominated comb attached to U and has, at the same time, a tree-decomposition (T, \mathcal{V}) as in (ii). Let R be the comb's spine. Since every dominated end of G lives in a part at a leaf of T , and since all parts at non-leaves are finite, we find without loss of generality a leaf ℓ of T with $R \subseteq G[V_\ell]$. But each part contains at most finitely many vertices from U . In particular, V_ℓ contains at most finitely many vertices from U . Therefore, the comb must send some infinitely many pairwise disjoint paths to vertices in $U \setminus V_\ell$. But the separator of G that is associated with the edge $\ell t \in T$ at ℓ is contained in the intersection $V_\ell \cap V_t \subseteq V_t$ which is finite since V_t is, a contradiction.

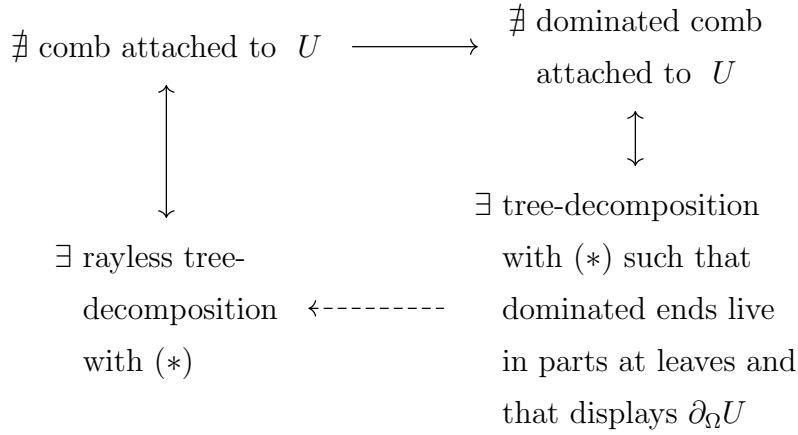


Figure 12.2.2.: The first two columns of Figure 12.2.1 with Theorem 12.2.3 (ii) including the theorem’s ‘moreover’ part inserted for ‘?’.

Condition $(*)$ says that parts contain at most finitely many vertices from U , that parts at non-leaves are finite and that the separators are connected.

Now, to show that at least one of (i) and (ii) holds, we show $\neg(\text{i}) \rightarrow (\text{ii})$. By Theorem 12.1 we find a normal tree $T_{\text{NT}} \subseteq G$ containing U cofinally all whose rays are undominated in G and such that every component of $G - T_{\text{NT}}$ has finite neighbourhood. We construct the desired tree-decomposition from T_{NT} .

Given a component C of $G - T_{\text{NT}}$ the neighbourhood of C is a finite chain in the tree-order of T_{NT} , and hence has a maximal element $t_C \in T_{\text{NT}}$. We obtain the tree T from T_{NT} by adding each component C of $G - T_{\text{NT}}$ as a new vertex and joining it precisely to t_C .

Having defined the decomposition tree T it remains to define the parts of the desired tree-decomposition. For nodes $t \in T_{\text{NT}} \subseteq T$ we let V_t consist of the down-closure $[t]_{T_{\text{NT}}}$ of t in the normal tree T_{NT} . And for newly added nodes C we let V_C be the union of V_{t_C} and the vertex set of the component C , i.e., we put $V_C := [t]_{T_{\text{NT}}} \cup V(C)$.

Since T_{NT} is normal and contains U cofinally, it follows by standard arguments employing Lemma 12.1.4 and Lemma 11.1.11 that (T, \mathcal{V}) displays $\partial_\Omega U$. Conditions (a) and (b) hold by construction. Combining (b) with (T, \mathcal{V}) displaying $\partial_\Omega U$ gives (c), which in turn is—as the rest of the ‘moreover’ part—a direct consequence of how the parts are defined. \square

Example 12.2.4. The tree-decomposition in Theorem 12.2.3 (ii) cannot be chosen to additionally have pairwise disjoint separators, which shows that the theorem does not make the right dashed arrow in Figure 12.2.1 true. To see this suppose that G consists of the first three levels of T_{\aleph_0} , the tree all whose vertices have countably infinite degree, and let $U = V(G)$. Then G contains no comb attached to U . Suppose for a contradiction that G has a tree-decomposition (T, \mathcal{V}) as in Theorem 12.2.3 (ii) which additionally has pairwise disjoint separators. The graph G being rayless and U being the whole vertex set of G together with our assumption that (T, \mathcal{V}) has pairwise disjoint separators makes sure that (T, \mathcal{V}) also displays $\partial_\Omega U$. In particular, by our argumentation in the text below Theorem 12.2.3, (T, \mathcal{V}) is also a tree-decomposition of G complementary to combs as in Theorem 12.2.1. But then (T, \mathcal{V}) cannot have pairwise disjoint separators, as pointed out in Example 11.2.7.

12.2.2. A duality theorem related to stars

Here we present a duality theorem for dominated combs in terms of tree-decompositions making the right but not the left dashed arrow in Figure 12.2.1 true.

Theorem 12.2.5. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains a dominated comb attached to U ;
- (ii) G has a tree-decomposition with pairwise disjoint finite separators that displays $\partial_\Omega U$.

Moreover, the tree-decomposition in (ii) can be chosen with connected separators and rooted so that it covers U cofinally.

Before we prepare the proof of our theorem, let us deduce the right dashed arrow of Figure 12.2.1 from it (also see Figure 12.2.3 which shows the last two columns of Figure 12.2.1 in greater detail and where Theorem 12.2.5 (ii) including the theorem's 'moreover' part is inserted for '?'): If G does not contain a star attached to U , then in particular it does not contain a dominated comb attached to U . Hence Theorem 12.2.5 yields a tree-decomposition (T, \mathcal{V}) of G which we choose so that it also satisfies the theorem's 'moreover' part; in particular (T, \mathcal{V})

is rooted so that it covers U cofinally. By assumption, the star-comb lemma yields a comb in G attached to U' for every infinite subset U' of U . Since (T, \mathcal{V}) displays $\partial_\Omega U$ this means that no part can meet U infinitely. And additionally employing the pairwise disjoint finite separators plus U being cofinally covered by the tree-decomposition, we deduce that no node of T can have infinite degree: Suppose for a contradiction that $t \in T$ is a vertex of infinite degree. For every up-neighbour t' of t we choose a vertex from U that is contained in a part $V_{t''}$ with $t'' \geq t'$ in T . Then applying the star-comb lemma in G to the infinitely many chosen vertices from U yields a comb. The end of the comb's spine must then live at t because the separators of (T, \mathcal{V}) are all finite and pairwise disjoint. But this contradicts the fact that (T, \mathcal{V}) displays $\partial_\Omega U$ which contains the end of the comb's spine. Finally, (T, \mathcal{V}) inherits the properties of the 'moreover' part of Theorem 12.2.2 from the identical properties of Theorem 12.2.5 (ii) including that theorem's 'moreover' part.

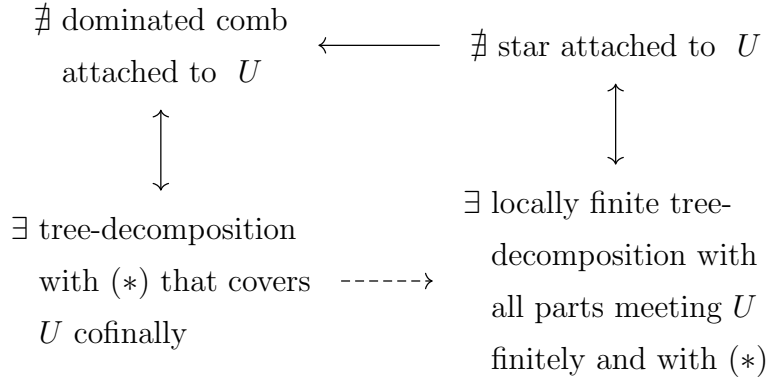


Figure 12.2.3.: The last two columns of Figure 12.2.1 with Theorem 12.2.5 (ii) including the theorem's 'moreover' part inserted for '?'.
Condition $(*)$ says that the tree-decomposition displays $\partial_\Omega U$ and has pairwise disjoint finite connected separators.

In order to prove Theorem 12.2.5, we will employ the following result by Carmesin. Recall that a rooted S_{\aleph_0} -tree (T, α) has *upwards disjoint* separators if for every two edges $\vec{e} < \vec{f}$ pointing away from the root r of T the separators of $\alpha(\vec{e})$ and $\alpha(\vec{f})$ are disjoint. And (T, α) is *upwards connected* if for every edge \vec{e} pointing away from the root r the induced subgraph $G[B]$ stemming from $(A, B) = \alpha(\vec{e})$ is connected. A rooted tree-decomposition has *upwards disjoint* separators or is

upwards connected if its corresponding S_{\aleph_0} -tree is.

Theorem 12.2.6 (Theorem 11.1.17). *Every connected graph G has an upwards connected rooted tree-decomposition with upwards disjoint finite separators that displays the undominated ends of G .*

Carmesin's proof of this theorem in [15] is long and complex. However, in this paper we need his theorem only for normally spanned graphs. This is why we will provide a substantially shorter proof for this class of graphs (cf. Theorem 12.2.9). Furthermore, we prove that the separators of the tree-decomposition in Theorem 12.2.6 can be chosen pairwise disjoint and connect, which makes it easier for us to apply the theorem. The latter is essentially accomplished by the following lemma:

Lemma 12.2.7. *Let G be any connected graph and let Ψ be any set of ends of G . Then the following assertions are equivalent:*

- (i) *G has an upwards connected rooted tree-decomposition with upwards disjoint finite separators that displays Ψ ;*
- (ii) *G has a tree-decomposition with pairwise disjoint finite connected separators that displays Ψ .*

Indeed, this lemma together with Theorem 12.2.6 yields the following theorem:

Theorem 12.2.8. *Every connected graph G has a tree-decomposition with pairwise disjoint finite connected separators that displays the undominated ends of G .*

□

For the proof of Lemma 12.2.7 we need the following lemma from the first chapter of this part:

Lemma 11.1.16. *Let G be any graph. Every upwards connected rooted S_{\aleph_0} -tree (T, α) with upwards disjoint separators displays the ends of G that correspond to the ends of T .*

Proof of Lemma 12.2.7. The implication (ii)→(i) is immediate, we prove (i)→(ii).

Let (T, \mathcal{V}) be an upwards connected rooted tree-decomposition of G with upwards disjoint finite separators that displays Ψ . We consider the S_{\aleph_0} -tree (T, α)

corresponding to (T, \mathcal{V}) . For every edge $e = t_1 t_2$ of T with $t_1 \leq t_2$ and $\alpha(t_1, t_2) = (A, B)$ we use that (T, α) is upwards connected to find a finite connected subgraph H_e of $G[B]$ that contains $A \cap B$. We define $A' := A \cup V(H_e)$ and $B' := B$ so that the separator $A' \cap B' = V(H_e)$ is connected. Then we define $\alpha'(t_1, t_2) := (A', B')$ and $\alpha'(t_2, t_1) := (B', A')$ to obtain another map $\alpha': \vec{E}(T) \rightarrow \vec{S}_{\aleph_0}$. The pair (T, α') does not need to be an S_{\aleph_0} -tree, for some of its separations might cross. To fix this, we will carefully ‘thin out’ the tree and, consequently, the set of separations associated with it via α' . This will result in a contraction minor \tilde{T} of T such that $(\tilde{T}, \tilde{\alpha}')$ with $\tilde{\alpha}' := \alpha' \upharpoonright E(\tilde{T})$ is an S_{\aleph_0} -tree with upwards disjoint finite connected separators that still displays Ψ . Then, in order to obtain the desired tree-decomposition, we just have to contract all the edges of \tilde{T} that are at an even distance from the root, and restrict $\tilde{\alpha}'$ to the smaller edge set of the resulting contraction minor of \tilde{T} .

To begin the construction of \tilde{T} , we partially order $E(T)$ by letting $e \leq f$ whenever e precedes f on a path in T starting at the root. For every edge e of T we do the following. We write T_e for the component of $T - e$ that does not contain the root. Then, we let $F_e \subseteq E(T_e)$ consist of the down-closure in $E(T_e)$ of those edges whose α' -separator (the separator of the separation that α' associates with the edge) meets the α' -separator of e . A distance argument employing the original upwards disjoint α -separators ensures that F_e induces a rayless down-closed subtree of T_e .

In order to reasonably name edges of T whose contraction leads to \tilde{T} , we recursively construct a sequence E_0, E_1, \dots of pairwise disjoint subsets of $E(T)$ such that their overall union $E' := \bigsqcup_{n \in \mathbb{N}} E_n$ induces a partition $\{ \{e\}, F_e \mid e \in E' \}$ of $E(T)$. The construction goes as follows. Take E_0 to be the set of minimal edges of $E(T)$, i.e., take E_0 to be the set of edges of T at the root. Then at step $n > 0$ consider the edges of $E(T)$ that are not contained in the down-closed edge set $\bigcup \{ \{e\}, F_e \mid e \in E_0 \cup \dots \cup E_{n-1} \}$, and take the minimal ones to form E_n .

Once we have constructed E' , we take \tilde{T} to be the contraction minor of T that is obtained by contracting all the edges occurring in some F_e with $e \in E'$. Then $(\tilde{T}, \tilde{\alpha}')$ has upwards disjoint finite connected separators and displays Ψ , as we verify now. Consider any distinct two edges e and f of \tilde{T} , that is, edges $e, f \in E'$. If the two edges are comparable with $e < f$, say, then their α' -separators are disjoint as f is not in F_e , and so in particular their α' -separations are nested.

Otherwise e and f are incomparable, and then their α' -separations are nested by the construction of α' from α . Therefore, the separators of $(\tilde{T}, \tilde{\alpha}')$ are finite, connected and pairwise disjoint. It remains to show that $(\tilde{T}, \tilde{\alpha}')$ displays Ψ .

Since all F_e are rayless, we deduce that every ray of T meets E' infinitely. Consequently, the rooted rays of T correspond bijectively to the rooted rays of \tilde{T} via the map $R \mapsto \tilde{R}$ satisfying $E(R) \supseteq E(\tilde{R})$. Now to see that $(\tilde{T}, \tilde{\alpha}')$ displays Ψ , consider any end ω of G . If ω is not contained in Ψ , then ω lives at a node $t \in T$ (with regard to (T, α)), and hence ω lives at the node $\tilde{t} \in \tilde{T}$ (with regard to $(\tilde{T}, \tilde{\alpha}')$) that contains t . Otherwise ω lies in Ψ . Then ω corresponds to an end of T . This end is uniquely represented by a rooted ray R of T . And then from $E(\tilde{R}) \subseteq E(R)$ it follows that ω corresponds to the end of \tilde{R} in \tilde{T} . So the ends in Ψ correspond to ends of \tilde{T} while all ends in $\Omega \setminus \Psi$ live at nodes. Then by Lemma 11.1.16 this correspondence is bijective, and hence $(\tilde{T}, \tilde{\alpha}')$ displays Ψ as desired. \square

Theorem 12.2.9. *Let G be any connected graph. If $T_{\text{NT}} \subseteq G$ is a normal tree such that every component of $G - T_{\text{NT}}$ has finite neighbourhood, then G has a rooted tree-decomposition (T, \mathcal{V}) with the following three properties:*

- *the separators are pairwise disjoint, finite and connected;*
- *(T, \mathcal{V}) displays the undominated ends in the closure of T_{NT} ;*
- *(T, \mathcal{V}) covers $V(T_{\text{NT}})$ cofinally.*

Proof. Given the normal tree T_{NT} , by Lemma 12.2.7 it suffices to find an upwards connected rooted tree-decomposition (T, \mathcal{V}) of G that displays the undominated ends in the closure of T_{NT} and that has upwards disjoint finite separators all of which meet $V(T_{\text{NT}})$.

Let us write r for the root of T_{NT} . Recall that every component of $G - T_{\text{NT}}$ has finite neighbourhood by assumption. Hence every end $\omega \in \Omega \setminus \partial_{\Omega} T_{\text{NT}}$ lives in a unique component of $G - T_{\text{NT}}$; we define the *height* of ω to be the height of the maximal neighbour of this component in T_{NT} .

Starting with $T_0 = r$ and $\alpha_0 = \emptyset$ we recursively construct an ascending¹

¹Here, we mean ascending in both entries with regard to inclusion, i.e., $T_n \subseteq T_{n+1}$ and $\alpha_n \subseteq \alpha_{n+1}$ for all $n \in \mathbb{N}$.

sequence of S_{\aleph_0} -trees (T_n, α_n) all rooted in r and satisfying the following conditions:

- (i) the separators of (T_n, α_n) are upwards disjoint and they are vertex sets of ascending paths in T_{NT} ;
- (ii) T_n arises from T_{n-1} by adding edges to its $(n-1)$ th level;
- (iii) undominated ends in the closure of T_{NT} live at nodes of the n th level of T_n with regard to (T_n, α_n) ;
- (iv) if $\omega \in \Omega \setminus \partial_\Omega T_{\text{NT}}$ has height $< n$, then ω lives at a node of T_n of height $< n$ with regard to (T_n, α_n) .

Before pointing out the details of our construction, let us see how to complete the proof once the (T_n, α_n) are defined. Consider the S_{\aleph_0} -tree (T, α) defined by letting $T := \bigcup_{n \in \mathbb{N}} T_n$ and $\alpha := \bigcup_{n \in \mathbb{N}} \alpha_n$, and let (T, \mathcal{V}) be the corresponding tree-decomposition of G . By (i) we have that (T, \mathcal{V}) is indeed a rooted tree-decomposition with upwards disjoint finite connected separators all of which meet the vertex set of T_{NT} . It remains to prove that (T, \mathcal{V}) displays the undominated ends in the closure of T_{NT} .

By Lemma 11.1.16 it suffices to show that the undominated ends in the closure of T_{NT} are precisely the ends of G that correspond to the ends of T . For the forward inclusion, consider any undominated end ω in the closure of T_{NT} . By (iii), it follows that ω lives at a node t_n of T_n (with regard to (T_n, α_n)) at level n for every $n \in \mathbb{N}$, and these nodes form a ray $R = t_0 t_1 \dots$ of T . Then ω corresponds to the end of T containing R .

For an indirect proof of the backward inclusion, consider any end ω of G that is either dominated or not contained in the closure of T_{NT} . We show that ω does not correspond to any end of T . If ω is dominated, then this follows from the fact that (T, \mathcal{V}) has upwards disjoint finite separators. Otherwise ω is not contained in the closure of T_{NT} . Let $n \in \mathbb{N}$ be strictly larger than the height of ω . By (iv), it follows that ω lives at a node t_ω of T_n of height $< n$ with regard to (T_n, α_n) . And by (ii), the tree T_n consists precisely of the first n levels of T . We conclude that ω lives in the part of (T, \mathcal{V}) corresponding to t_ω .

Now, we turn to the construction of the (T_n, α_n) , also see Figure 12.2.4. At step $n+1$ suppose that (T_n, α_n) has already been defined and recall that the




Figure 12.2.4.: The construction of the (T_n, α_n) in the proof of Theorem 12.2.9. Here the vertex set Z consists of all vertices that are contained in some Z_y with $y \in Y$. The depicted tree is T_{NT} .

separators of (T_n, α_n) are vertex sets of ascending paths in T_{NT} by (i). Let L be the n th level of T_n . To obtain (T_{n+1}, α_{n+1}) from (T_n, α_n) , we will add for each $\ell \in L$ new vertices (possibly none) to T_n that we join exactly to ℓ and define the image of the so emerging edges under α_{n+1} . So fix $\ell \in L$. Let X be the separator of the separation corresponding to the edge between ℓ and its predecessor in T_n (if $n = 0$ put $X = \emptyset$). Recall that X is the vertex set of an ascending path in T_{NT} by (i). In T_{NT} , let Y be the set of up-neighbours of the maximal vertices in X (for $n = 0$ let $Y := \{r\}$). For each $y \in Y$ let Z_y be the set of those $z \in [y]_{T_{NT}}$ that are minimal with the property that G contains no T_{NT} -path starting in $[y]_{T_{NT}}$ and ending in $[z]_{T_{NT}}$. (Note that a normal ray of T_{NT} that contains y meets Z_y if and only if it is not dominated by any of the vertices in $[y]_{T_{NT}}$; this fact together with (i) will guarantee (iii) for $n+1$.) Then the vertex set of $yT_{NT}z$ separates the connected sets $A_{yz} := (V \setminus \llbracket z \rrbracket_{T_{NT}}) \cup V(yT_{NT}z)$ and $B_{yz} := V(yT_{NT}z) \cup \llbracket z \rrbracket_{T_{NT}}$ whenever $y \in Y$ and $z \in Z_y$. Join a node t_{yz} to ℓ for every pair (y, z) with $y \in Y$ and $z \in Z_y$, and put $\alpha_{n+1}(t_{yz}) := (A_{yz}, B_{yz})$. Then the S_{N_0} -tree (T_{n+1}, α_{n+1}) clearly satisfies (i) and (ii). That it satisfies (iii) was already argued in the construction and (iv) follows from (i) and the definition of $\alpha_{n+1}(t_{yz})$. \square

With Theorem 12.2.9 at hand, we are finally able to prove Theorem 12.2.5:

Proof of Theorem 12.2.5. First, we show that (i) and (ii) cannot hold at the same time. For this, assume for a contradiction that G contains a dominated comb attached to U and has a tree-decomposition (T, \mathcal{V}) with pairwise disjoint finite separators that displays $\partial_\Omega U$. We write ω for the end of G containing the comb's spine. Then ω lies in the closure of U , and since (T, \mathcal{V}) displays $\partial_\Omega U$ there is a unique end η of T to which ω corresponds. But as the finite separators of (T, \mathcal{V}) are pairwise disjoint, it follows that ω is undominated in G , contradicting that ω contains the spine of a dominated comb.

Now, to show that at least one of (i) and (ii) holds, we prove $\neg(\text{i}) \rightarrow (\text{ii})$. Using Theorem 12.1 we find a normal tree $T_{\text{NT}} \subseteq G$ that contains U cofinally and all whose rays are undominated in G . Furthermore, by the 'moreover' part of Theorem 12.1 we may assume that every component of $G - T_{\text{NT}}$ has finite neighbourhood, and by Lemma 12.1.4 we have $\partial_\Omega U = \partial_\Omega T_{\text{NT}}$. Then Theorem 12.2.9 yields a rooted tree-decomposition (T', \mathcal{V}') of G as in (ii) that has connected separators and covers $V(T_{\text{NT}})$ cofinally. It remains to show that (T', \mathcal{V}') can be chosen so as to cover U cofinally. For this, consider the nodes of T' whose parts meet U , and let $T \subseteq T'$ be induced by their down-closure in T' . Then let (T', α') be the S_{\aleph_0} -tree of G that corresponds to (T', \mathcal{V}') and consider the rooted tree-decomposition (T, \mathcal{V}) of G that corresponds to $(T, \alpha' \upharpoonright \vec{E}(T))$. Now (T, \mathcal{V}) is as in (ii) and satisfies the theorem's 'moreover' part. \square

12.2.3. A duality theorem related to stars and combs

Finally, we present a duality theorem for dominated combs in terms of tree-decompositions that makes both the left and the right dashed arrow in Figure 12.2.1 true. In order to state the theorem, we need one more definition. A tree-decomposition (T, \mathcal{V}) of a graph G has *essentially disjoint separators* if there is an edge set $F \subseteq E(T)$ meeting every ray of T infinitely often such that the separators of (T, \mathcal{V}) associated with the edges in F are pairwise disjoint.

Theorem 12.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) *G contains a dominated comb attached to U ;*
- (ii) *G has a tree-decomposition (T, \mathcal{V}) such that:*
 - *each part contains at most finitely many vertices from U ;*
 - *all parts at non-leaves of T are finite;*
 - *(T, \mathcal{V}) has essentially disjoint connected separators;*
 - *(T, \mathcal{V}) displays the ends in the closure of U .*

Before we provide a proof of this theorem, let us see that it relates to the duality theorems for stars and combs in terms of tree-decompositions as desired (also see Figure 12.2.5, which shows Figure 12.2.1 in greater detail and where Theorem 12.2 (ii) including the theorem’s ‘moreover’ part is inserted for ‘?’).

On the one hand, if G does not contain a comb attached to U , then in particular it does not contain a dominated comb attached to U . Hence Theorem 12.2 returns a tree-decomposition (T, \mathcal{V}) . By our assumption that there is no comb attached to U , and since (T, \mathcal{V}) displays $\partial_\Omega U$, it follows that the decomposition-tree T is rayless. We conclude that (T, \mathcal{V}) is as in Theorem 12.2.1 (ii) including the theorem’s ‘moreover’ part.

On the other hand, if G does not contain a star attached to U , then in particular it does not contain a dominated comb attached to U . Hence Theorem 12.2 returns a tree-decomposition (T, \mathcal{V}) that, in particular, has essentially disjoint finite connected separators and displays $\partial_\Omega U$. Write (T, α) for the S_{\aleph_0} -tree that corresponds to (T, \mathcal{V}) . Let $F \subseteq E(T)$ witness that (T, \mathcal{V}) has essentially disjoint separators and root T arbitrarily. By possibly thinning out F , we may assume that each edge in F meets a rooted ray of T . Consider the tree \tilde{T} that is obtained from T by contracting all the edges of T that are not in F and let $\tilde{\alpha}$ be the restriction of α to $\vec{F} = \vec{E}(\tilde{T})$. Then $(\tilde{T}, \tilde{\alpha})$ corresponds to a tree-decomposition (\tilde{T}, \mathcal{W}) of G with pairwise disjoint finite connected separators that displays $\partial_\Omega U$. Thus, the tree-decomposition (\tilde{T}, \mathcal{W}) is one of the tree-decompositions of G that are complementary to dominated combs as in Theorem 12.2.5 (ii) including the theorem’s ‘moreover’ part (it covers U cofinally as F meets every rooted ray of

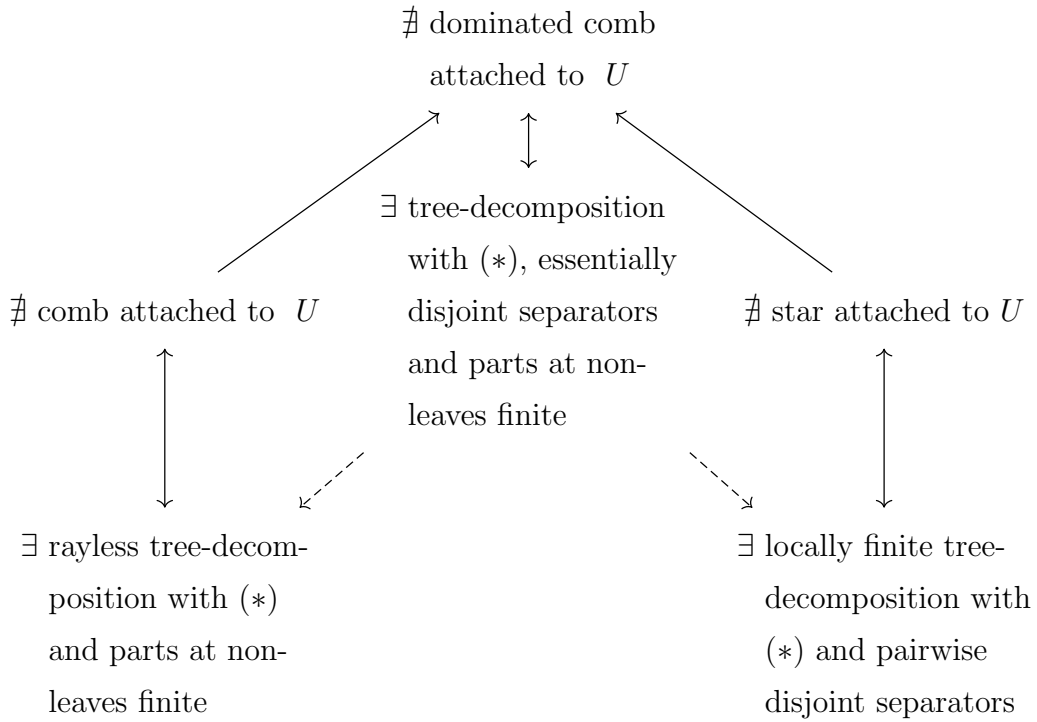


Figure 12.2.5.: The relation between the duality theorems for combs, stars and the final duality theorem for the dominated combs in terms of tree-decompositions.

Condition $(*)$ says that parts contain at most finitely many vertices from U , that the separators are finite and connected, and that the tree-decomposition displays $\partial_\Omega U$.

T while (T, \mathcal{V}) displays $\partial_\Omega U$). Then, as we have already argued below Theorem 12.2.5, the tree-decomposition (\tilde{T}, \mathcal{W}) must be locally finite and each part may contain at most finitely many vertices of U . That is to say that (\tilde{T}, \mathcal{W}) is as in Theorem 12.2.2 (ii) including the theorem's ‘moreover’ part.

As we work with contraction minors in the proof of Theorem 12.2 we need some preparation. Let H and G be any two graphs. We say that H is a contraction minor of G with *fixed branch sets* if an indexed collection of branch sets $\{V_x \mid x \in V(H)\}$ is fixed to witness that G is an IH . In this case, we write $[v] = [v]_H$ for the branch set V_x containing a vertex v of G and also refer to x by $[v]$. Similarly, we write $[U] = [U]_H := \{[u] \mid u \in U\}$ for vertex sets $U \subseteq V(G)$.

The following notation will help us to translate between the endspace of G and that of H . Consider a contraction minor H of a graph G with fixed finite branch

sets. Every direction f of G defines a direction $[f]$ of H by letting $[f](X) := [f(\bigcup X)]$ for every finite vertex set $X \subseteq V(H)$. In fact, it is straightforward to check that every direction of H is defined by a direction of G in this way:

Lemma 12.2.10. *Let H be a contraction minor of a graph G with fixed finite branch sets. Then the map $f \mapsto [f]$ is a bijection between the directions of G and the directions of H . \square*

This one-to-one correspondence then combines with the well-known one-to-one correspondence between the directions and ends of a graph (see Theorem 11.1.7), giving rise to a bijection $\omega \mapsto [\omega]$ between the ends of G and the ends of H . The natural one-to-one correspondence between the two end spaces extends to other aspects of the graphs and their ends:

Lemma 12.2.11. *Let H be a contraction minor of a graph G with fixed finite branch sets, let ω be an end of G and let $U \subseteq V(G)$ be any vertex set. Then ω lies in the closure of U in G if and only if $[\omega]$ lies in the closure of $[U]$ in H ; and ω is dominated in G if and only if $[\omega]$ is dominated in H .*

We remark that this extends [20, Exercise 82 (i)].

Proof. Write f_ω for the direction of G that corresponds to ω . Then the following statements are equivalent:

- (i) ω lies in the closure of U in G ;
- (ii) $f_\omega(X)$ meets U for every finite vertex set $X \subseteq V(G)$;
- (iii) $[f_\omega](X)$ meets $[U]$ for every finite vertex set $X \subseteq V(H)$;
- (iv) $[\omega]$ lies in the closure of $[U]$ in H .

Indeed, one easily verifies (i) \leftrightarrow (ii) \leftrightarrow (iii) \leftrightarrow (iv).

This establishes that the end ω of G lies in the closure of U in G if and only if $[\omega]$ lies in the closure of $[U]$ in H . Similarly, it is straightforward to check that the following statements are equivalent for any vertex v of G (except for (iii) \rightarrow (ii) which we will verify in detail):

- (i) there is a vertex $z \in [v]$ that dominates ω in G ;
- (ii) there is a vertex $z \in [v]$ such that $z \in f_\omega(X)$ for every finite vertex set $X \subseteq V(G) \setminus \{z\}$;
- (iii) $[v] \in [f_\omega](X)$ for every finite vertex set $X \subseteq V(H) \setminus \{[v]\}$;
- (iv) $[v]$ dominates $[\omega]$ in H .

To see (iii) \rightarrow (ii) we show $\neg(\text{ii}) \rightarrow \neg(\text{iii})$. Since (ii) fails, there is for every vertex $z \in [v]$ a finite vertex set $X_z \subseteq V(G) \setminus \{z\}$ such that z is not contained in $f_\omega(X_z)$. Consider the finite vertex set $X := \bigcup_z X_z$. Then no $z \in [v]$ is contained in the component $f_\omega(X)$ or is one of its neighbours, because $f_\omega(X) \subseteq f_\omega(X_z)$ and $z \notin X_z \cup f_\omega(X_z)$. Hence $[v] \notin [f_\omega](X')$ for the neighbourhood X' of $f_\omega(X)$ in G that avoids $[v]$. Therefore the end ω of G is dominated in G if and only if $[\omega]$ is dominated in H . \square

Suppose that (T, \mathcal{V}) is a tree-decomposition of a given graph G and that H is a contraction minor of G with fixed branch sets. The tree-decomposition of H that is obtained by *passing on* (T, \mathcal{V}) to H is the tree-decomposition $(T, ([V_t])_{t \in T})$. Note that this is indeed a tree-decomposition, cf. [20, Lemma 12.3.3].

Lemma 12.2.12. *Let G be any graph, let $U \subseteq V(G)$ be any vertex set, and let (T, \mathcal{V}) be any tree-decomposition of G with finite separators. Furthermore, let H be any contraction minor of G with fixed finite branch sets. Then (T, \mathcal{V}) displays the ends of G in the closure of U if and only if the tree-decomposition of H that is obtained by passing on (T, \mathcal{V}) to H displays the ends of H in the closure of $[U]$.*

Proof. Let (T, α) be the S_{\aleph_0} -tree corresponding to (T, \mathcal{V}) and let (T, α') be the S_{\aleph_0} -tree corresponding to the tree-decomposition of H that is obtained by passing on (T, \mathcal{V}) to H . The ends of G correspond bijectively to the ends of H through the bijection $\Omega(G) \rightarrow \Omega(H)$ that maps ω to $[\omega]$. By Lemma 12.2.11, this bijection restricts to a bijection between the ends of G in the closure of U and the ends of H in the closure of $[U]$. Hence it suffices to show that every end ω of G induces the same orientation on $\vec{E}(T)$ with regard to (T, α) as $[\omega]$ does with regard to (T, α') . For this, let ω be any end of G and write f_ω for the direction of G that corresponds to ω . The following statements are equivalent for every oriented edge $(e, s, t) \in \vec{E}(T)$ and $\alpha(s, t) = (A, B)$:

- (i) (e, s, t) is contained in the orientation of $\vec{E}(T)$ induced by ω ;
- (ii) every ray in ω has a tail in $G[B]$;
- (iii) $f_\omega(A \cap B)$ is included in $G[B]$;
- (iv) $[f_\omega]([A] \cap [B])$ is included in $H[[B]]$;
- (v) every ray in $[\omega]$ has a tail in $H[[B]]$;
- (vi) (e, s, t) is contained in the orientation of $\vec{E}(T)$ induced by $[\omega]$.

Indeed, having in mind that $\alpha'(s, t) = ([A], [B])$ one easily verifies the implications (i) \leftrightarrow (ii) \leftrightarrow (iii) \leftrightarrow (iv) \leftrightarrow (v) \leftrightarrow (vi) in the given order. \square

Lemma 12.2.13. *Let G be any graph, let $U \subseteq V(G)$ be any vertex set and let H be any contraction minor of G with fixed finite branch sets. If assertion (ii) of Theorem 12.2 holds with G and U replaced by H and $[U]$ respectively, then assertion (ii) also holds for G and U .*

Proof. Let (T, \mathcal{W}) be any tree-decomposition of H that witnesses that assertion (ii) holds with G and U replaced by H and $[U]$. Then the tree-decomposition (T, \mathcal{W}) of H gives rise to a tree-decomposition (T, \mathcal{V}) of G by replacing every part with the union of the branch sets that correspond to its vertices. We claim that (T, \mathcal{V}) witnesses that assertion (ii) holds for G and U . For this, we have to show that (T, \mathcal{V}) satisfies four conditions, of which only the fourth condition—that (T, \mathcal{V}) displays the ends of G in the closure of U —is not immediate. This fourth condition, however, is covered by Lemma 12.2.12. \square

Proof of Theorem 12.2. Since the tree-decomposition from (ii) displays $\partial_\Omega U$ and has essentially disjoint finite separators, it follows by standard arguments that not both (i) and (ii) can hold at the same time.

In order to show that at least one of (i) and (ii) holds, we prove $\neg(\text{i}) \rightarrow (\text{ii})$. For this, suppose that G contains no dominated comb attached to U . Using Theorem 12.2.5 we find a tree-decomposition $\mathcal{T}_{\text{DISJ}} = (T_{\text{DISJ}}, \mathcal{V}_{\text{DISJ}})$ of G with pairwise disjoint connected finite separators that displays the ends of G in the closure of U . Then the contraction minor H of G that is obtained from G by contracting every separator of $\mathcal{T}_{\text{DISJ}}$ does not contain any dominated comb attached to $[U]$ by Lemma 12.2.11. By Lemma 12.2.13 it suffices to show assertion (ii) with G and

U replaced by H and $[U]$. That is why in order to show assertion (ii) for G and U we may assume that the separators of $\mathcal{T}_{\text{DISJ}}$ are singletons.

By Theorem 12.1 we find a normal tree $T_{\text{NT}} \subseteq G$ that contains U cofinally and all whose rays are undominated. Furthermore, by the theorem's 'moreover' part we may choose T_{NT} so that every component of $G - T_{\text{NT}}$ has finite neighbourhood. As the nodes of T_{DISJ} whose parts meet T_{NT} induce a subtree T'_{DISJ} of T_{DISJ} , we may additionally assume that T_{NT} meets every part of $\mathcal{T}_{\text{DISJ}}$: we may replace $\mathcal{T}_{\text{DISJ}}$ with the tree-decomposition of G that corresponds to the S_{\aleph_0} -tree $(T'_{\text{DISJ}}, \alpha \upharpoonright \vec{E}(T'_{\text{DISJ}}))$ where $(T_{\text{DISJ}}, \alpha)$ is the S_{\aleph_0} -tree corresponding to $\mathcal{T}_{\text{DISJ}}$ (here Lemma 12.1.4 ensures that the new tree-decomposition still displays $\partial_{\Omega}U$).

As T_{NT} is normal, the neighbourhood of every such component C is a chain in T_{NT} and thus has a maximal element t_C . Now, let T' be the tree that is obtained from T_{NT} by adding every component C of $G - T_{\text{NT}}$ as a new vertex and joining it precisely to t_C . We define a tree-decomposition (T', \mathcal{V}') of G that is almost as desired.

Before we do that, let us have a closer look at how T_{NT} interacts with the tree-decomposition $\mathcal{T}_{\text{DISJ}}$, also see Figure 12.2.6. For every node $x \in T_{\text{DISJ}}$ the normal tree T_{NT} restricts to a normal tree $T_{\text{NT}}^x := T_{\text{NT}} \cap G[V_x]$ in $G[V_x]$ that contains all the vertices of U in the part V_x from $\mathcal{V}_{\text{DISJ}}$ cofinally. We write r_x for the root of T_{NT}^x . As the tree-decomposition $\mathcal{T}_{\text{DISJ}}$ of G displays all the ends in the closure of U , each T_{NT}^x must be rayless. The normal trees T_{NT}^x intersect each other as follows. For every two distinct nodes $x, y \in T_{\text{DISJ}}$ the normal trees T_{NT}^x and T_{NT}^y avoid each other if xy is not an edge of T_{DISJ} , and they intersect precisely in the single vertex of the separator associated with the edge xy if xy is an edge of T_{DISJ} .

Now let us define the parts V'_t of (T', \mathcal{V}') for every node $t \in T'$. For this, we choose for every node $t \in T_{\text{NT}}$ a root $r(t)$ of some of the normal trees T_{NT}^x with $x \in T_{\text{DISJ}}$ as follows. If just one of the normal trees T_{NT}^x contains t , then we let $r(t)$ be the root r_x of T_{NT}^x . Otherwise there are two normal trees T_{NT}^x and T_{NT}^y with $xy \in T_{\text{DISJ}}$ and we choose the smaller node of r_x and r_y with regard to the tree-order of T_{NT} as $r(t)$ (in particular, if $r_x < r_y$ then $r(r_y) = r_x$). For all nodes $t \in T_{\text{NT}} \subseteq T'$ we let V'_t be the vertex set of the decreasing path $tT_{\text{NT}}r(t)$ in T_{NT} . For newly added nodes $C \in T' - T_{\text{NT}}$ coming from components of $G - T_{\text{NT}}$ we let V'_C be the union of V'_{t_C} and the vertex set of the component C .

In a final construction, we obtain the desired tree-decomposition (T, \mathcal{V}) from




Figure 12.2.6.: The construction of (T', \mathcal{V}') in the proof of Theorem 12.2. The tree depicted is the normal tree T_{NT} and the grey disks are the parts of $\mathcal{T}_{\text{DISJ}}$. Here the root r_x of T_{NT}^x agrees with the root of T_{NT} . Also we have $r(r_y) = r(r_z) = r_x$ and $r(t) = r_y$.

the tree-decomposition (T', \mathcal{V}') . For every vertex $x \in T_{\text{DISJ}}$ let T_x be the tree that is obtained from T_{NT}^x as follows: Take a copy s_x of r_x (making sure that $s_x \notin T_{\text{NT}}$ and $s_x \neq s_y$ for all $x \neq y \in T_{\text{DISJ}}$) and join it precisely to the neighbours of r_x in T_{NT}^x and to r_x . Then delete all edges incident to r_x other than $r_x s_x$. We let T be the union of all the trees T_x and define the parts of (T, \mathcal{V}) as follows. For every node $t \in V(T') \subseteq V(T)$ we let $V_t := V_t'$ and for all vertices $s_x \in T - T'$ we let V_{s_x} be the singleton consisting only of r_x . Let us prove that (T, \mathcal{V}) is as desired. Each part contains at most finitely many vertices from U because $U \subseteq V(T_{\text{NT}})$ and $V_t \cap T_{\text{NT}}$ is the vertex set of a finite path (or a singleton) for every node $t \in T$. Quite similarly, all parts at non-leaves of T' are finite because they are vertex sets of finite paths of T_{NT} .

To see that (T, \mathcal{V}) has essentially disjoint separators, let $F \subseteq E(T)$ be the set of all edges $r_x s_x$ with $x \in T_{\text{DISJ}}$ and r_x distinct from the root of T_{NT} . The latter requirement becomes necessary when the root of T_{NT} forms a separator Z of $\mathcal{T}_{\text{DISJ}}$: then the root is chosen as $r_x = r_y$ for the edge $xy \in T_{\text{DISJ}}$ with which the separator

Z is associated in $\mathcal{T}_{\text{DISJ}}$, meaning that both edges $r_x s_x$ and $r_y s_y$ of T have the same separator $\{r_x\} = \{r_y\}$ associated with them in (T, \mathcal{V}) . In particular, the requirement affects at most two edges of T . Now, let us see that F witnesses that (T, \mathcal{V}) has essentially disjoint separators. On the one hand, the separators of (T, \mathcal{V}) associated with edges $r_x s_x \in F$ are singletons of the form $\{r_x\}$ and thus are pairwise disjoint. On the other hand, using that the trees T_{NT}^x with $x \in T_{\text{DISJ}}$ are rayless, it is easy to see that every ray $R \subseteq T$ passes through infinitely many edges from F .

In order to see that (T, \mathcal{V}) displays the ends in the closure of U it suffices to show that (T', \mathcal{V}') displays the ends in the closure of U . For this in turn, by Lemma 12.1.4, it suffices to show that (T', \mathcal{V}') displays the ends in the closure of T_{NT} , which follows from standard arguments. \square

Example 12.2.14. The tree-decomposition in Theorem 12.2 (ii) cannot be chosen with pairwise disjoint separators instead of essentially disjoint separators: Suppose that G consists of the first three levels of T_{\aleph_0} and let $U := V(G)$. Then G contains no comb attached to U . In particular, as we have already argued in the text below Theorem 12.2, every tree-decomposition (T, \mathcal{V}) of G complementary to dominated combs as in Theorem 12.2 is also a tree-decomposition of G complementary to combs as in Theorem 12.2.1. But then (T, \mathcal{V}) cannot be chosen with pairwise disjoint separators, as pointed out in Example 11.2.7.

13. Undominated combs

In the first Chapter 11 of this part we found structures whose existence is complementary to the existence of a star or a comb attached to a given set U of vertices, and two types of these structures turned out to be relevant for both stars and combs: normal trees and tree-decompositions. A *comb* is the union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R . The last vertices of those paths are the *teeth* of this comb. Given a vertex set U , a *comb attached to U* is a comb with all its teeth in U , and a *star attached to U* is a subdivided infinite star with all its leaves in U . Then the set of teeth is the *attachment set* of the comb, and the set of leaves is the *attachment set* of the star. Given a graph G , a rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T , cf. [20]. For the definition of tree-decompositions see [20].

As stars and combs can interact with each other, this is not the end of the story. For example, a given vertex set U might be connected in a graph G by both a star and a comb, even with infinitely intersecting sets of leaves and teeth. To formalise this, let us say that a subdivided star S *dominates* a comb C if infinitely many of the leaves of S are also teeth of C . A *dominating star* in a graph G then is a subdivided star $S \subseteq G$ that dominates some comb $C \subseteq G$; and a *dominated comb* in G is a comb $C \subseteq G$ that is dominated by some subdivided star $S \subseteq G$. Thus, a comb $C \subseteq G$ is undominated in G if it is not dominated in G . Recall that a vertex v of G *dominates* a ray $R \subseteq G$ if there is an infinite v – $(R - v)$ fan in G , see [20]. A ray $R \subseteq G$ is *dominated* if some vertex of G dominates it. Rays not dominated by any vertex of G are *undominated*. Dominated combs are related to dominated rays in that a comb is dominated in G if and only if its spine is dominated in G .

In Chapter 12 we determined structures whose existence is complementary to the existence of dominating stars or dominated combs—again in terms of normal trees or tree-decompositions.

Here, in this chapter, we determine structures whose existence is complementary

to the existence of undominated combs. A candidate for a normal tree that is complementary to an undominated comb in G attached to a given set U of vertices is a normal tree $T \subseteq G$ that contains U and all whose rays are dominated in G , for if $U = V(G)$ then T is spanning and hence its (dominated) rooted rays are in a natural one-to-one correspondence to the ends of G . Such normal trees T are easily seen to be complementary structures for undominated combs whenever G happens to contain some normal tree that contains U . But in general, normal trees $T \subseteq G$ containing U all whose rays are dominated in G are not complementary to undominated combs, because the absence of an undominated comb does not imply the existence of such a normal tree: for example if G is an uncountable complete graph and $U = V(G)$, then every normal tree in G containing U must be spanning but G does not have any normal spanning tree.

As our first main result, we show that if U is contained in any normal tree $T \subseteq G$, there is a more elementary structure that is complementary to undominated combs attached to U and which obstructs undominated combs attached to U immediately: a rayless tree containing U . Call a set $U \subseteq V(G)$ of vertices of a graph G *normally spanned* in G if U is contained in a tree $T \subseteq G$ that is normal in G . The graph G is *normally spanned* if $V(G)$ is normally spanned in G , i.e., if G has a normal spanning tree.

Theorem 13.1. *Let G be any graph and let $U \subseteq V(G)$ be normally spanned in G . Then the following assertions are complementary:*

- (i) G contains an *undominated comb* attached to U ;
- (ii) there is a *rayless tree* $T \subseteq G$ that contains U .

This extends results of Polat [56, 57] and Širáň [73], who proved the case $U = V(G)$ for countable G : *A countable connected graph has a rayless spanning tree if and only if all its rays are dominated.*

There are uncountable graphs G for which this duality fails, even for $U = V(G)$. By Theorem 13.1, such graphs G cannot have a normal spanning tree. There are two known constructions of such graphs, by Seymour and Thomas [65] and by Thomassen [70]. Both these constructions are involved.

As a corollary of Theorem 13.1 we obtain a full characterisation of the graphs that contain a rayless tree containing a given set U of vertices: they are precisely

the graphs G that have a subgraph H in which U is normally spanned and all whose rays are dominated in H . In particular, we obtain the following corollary:

Corollary 13.2. *Graphs with a normal spanning tree have a rayless spanning tree if and only if all their rays are dominated.*

The graphs with a normal spanning tree are well studied and are quite well known: see [29, 43].

Our duality theorem for undominated combs in terms of rayless trees, Theorem 13.1, has two applications, Theorems 13.3 and 13.5 below. In order to state our first application we need the following notation for arbitrary graphs G . Suppose that H is any subgraph of G and $\varphi: \Omega(H) \rightarrow \Omega(G)$ is the natural map satisfying $\eta \subseteq \varphi(\eta)$ for every end η of H . Furthermore suppose that a set $\Psi \subseteq \Omega(G)$ of ends of G is given. We say that H is *end-faithful* for Ψ if $\varphi \upharpoonright \varphi^{-1}(\Psi)$ is injective and $\text{im}(\varphi) \supseteq \Psi$. And H *reflects* Ψ if φ is injective with $\text{im}(\varphi) = \Psi$. An end of G is *dominated* and *undominated* if one (equivalently: each) of its rays is dominated and undominated, respectively (see [20]).

Carmesin [15] proved that every connected graph G has a spanning tree that is end-faithful for the undominated ends of G . He also pointed out that his result becomes false when ‘end-faithful’ is replaced with ‘reflecting’. As our first application of Theorem 13.1 we characterise the graphs that have spanning trees reflecting their undominated ends. An end ω of G is contained *in the closure* of a vertex set $U \subseteq V(G)$ in G if G contains a comb attached to U whose spine lies in ω .

Theorem 13.3. *Let G be any graph and let $U \subseteq V(G)$ be any vertex set. Then the following assertions are equivalent:*

- (i) *There exists a tree $T \subseteq G$ that contains U and reflects the undominated ends of G in the closure of U in G ;*
- (ii) *G has a subgraph H with $U \subseteq V(H)$ normally spanned in H and all whose undominated ends are included in distinct undominated ends of G .*

Corollary 13.4. *Every graph that has a normal spanning tree does have a spanning tree reflecting its undominated ends.*

As a consequence of the star-comb lemma, every spanning tree of a graph G contains a ray from every undominated end of G . Thus, rayless spanning trees

always reflect the undominated ends of the graphs they span. In this sense, spanning trees reflecting the undominated ends can be seen as a generalisation of rayless spanning trees.

Spanning trees reflecting the undominated ends are particularly interesting for finitely separable graphs. A graph is *finitely separable* if every two of its vertices can be separated by finitely many edges, cf. [2]. Our second application of Theorem 13.1 reads as follows:

Theorem 13.5. *Let G be any graph and let $T \subseteq G$ be any spanning tree.*

- (i) *All the fundamental cuts of T are finite if and only if G is finitely separable and T reflects the undominated ends of G .*
- (ii) *If G is finitely separable, then it has a spanning tree all whose fundamental cuts are finite.*

For a finitely separable graph G , the spanning trees of G all whose fundamental cuts are finite are precisely the spanning trees of G whose closure in $\tilde{G} = (\tilde{G}, \text{ITOP})$ contains no (topological) cycle, see [2] for definitions. The space \tilde{G} was used by Bruhn and Diestel [2] to extend Whitney’s theorem [20, 71]—which states that a finite graph is planar if and only if it has an abstract dual—to finitely separable infinite graphs. Bruhn and Diestel also showed that \tilde{G} permits the extension of another well known duality theorem for finite graphs: that the complement of the edge set of any spanning tree of G defines a spanning tree in any abstract dual of G , and conversely that any two graphs with the same edge sets so that their spanning trees complement each other form a pair of abstract duals. Their latter extension speaks of spanning trees whose closure in \tilde{G} contains no (topological) cycle instead of arbitrary spanning trees. Solving a problem of Diestel and Kühn [28, Problem 7.9], they showed that such spanning trees always exist in connected finitely separable graphs. Our Theorem 13.5 provides an alternative proof:

Corollary 13.6. *Every connected finitely separable graph G has a spanning tree whose closure in \tilde{G} contains no topological cycle.*

In contrast to Bruhn and Diestel’s proof, ours is rather methodic in that it combines various structural results.

Let us return to our initial problem of finding complementary structures for undominated combs. While it is not always possible to find normal trees or rayless trees that are complementary to undominated combs, it turns out that suitable tree-decompositions still serve as complementary structures:

Theorem 13.7. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains an *undominated comb* attached to U ;
- (ii) G has a *star-decomposition* with finite adhesion sets such that U is contained in the central part and all undominated ends of G live in the leaves' parts.

Moreover, we may assume that the adhesion sets of the tree-decomposition in (ii) are pairwise disjoint and connected.

As discussed above, rayless trees are in general too strong to serve as complementary structures for undominated combs. It turns out that less specific structures than rayless trees, subgraphs all of whose rays are dominated, yield another complementary structure for undominated combs:

Theorem 13.8. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains an *undominated comb* attached to U ;
- (ii) G has a connected *subgraph* that contains U and all whose rays are dominated in it.

Moreover, the subgraph H in (ii) can be chosen so as to reflect the ends in the closure of H .

This chapter is organised as follows. In Section 13.1, we prove our duality theorem for undominated combs in terms of rayless trees, Theorem 13.1. In Section 13.2, we discuss our applications of this duality theorem, i.e., we prove Theorem 13.3 and Theorem 13.5. In Section 13.3, we provide our two full duality theorems for undominated combs: Theorem 13.7 and Theorem 13.8.

We assume familiarity with the tools and terminology described in the first chapter of this part.

13.1. Rayless trees

In this section, we will consider rayless trees as structures that are complementary to undominated combs. As usual, let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. There are three reasons why rayless trees containing U are good candidates. First, an undominated comb attached to U is more specific than a comb attached to U and in Chapter 11 (Theorem 11.1) we proved that rayless normal trees $T \subseteq G$ that contain U are complementary to combs. Therefore, structures that are complementary to undominated combs should be less specific than such normal trees.

Second, by the star-comb lemma, G containing no undominated comb attached to U can be rephrased as follows: for every infinite subset $U' \subseteq U$ the graph G contains a star attached to U' . So combining such stars in a clever way might lead to a rayless tree containing U .

Finally, a graph cannot contain both an undominated comb attached to U and a rayless tree containing U at the same time:

Lemma 11.1.4. *If U is an infinite set of vertices in a rayless rooted tree T , then T contains a star attached to U which is contained in the up-closure of its central vertex in the tree-order of T .*

For $U = V(G)$, Širáň [73] conjectured that G having a rayless spanning tree is complementary to G containing an undominated comb attached to U . Surprisingly, his conjecture has turned out to be false, as shown by Seymour and Thomas [65]. The counterexample they have found is also a big surprise. Recall that T_κ for a cardinal κ denotes the tree all whose vertices have degree κ .

Theorem 13.1.1 ([65, Theorem 1.6]). *There is an infinitely connected, in particular one-ended, graph G of order 2^{\aleph_0} which does not contain a subdivided K^{\aleph_1} , such that every spanning tree of G contains a subdivision of T_{\aleph_1} .*

Indeed, the end of a graph G as in Theorem 13.1.1 is dominated as G is infinitely connected, but for $U = V(G)$ the graph does not contain a rayless tree that contains U .

A similar counterexample has been obtained independently by Thomassen [70]. Set-theoretic points of view are presented in [65] and Komjáth's [45]. Komjáth even gives a positive consistency result under Martin's axiom for graphs G with

$< 2^{\aleph_0}$ many vertices: If $\kappa < 2^{\aleph_0}$ is a cardinal, $MA(\kappa)$ holds, and G is infinitely connected with $|V(G)| \leq \kappa$, then G has a rayless spanning tree.

Nevertheless, it is known that requiring G to be countable does suffice to ensure the existence of a rayless spanning tree when G is connected and every end is dominated, giving the following duality:

Theorem 13.1.2. *Let G be any connected countable graph. Then the following assertions are complementary:*

- (i) G contains an undominated comb attached to $V(G)$;
- (ii) G has a rayless spanning tree.

Proofs are due to Polat [56, 57] and Širáň [73]. Our main result in this section extends Theorem 13.1.2:

Theorem 13.1. *Let G be any graph and let $U \subseteq V(G)$ be normally spanned in G . Then the following assertions are complementary:*

- (i) G contains an undominated comb attached to U ;
- (ii) there is a rayless tree $T \subseteq G$ that contains U .

Note that this extends Theorem 13.1.2 twofold: On the one hand, we localise the statement to an arbitrary vertex set $U \subseteq V(G)$. On the other hand, we extend the statement to the class of all graphs in which U is normally spanned.

While our focus in this chapter is to find duality theorems for undominated combs, Polat and Širáň were rather interested in a characterisation of those graphs that have rayless spanning trees. The strongest sufficient condition for the existence of a rayless spanning tree, other than Theorem 13.1 (to the knowledge of the authors), is due to Polat [60]: *If every end of a connected graph G is dominated and G contains no subdivided T_{\aleph_1} , then G has a rayless spanning tree.* His result does not imply our Theorem 13.1, for example consider G to be the graph obtained from T_{\aleph_1} by completely joining an arbitrarily chosen root to all other nodes, and $U = V(G)$. However, as a corollary of Theorem 13.1, we obtain a full characterisation of the graphs that have rayless spanning trees. Our characterisation even takes an arbitrary vertex set $U \subseteq V(G)$ into account:

Corollary 13.1.3. *Let G be any graph. Then the following assertions are equivalent:*

- (i) *There is a rayless tree $T \subseteq G$ that contains U ;*
- (ii) *G has a subgraph H in which $U \subseteq V(H)$ is normally spanned and all whose rays are dominated in H .* □

If the graph G itself has a normal spanning tree, then our characterisation simplifies as follows:

Corollary 13.2. *Graphs with a normal spanning tree have a rayless spanning tree if and only if all their rays are dominated.* □

This section is organised as follows. In Section 13.1.1 we will prove Theorem 13.1 for normally spanned graphs. Then, in Section 13.1.2, we will deduce Theorem 13.1.

13.1.1. Proof for normally spanned graphs

As a first approximation to Theorem 13.1 we prove the following:

Theorem 13.1.4. *Let G be any normally spanned graph and let $U \subseteq V(G)$ be any vertex set. Then the following assertions are complementary:*

- (i) *G contains an undominated comb attached to U ;*
- (ii) *G contains a rayless tree that contains U .*

Our proof consists of three key ideas, organised in the following three lemmas: Lemma 11.1.13, Lemma 13.1.5 and Lemma 13.1.6.

Lemma 11.1.13. *Let G be any graph. If $T \subseteq G$ is a rooted tree that contains a vertex set W cofinally, then $\partial_\Omega T = \partial_\Omega W$.*

Lemma 13.1.5. *Let G be any graph and let $U \subseteq V(G)$ be any vertex set. If \hat{U} is the superset of U also containing all the vertices dominating an end in the closure of U , then $\partial_\Omega \hat{U} = \partial_\Omega U$. In particular, $\partial_\Omega U' = \partial_\Omega U$ for all vertex sets U' with $U \subseteq U' \subseteq \hat{U}$ and \hat{U} contains all the vertices dominating an end in the closure of \hat{U} .*

Proof. Every end in the closure of U is contained in the closure of \hat{U} because \hat{U} contains U . For the other inclusion consider any end ω in the closure of \hat{U} . Given a finite vertex set $X \in \mathcal{X}$ we show that $C(X, \omega)$ contains a vertex from U . Fix a comb attached to \hat{U} and with spine in ω , and pick any tooth v of the comb in the component $C(X, \omega)$ of $G - X$. Then either v is contained in U , or v dominates an end ω' in the closure of U in which case U must meet $C(X, \omega') = C(X, \omega)$. Therefore, $C(X, \omega)$ meets U for all $X \in \mathcal{X}$, and so ω lies in the closure of U . \square

For our last key lemma, we shall need the following result of Jung, which we recall here:

Theorem 11.2.5. *Let G be any graph. A vertex set $W \subseteq V(G)$ is normally spanned in G if and only if it is a countable union of dispersed sets. In particular, G is normally spanned if and only if $V(G)$ is a countable union of dispersed sets.*

Lemma 13.1.6. *Let G be any graph and let $U \subseteq V(G)$ be normally spanned. If every end in the closure of U is dominated by some vertex in U , then there is a rayless tree $T \subseteq G$ containing U .*

Normal trees follow the concept of depth-first search trees. Speaking informally, all ends of G are ‘far away’ from the perspective of any fixed vertex. This is why normal spanning trees grow towards the ends of the underlying graph in the sense that they contain (precisely) one normal ray from every end. We, however, seek to avoid having any rays in our tree. This is why our construction of a rayless tree containing U will follow the opposite concept of depth-first search trees, namely that of breadth-first search trees.

Proof of Lemma 13.1.6. First we choose a well-ordering of U all whose proper initial segments are dispersed: By Theorem 11.2.5, we have that U is a countable union $\bigcup_{n \in \mathbb{N}} U_n$ of, say pairwise disjoint, dispersed sets U_n . Choose a well-ordering \preceq_n of every vertex set U_n . Given $u, u' \in U$ with $u \in U_m$ and $u' \in U_n$, we write $u \preceq u'$ if either $m < n$ or $m = n$ with $u \preceq_m u'$ holds. It is straightforward to show that \preceq defines a well-ordering of U that is as desired. From now on we view U as well-ordered set (U, \preceq) .

We recursively construct an ascending sequence $(T_\alpha)_{\alpha < \kappa}$ of rooted trees T_α sharing their root and satisfying that the overall union of the T_α is a rayless tree containing U . Let T_0 be the tree consisting of and rooted in the smallest

vertex of U . In a limit step $\beta > 0$ we let T_β be the tree $\bigcup \{T_\alpha \mid \alpha < \beta\}$. In a successor step $\beta = \alpha + 1$ we terminate and set $\kappa = \beta$ if U is included in T_α . Otherwise we let u be the smallest vertex in $U \setminus V(T_\alpha)$. Following the concept of a breadth-first search tree, among all u - T_α paths fix one P_β whose endvertex in T_α has minimal height in T_α . We obtain T_β from T_α by adding the path P_β .

Let T be the overall union of the trees T_α , i.e., $T := \bigcup \{T_\alpha \mid \alpha < \kappa\}$. Then T is a rooted tree that contains U cofinally. It remains to check that T is rayless. Suppose for a contradiction that R is a ray in T starting in the root, say. By Lemma 11.1.13 the end of the ray R is contained in the closure of U . As all ends in $\partial_\Omega U$ are dominated by vertices in U , we find a vertex $u^* \in U$ dominating R . Let P_{α^*} be the path from the construction of T that added u^* .

We claim that every tree T_α meets R in a finite initial subpath. This can be seen as follows. Since all proper initial segments of U are dispersed, by Lemma 11.1.13 it suffices to show that every T_α with $\alpha > 0$ contains a subset of such a segment cofinally. A transfinite induction on α shows that for T_α this subset may be chosen as the set of starting vertices of the paths P_ξ with $\xi \leq \alpha$ a successor ordinal while the proper initial segment may be chosen as the down-closure in U of the starting vertex of $P_{\alpha+1}$. Here we remark that $\alpha + 1 < \kappa$ for all $\alpha < \kappa$ (i.e. κ is a limit ordinal): indeed, by our assumption that $R \subseteq T$ we know that the vertex set U is not dispersed and, therefore, meets infinitely many U_n .

Finally, we derive the desired contradiction. Fix $\beta > \alpha^*$ so that the endvertex x of $P_{\beta+1}$ in T_β has larger height than u^* has in T_β and so that $P_{\beta+1}$ contains an edge of R . Let u be the first vertex of $P_{\beta+1}$, i.e., the smallest vertex in $U \setminus V(T_\beta)$. Note that the first vertex w of $P_{\beta+1}$ that is contained in R is distinct from x . (Also see Figure 13.1.1.) As u^* dominates R we find an infinite set \mathcal{Q} of u^* - R paths in G such that distinct paths in \mathcal{Q} only meet in u^* . All but finitely many paths in \mathcal{Q} meet $T_{\beta+1}$ precisely in u^* : Otherwise the end of R is contained in the closure of $T_{\beta+1}$ contradicting that the vertex set of $T_{\beta+1}$ is dispersed. Fix a path $Q \in \mathcal{Q}$ meeting $T_{\beta+1}$ precisely in u^* and having its endvertex v in $\mathring{w}R$. We conclude that $uP_{\beta+1}wRvQu^*$ would have been a better choice than $P_{\beta+1}$ in the construction of $T_{\beta+1}$ (contradiction). \square

Proof of Theorem 13.1.4. By Lemma 11.1.4 at most one of (i) and (ii) holds at a time. To verify that least one of (i) and (ii) holds, we show $\neg(i) \rightarrow (ii)$. By




Figure 13.1.1.: The situation in the last paragraph of the proof of Lemma 13.1.6.

Lemma 13.1.5 we may assume that U contains all vertices dominating an end in the closure of U , and by Lemma 13.1.6 there is a rayless tree $T \subseteq G$ that contains U . \square

13.1.2. Deducing our duality theorem in terms of rayless trees

Let us analyse why the proof of our duality theorem for undominated combs in terms of rayless trees for normally spanned graphs, Theorem 13.1.4, does not immediately give a proof for arbitrary graphs. For this, consider any graph G and let $U \subseteq V(G)$ be any vertex set. Furthermore, suppose that there is a normal tree $T \subseteq G$ that contains U and that G contains no undominated comb attached to U . In the proof of Theorem 13.1.4 we assume without loss of generality that U contains all the vertices dominating an end in the closure of U . This is possible because, by Lemma 13.1.5, adding all the vertices to U that dominate an end in the closure of U does not change the set $\partial_{\Omega}U$ of ends in the closure of U . However, after adding all these vertices it may happen—in contrast to the situation in the proof of Theorem 13.1.4 where G has a normal spanning tree—that U is no longer normally spanned in G (e.g. consider any countably infinite set U of vertices in an uncountable complete graph). And U being normally spanned in G is a crucial

requirement of the lemma that yields the desired rayless tree, Lemma 13.1.6.

But maybe adding all the vertices that dominate an end in the closure of U and maintaining that U is normally spanned was too much to ask. Indeed, Lemma 13.1.6 only requires that U contains for every end $\omega \in \partial_\Omega U$ at least one vertex dominating ω , and adding just one dominating vertex for every end ω might preserve the property of U being normally spanned in G . The following example shows that this is in general false:

Example 13.1.7. Let G be a *binary tree with tops*, i.e., let G be obtained from the rooted infinite binary tree T_2 by adding for every normal ray R of T_2 a new vertex v_R , its *top*, that is joined completely to R (cf. Diestel and Leader's [29]). Let U be the vertex set of T_2 . Then $\partial_\Omega U = \Omega(G)$ and every end ω is dominated precisely by the top that was added for the unique normal ray of T_2 that is contained in ω . Hence adding for every end in $\partial_\Omega U$ a vertex dominating it to U results in the whole vertex set of G . However, as pointed out in [29], the graph G does not have a normal spanning tree.

Our way out is to work in a suitable contraction minor, which requires some preparation: Let H and G be any two graphs. We say that H is a contraction minor of G with *fixed branch sets* if an indexed collection $\{V_x \mid x \in V(H)\}$ of branch sets is fixed to witness that G is an IH . In this case, we write $[v] = [v]_H$ for the branch set V_x containing a vertex v of G and also refer to x by $[v]$. Similarly, we write $[U] = [U]_H := \{[u] \mid u \in U\}$ for vertex sets $U \subseteq V(G)$.

Lemma 13.1.8. *Let G be any graph and let H be any contraction minor of G with fixed branch sets that induce subgraphs of G with rayless spanning trees. Furthermore, let $U \subseteq V(G)$ be any vertex set. If H contains a rayless tree that contains $[U]$, then G contains a rayless tree that contains U .*

Proof. Let $T \subseteq H$ be a rayless tree that contains $[U]$. Fix for every branch set $W \in [V(T)]$ a rayless spanning tree T_W in the subgraph that G induces on W . Furthermore, select one edge $e_f \in E_G(t_1, t_2)$ for every edge $f = t_1 t_2 \in T$. It is straightforward to show that the union of all the trees T_W plus all the edges e_f is a rayless tree in G that contains U . \square

Let H be a contraction minor of a graph G with fixed branch sets. A subgraph $G' = (V', E')$ of G can be passed on to H as follows. Take as vertex set the set

$[V']$ and declare W_1W_2 to be an edge whenever E' contains an edge between W_1 and W_2 . We write $[G'] = [G']_H$ for the resulting subgraph of H and call it the graph that is obtained by *passing on* G' to H . If every vertex $W \in [V']$ meets V' in precisely one vertex, then we say that G' is *properly passed on* to H . Note that if G' is properly passed on to H , then $[G']$ and G' are isomorphic.

Lemma 13.1.9. *Let H be a contraction minor of a graph G with fixed branch sets and let $T \subseteq G$ be a tree that is normal in G . If T is properly passed on to H , then $[T] \subseteq H$ is a tree that is normal in H .*

Proof. Since T is properly passed on to G we have that T and $[T]$ are isomorphic as witnessed by the bijection φ that maps every vertex $t \in T$ to $[t]$. In order to see that $[T]$ is normal in H when it is rooted in $[r]$ for the root r of T , consider any $[T]$ -path $W_0 \dots W_k$ in $[H]$. Using that branch sets are connected, it is straightforward to show that there is T -path in G between the two vertices $\varphi^{-1}(W_0)$ and $\varphi^{-1}(W_k)$ of T . Hence W_0 and W_k must be comparable in $[T]$. \square

We need two more lemmas for the proof of Theorem 13.1. Recall that the *generalised up-closure* $\llbracket x \rrbracket$ of a vertex $x \in T$ is the union of $[x]$ with the vertex set of $\bigcup \mathcal{C}(x)$, where the set $\mathcal{C}(x)$ consists of those components of $G - T$ whose neighbourhoods meet $[x]$.

Lemma 11.1.10. *Let G be any graph and $T \subseteq G$ any normal tree.*

- (i) *Any two vertices $x, y \in T$ are separated in G by the vertex set $[x] \cap [y]$.*
- (ii) *Let $W \subseteq V(T)$ be down-closed. Then the components of $G - W$ come in two types: the components that avoid T ; and the components that meet T , which are spanned by the sets $\llbracket x \rrbracket$ with x minimal in $T - W$.*

Lemma 11.1.11. *If G is any graph and $T \subseteq G$ is any normal tree, then every end of G in the closure of T contains exactly one normal ray of T . Moreover, sending these ends to the normal rays they contain defines a bijection between $\partial_\Omega T$ and the normal rays of T .*

Proof of Theorem 13.1. Given a normally spanned vertex set $U \subseteq V(G)$ we have to show that the following assertions are complementary:

- (i) G contains an undominated comb attached to U ;
- (ii) G contains a rayless tree that contains U .

By Lemma 11.1.4 at most one of (i) and (ii) holds at a time. To verify that at least one of (i) and (ii) holds, we show $\neg(\text{i}) \rightarrow (\text{ii})$. For this, we may assume by Lemma 11.1.13 that U is the vertex set of a normal tree $T \subseteq G$. In the following we will find a contraction minor H of G with fixed branch sets V_x such that:

- all $G[V_x]$ have rayless spanning trees;
- T is properly passed on to H ;
- and every end of H in the closure of $[T] \subseteq H$ is dominated in H by some vertex of $[T]$.

Before we prove that such H exists, let us see how to complete the proof once H is found. By Lemma 13.1.9, the tree $[T]$ is normal in H , and it has vertex set $[U]$ because $V(T) = U$. So, by Lemma 13.1.6, the graph H contains a rayless tree that contains $[U]$. Finally, by Lemma 13.1.8, this rayless tree in H containing $[U]$ gives rise to a rayless tree in G containing U as desired.

In order to construct H , fix for every normal ray R of T a vertex v_R dominating R in G . Let \mathcal{R} be the set of all normal rays R of T for which v_R is contained in a component C_R of $G - T$. Note that the down-closure of the neighbourhood of each C_R is $V(R)$ due to the separation properties of normal trees (Lemma 11.1.10). Thus, we have $C_R \neq C_{R'}$ for distinct normal rays $R, R' \in \mathcal{R}$. Fix a v_R - R path P_R for every $R \in \mathcal{R}$. Then the overall union of the paths P_R is a forest of subdivided stars, each having its centre on T . Let us refer by S_R to the subdivided star that contains v_R for $R \in \mathcal{R}$, i.e., S_R is the union of all the paths $P_{R'}$ that contain the last vertex of P_R and this last vertex is the centre of S_R . Let H be the contraction minor of G with fixed branch sets defined as follows: if v is contained on a path P_R , then put $[v] := S_R$; otherwise let $[v] := \{v\}$. Then, in particular, every branch set of H induces a subgraph of G that has a rayless spanning tree.

As every star S_R meets T precisely in its centre, the tree T is properly passed on to H . By Lemma 13.1.9, the tree $[T] \subseteq H$ is normal in H and $V([T]) = [U]$ since $V(T) = U$. And by Lemma 11.1.11 it remains to show that every normal ray of $[T]$ is dominated in H by some vertex of $[T]$. For this, we consider three cases.

In all three cases, fix any normal ray $R \subseteq T$ and some collection \mathcal{P} of infinitely many v_R - R paths in G meeting precisely in v_R .

First assume that $R \in \mathcal{R}$. Note that only finitely many of the paths in \mathcal{P} meet $v_R^\circ P_R$, without loss of generality none. Then all graphs $[P] \subseteq H$ with $P \in \mathcal{P}$ are $[v_R]$ - $[R]$ paths that meet only in $[v_R]$. This shows that $[v_R] \in [T]$ dominates $[R]$ in H .

Second, suppose that $R \notin \mathcal{R}$ and that every branch set of H other than $[v_R]$ meets only finitely many of the paths in \mathcal{P} . By thinning out \mathcal{P} we may assume that every branch set other than $[v_R]$ meets at most one of the paths in \mathcal{P} . Then the connected graphs $[P]$ with $P \in \mathcal{P}$ pairwise meet in $[v_R]$ but nowhere else and all contain a vertex of $[R]$ other than $[v_R]$. Taking one $[v_R]$ - $([R] - [v_R])$ path inside each $[P]$ yields a fan witnessing that $[v_R] \in [T]$ dominates $[R]$ in H .

Finally, suppose that $R \notin \mathcal{R}$ and that some branch set $S \neq [v_R]$ of H meets infinitely many of the paths in \mathcal{P} , say all of them. We write c for the centre of S . Without loss of generality none of the paths in \mathcal{P} contains c . Also note that c is contained in $V(R)$ as otherwise all the paths in \mathcal{P} need to pass through the finite down-closure of c in T in vertices other than v_R . Let \mathcal{R}' be the collection of normal rays of T that satisfies $S = \bigcup \{V(P_{R'}) \mid R' \in \mathcal{R}'\}$. For every v_R - R path $P \in \mathcal{P}$ let v_P be the last vertex on P that is contained in S , let w_P be the first vertex on P after v_P in which P meets T and let Q_P be the unique w_P - R path in T . (See Figure 13.1.2.) For every path $P \in \mathcal{P}$ let $P' = P'(P) := v_P P w_P Q_P$, and let $\mathcal{P}' = \mathcal{P}'(\mathcal{P}) := \{P' \mid P \in \mathcal{P}\}$.

Each path $P_{R'} \subseteq S$ with $R' \in \mathcal{R}'$ meets only finitely many paths from \mathcal{P}' , and these latter paths are precisely the paths in \mathcal{P}' that meet $C_{R'}$: This is because every path in \mathcal{P}' that meets $C_{R'}$ starts in a vertex $v_P \in C_{R'}$ and after leaving $C_{R'}$ only traverses through vertices of T . Therefore, by replacing \mathcal{P} with an infinite subset of \mathcal{P} , we can see to it that every component $C_{R'}$ with $R' \in \mathcal{R}'$ meets at most one of the paths in the then smaller set $\mathcal{P}' = \mathcal{P}'(\mathcal{P})$. In countably many steps we fix paths P'_1, P'_2, \dots in \mathcal{P}' so that their last vertices are pairwise distinct: In order to see that this is possible suppose for a contradiction that $t \in R$ is maximal in the tree order of T so that t is the last vertex of a path in \mathcal{P}' . Note that R together with the paths $v_P P$ with $P \in \mathcal{P}$ forms a comb in G . Hence infinitely many of the paths $v_P P$ are contained in the same component of $G - [t]$ as some tail of R . By Lemma 11.1.10, this component is of the form $\llbracket t' \rrbracket$ for the




Figure 13.1.2.: The final case in the proof of our duality theorem for undominated combs in term of rayless trees.

successor t' of t on R . In particular, we find some $P \in \mathcal{P}$ so that w_P lies above t' in the tree order of T . But then the endvertex of Q_P in R lies above t' and, in particular, above t , contradicting the choice of t .

So let P'_1, P'_2, \dots be paths in \mathcal{P}' with pairwise distinct last vertices. We show that the paths P'_i give rise to S - $[R]$ paths $[P'_i]$ in H that form an infinite S - $[R]$ fan witnessing that S dominates $[R]$ in H . Every path P'_i is an S - R path because every path in \mathcal{P}' is an S - R path by the choice of the vertices v_P . Moreover, the paths P'_i are pairwise disjoint: Every path P'_i starts in a component $C_{R'}$. Using the choice of the vertices v_P with $P \in \mathcal{P}$ as the last vertex on P that is contained in S we have that the $[P'_i]$ are S - $[R]$ paths of H that only share their first vertex S . Hence the $[P'_i]$ form an infinite S - R fan in H and we conclude that $S \in [T]$ dominates $[R]$ in H . \square

13.2. Spanning trees reflecting the undominated ends

In [37], Halin conjectured that every connected graph has a spanning tree that is end-faithful for all its ends. However, Seymour and Thomas' counterexample in Theorem 13.1.1 shows that his conjecture is in general false. Recently, Carmesin [15] amended Halin's conjecture by proving the following:

Theorem 13.2.1 (Carmesin 2014). *Every connected graph G has a spanning tree that is end-faithful for the undominated ends of G .*

Carmesin pointed out that his theorem is best possible in that it becomes false when one replaces 'is end-faithful for' with the more specific 'reflects' in its wording: by Theorem 13.1.1 there are connected graphs without rayless spanning trees all whose rays are dominated. Characterising the graphs that have spanning trees reflecting their undominated ends has remained an open problem, until today.

Our aim in this section is threefold. Our first goal is to prove Theorem 13.3 below which characterises the graphs that have spanning trees reflecting their undominated ends. Thereafter, we will characterise in Theorem 13.5 (i) the spanning trees of finitely separable graphs that reflect the undominated ends, and we will establish in Theorem 13.5 (ii) that every connected finitely separable graph has such a tree. Finally, we will deduce Corollary 13.6 which states that every connected finitely separable graph G has a spanning tree whose closure in \tilde{G} contains no topological cycle.

Our characterisation of the graphs that have a spanning tree reflecting their undominated ends even takes an arbitrary vertex set U into account:

Theorem 13.3. *Let G be any graph and let $U \subseteq V(G)$ be any vertex set. Then the following assertions are equivalent:*

- (i) *There exists a tree $T \subseteq G$ that contains U and reflects the undominated ends in the closure of U ;*
- (ii) *G has a subgraph H with $U \subseteq V(H)$ normally spanned in H and all whose undominated ends are included in distinct undominated ends of G .*

Assume for a moment that Theorem 13.3 is already verified. If G is any graph and $U \subseteq V(G)$ is normally spanned in G , then statement (ii) of the theorem is satisfied with $H = G$. Hence the implication (ii) \rightarrow (i) yields the following theorem:

Theorem 13.2.2. *Let G be any graph and let $U \subseteq V(G)$ be normally spanned. Then there is a tree $T \subseteq G$ that contains U and reflects the undominated ends in the closure of U .*

Conversely, let us see that Theorem 13.3 can be deduced from Theorem 13.2.2. The implication (i) \rightarrow (ii) of Theorem 13.3 is immediate because any tree as in (i) serves as a subgraph $H \subseteq G$ that is sought in (ii).

For the reverse implication let H and U be as in Theorem 13.3 (ii). Then Theorem 13.2.2 yields a tree $T \subseteq H$ that contains U and reflects the undominated ends of H in the closure of U in H . Let Ψ_H be the set of undominated ends of H in the closure of U in H and let Ψ_G be the set of undominated ends of G in the closure of U in G . Furthermore, let $\phi: \Psi_H \rightarrow \Psi_G$ be the map satisfying $\eta \subseteq \phi(\eta)$ for every end $\eta \in \Psi_H$. By (ii) the map is injective and really has Ψ_G as its target set. Let us show that it is also onto. Given an undominated end ω of G in the closure of U it follows from the star-comb lemma and $U \subseteq T$ that T contains a ray $R \in \omega$ and that the end of T containing R lies in the closure of U in T . Since T is a subgraph of H , the end of H containing R lies in the closure of U in H , and so the map ϕ sends the undominated end of H that contains R to ω , establishing that ϕ is onto. Therefore, $\phi: \Psi_H \rightarrow \Psi_G$ is bijective.

Now consider the natural map $\varphi: \Omega(T) \rightarrow \Omega(H)$ that satisfies $\eta \subseteq \varphi(\eta)$ for every end η of T . Note that $\eta \subseteq (\phi \circ \varphi)(\eta)$ for every end η of T . Since T reflects the undominated ends of H in the closure of U and ϕ is bijective we conclude that the map $\phi \circ \varphi$ witnesses that T reflects the undominated ends of G in the closure of U , as required by (i).

Hence to prove Theorem 13.3 we may equivalently prove Theorem 13.2.2:

Proof of Theorem 13.3. Employ Theorem 13.2.2 as above. □

Furthermore, the case $U = V(G)$ of Theorem 13.2.2 establishes our second main corollary:

Corollary 13.4. *Every graph that has a normal spanning tree does have a spanning tree reflecting its undominated ends.*

Our proof of Theorem 13.2.2 requires some preparation. First, we need the following strengthening of a structural result by Carmesin. Recall from Chapter 11 that a tree-decomposition (T, \mathcal{V}) of a given graph G with finite separators *displays* a set Ψ of ends of G if τ restricts to a bijection $\tau \upharpoonright \Psi: \Psi \rightarrow \Omega(T)$ between Ψ and the end space of T and maps every end that is not contained in Ψ to some node of T , where $\tau: \Omega(G) \rightarrow \Omega(T) \sqcup V(T)$ maps every end of G to the end or node of T which it corresponds to or lives at, respectively.

Theorem 12.2.8. *Every connected graph G has a tree-decomposition with pairwise disjoint finite connected separators that displays the undominated ends of G .*

For our purposes we need to strengthen Carmesin's result further so as to take an arbitrary vertex set U into account. Recall that a rooted tree-decomposition (T, \mathcal{V}) of a graph G *covers* a vertex set $U \subseteq V(G)$ *cofinally* if the set of nodes of T whose parts meet U is cofinal in the tree-order of T .

Theorem 13.2.3. *Let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. Then G has a rooted tree-decomposition with pairwise disjoint finite connected separators that displays the undominated ends of G that lie in the closure of U . Moreover, the tree-decomposition can be chosen so that it covers U cofinally.*

Proof. By Theorem 12.2.8, we find a tree-decomposition (T, \mathcal{V}) of G with pairwise disjoint finite connected separators that displays the undominated ends of G . Consider T rooted in an arbitrary node. Let U' be the set of vertices of T whose parts meet U and let T' be the subtree of T obtained by taking the down-closure of U' in T . Then we let (T, α) be the S_{\aleph_0} -tree corresponding to (T, \mathcal{V}) , so $(T', \alpha \upharpoonright \vec{E}(T'))$ is an S_{\aleph_0} -tree that induces the desired tree-decomposition. \square

Our construction of a tree reflecting the undominated ends in the closure of a given set of vertices will employ a contraction minor H of the underlying graph G . The following notation will help us to translate between the endspace of G and that of H . Consider a contraction minor H of a graph G with fixed finite branch sets. Every direction f of G defines a direction $[f]$ of H by letting $[f](X) := [f(\bigcup X)]$ for every finite vertex set $X \subseteq V(H)$. In fact, it is straightforward to check that every direction of H is defined by a direction of G in this way:

Lemma 13.2.4. *Let H be a contraction minor of a graph G with fixed finite branch sets. Then the map $f \mapsto [f]$ is a bijection between the directions of G and the directions of H . \square*

This one-to-one correspondence then combines with the well-known one-to-one correspondence between the directions and ends of a graph (see Theorem 11.1.7), giving rise to a bijection $\omega \mapsto [\omega]$ between the ends of G and the ends of H . The natural one-to-one correspondence between the two end spaces extends to other aspects of the graphs and their ends:

Lemma 12.2.11. *Let H be a contraction minor of a graph G with fixed finite branch sets, let ω be an end of G and let $U \subseteq V(G)$ be any vertex set. Then ω lies in the closure of U in G if and only if $[\omega]$ lies in the closure of $[U]$ in H ; and ω is dominated in G if and only if $[\omega]$ is dominated in H .*

Lemma 13.2.5. *Let H be a contraction minor of a graph G with fixed branch sets and let $U \subseteq V(G)$ be any vertex set. If U is normally spanned in G , then $[U]$ is normally spanned in H .*

We remark that this is essentially [41, Lemma 7.2 (b)].

Proof. Without loss of generality both graphs G and H are connected. By Theorem 11.2.5, we have that U can be written as a countable union $\bigcup_{n \in \mathbb{N}} U_n$ with every U_n dispersed in G . Then every vertex set $[U_n]$ is dispersed in H , because every comb attached to $[U_n]$ in H would give rise to a comb attached to U_n in G , contradicting that U_n is dispersed in G . Hence $[U] = \bigcup_{n \in \mathbb{N}} [U_n]$ is normally spanned in H by Theorem 11.2.5. \square

We need one more lemma for the proof of Theorem 13.2.2:

Lemma 13.2.6. *Let G be any graph and let $U \subseteq V(G)$ be any vertex set. If (T, \mathcal{V}) is a rooted tree-decomposition of G with pairwise disjoint finite connected separators that displays the undominated ends in $\partial_\Omega U$ and covers U cofinally, then $\partial_\Omega U = \partial_\Omega \hat{U}$ for the superset \hat{U} of U that arises from U by adding all the vertices that lie in the separators of (T, \mathcal{V}) .*

Proof. The inclusion $\partial_\Omega U \subseteq \partial_\Omega \hat{U}$ holds because $U \subseteq \hat{U}$. For the backward inclusion, consider any end ω in the closure of \hat{U} , and assume for a contradiction

that ω does not lie in the closure of U . Then ω lives at a node $t \in T$ because (T, \mathcal{V}) displays the ends in the closure of U . Pick a comb in G attached to \hat{U} and with spine in ω . As ω does not lie in the closure of U we may assume that the comb avoids U . Furthermore, we may assume that every tooth of the comb lies in a separator of (T, \mathcal{V}) associated with an edge of T at and above t . Since the separators of (T, \mathcal{V}) are finite and pairwise disjoint, we may even ensure that no separator contains more than one tooth. As (T, \mathcal{V}) has connected separators and covers U cofinally, we find infinitely many disjoint paths from the comb to U , one starting in each tooth. Then the comb together with these paths witnesses that ω lies in the closure of U , a contradiction. \square

Proof of Theorem 13.2.2. Let G be any graph and let $U \subseteq V(G)$ be normally spanned. Without loss of generality, G is connected. By Theorem 13.2.3 we find a rooted tree-decomposition $(T_{\text{DEC}}, \mathcal{V})$ of G with pairwise disjoint finite connected separators such that $(T_{\text{DEC}}, \mathcal{V})$ displays the undominated ends in the closure of U and covers U cofinally. And by Lemma 13.2.6 we may assume that U contains all the vertices that are contained in the separators of $(T_{\text{DEC}}, \mathcal{V})$.

We construct a tree $T \subseteq G$ displaying the undominated ends in the closure of U as follows. For every separator X of $(T_{\text{DEC}}, \mathcal{V})$ we pick a spanning tree T_X of $G[X]$. As all X are finite and pairwise disjoint, so are the T_X . Next, we choose for every part V_t of $(T_{\text{DEC}}, \mathcal{V})$ a rayless tree T_t in $G[V_t]$ containing $U_t := V_t \cap U$ and extending all the trees T_X for which X is a separator corresponding to some edge incident with t , as follows. Given V_t , we first consider the contraction minor H_t of $G[V_t]$ with fixed branch sets that is obtained from $G[V_t]$ by contracting each $G[X]$ with X a separator induced by an edge of T_{DEC} at t to a single *dummy* vertex named X . As U is normally spanned in G it follows by Lemma 13.2.5 that $[U]_H$ is normally spanned in the contraction minor H obtained from G by contracting every $G[X]$ for every separator. It follows that the vertex sets $[U_t]_{H_t}$ are normally spanned in $H_t \subseteq H$. Furthermore, since $(T_{\text{DEC}}, \mathcal{V})$ has disjoint finite connected separators and displays the undominated ends of G in the closure of U , every end of $G[V_t]$ in the closure of U_t in the graph $G[V_t]$ is dominated in $G[V_t]$. Thus, by Lemma 12.2.11 every end of H_t in the closure of $[U_t]$ is dominated in H_t . Hence we may apply Theorem 13.1 to H_t and $[U_t]$ to obtain a rayless tree \tilde{T}_t in H_t containing $[U_t]$. Then by expanding each dummy vertex X of \tilde{T}_t to T_X we obtain

a rayless tree T_t in $G[V_t]$ that contains U_t and extends all these T_X .

Let T be spanned by the down-closure of U in the tree $\bigcup_{t \in T_{\text{DEC}}} T_t$ with regard to an arbitrary root. We claim that T contains U and reflects the undominated ends in the closure of U . Clearly, T is a tree in G that contains U even cofinally. By the star-comb lemma, every tree in G containing U contains for each undominated end in the closure of U a ray from that end. In particular, T contains a ray from every undominated end in the closure of U .

Next, the tree T contains at most one ray starting in the root for every undominated end in the closure of U : Indeed, if T contains two (say) vertex-disjoint rays from the same undominated end ω in the closure of U , then these give rise to a subdivided ladder in T via the trees T_X along any ray of T_{DEC} to which ω corresponds, and the ladder comes with infinitely many cycles, contradicting that T is a tree.

That T contains only rays from ends in the closure of U is a consequence of Lemma 11.1.13 and the fact that T contains U cofinally by construction.

Finally, the tree T contains no ray from dominated ends in the closure of U , for if T contains a ray from such an end, then the vertex set of that ray intersects some part V_t of (T, \mathcal{V}) infinitely often, and then Lemma 11.1.4 applied in the rayless tree T_t to that intersection yields infinitely many cycles in the tree T . \square

Now that we established the proof of Theorem 13.2.2, let us turn to an application.

Theorem 13.5. *Let G be any graph and let $T \subseteq G$ be any spanning tree.*

- (i) *The fundamental cuts of T are all finite if and only if G is finitely separable and T reflects the undominated ends of G .*
- (ii) *If G is finitely separable and connected, then it has a spanning tree all whose fundamental cuts are finite.*

Before we prove Theorem 13.5, we show a corollary for the topological space \tilde{G} (see [2] for definitions regarding \tilde{G}).

Corollary 13.6. *Every connected finitely separable graph G has a spanning tree whose closure in \tilde{G} contains no topological cycle.*

Proof. By Theorem 13.5 (ii) the graph G has a spanning tree all whose fundamental cuts are finite. We claim that the closure of T in \tilde{G} contains no topological cycles. Indeed, suppose for a contradiction that C is a topological cycle in \bar{T} and fix an edge e of T that is contained in C as a topological edge. Let F_e be the fundamental cut of e with respect to T and let us write V_1 and V_2 for the two sides of F_e . Then $C \setminus \mathring{e}$ is a topological arc A between V_1 and V_2 avoiding the interior of the edges in the finite cut F_e . But then A is a connected subset of $|G| \setminus \bigcup \{\mathring{f} \mid f \in F_e\}$ that is divided into the two closed disjoint sets $\overline{G[V_1]}$ and $\overline{G[V_2]}$ (contradiction). \square

Proof of Theorem 13.5. (i) For the forward implication suppose that the fundamental cuts of T are all finite. First let us see that G is finitely separable. For this consider any two distinct vertices $v, w \in V(G)$ and let e be an edge on the unique path between v and w in T . Then the fundamental cut of e with respect to T is finite and separates v from w in G .

Next, let us show that no ray of T is dominated. For this, consider any ray $R \subseteq T$ and any vertex $v \in V(G)$. Let C be the component of $T - v$ that contains a tail of R and let $e \in E(T)$ be the unique edge between C and v . As the fundamental cut of e with respect to T is finite, and as all the paths of any v - $(R - v)$ fan need to pass through this fundamental cut, the vertex v cannot dominate R .

The tree T contains a ray from every undominated end, because, by the star-comb lemma, every spanning tree of G does so. It remains to show that every distinct two ends of T are included in distinct ends of G . For this consider rays $R, R' \subseteq T$ that belong to distinct ends of T . Let e be an edge on a tail of R that does not meet R' . Then the endvertices of the edges in the finite fundamental cut of e form a finite vertex set that separates a tail of R from a tail of R' in G . Hence R and R' belong to distinct ends of G .

For the backward implication suppose that G is finitely separable and that T reflects the undominated ends of G . Consider any fundamental cut F_e of an edge $e \in E(T)$ with respect to T . Write T_1 and T_2 for the two components of $T - e$. Then F_e consists of the T_1 - T_2 edges of G . Suppose for a contradiction that F_e is infinite. Then F_e has infinitely many endvertices in at least one of T_1 and T_2 . Let us write X_i for the set of endvertices that F_e has in T_i for $i = 1, 2$. We consider

two cases and derive contradictions for both of them.

In the first case, some vertex $x \in X_i$ is incident with infinitely many edges of F_e , say for $i = 1$. Then, as G is finitely separable, applying the star-comb lemma in T_2 to the infinitely many endvertices that these edges have in T_2 must yield a comb whose spine is then dominated by x in G , contradicting that T reflects the undominated ends of G .

In the second case, every vertex of G is incident with at most finitely many edges from F_e . Then F_e contains an infinite matching of an infinite subset of $V(T_1)$ and an infinite subset of $V(T_2)$. First, we apply the star-comb lemma in T_1 to the endvertices of this matching. This yields either a star or a comb, and we write U_1 for its attachment set. Then we apply the star-comb lemma in T_2 to those vertices that are matched to U_1 . Since G is finitely separable, we cannot get two stars. Like in the first case, we cannot get one star and one comb. So we must get two combs. But then T contains two rays that are equivalent in G , contradicting that T reflects some set of ends of G .

(ii) Connected finitely separable graphs are normally spanned due to a result of Halin [40] which states: all connected graphs that do not contain a subdivided K^{\aleph_0} as a subgraph are normally spanned. But it is also possible to construct a normal spanning tree in a connected finitely separable graph directly, as follows. Every 2-connected finitely separable graph G is countable, cf. [69] or [2, Lemma 4.4]. Indeed, if G is 2-connected and uncountable, then G contains a vertex v of uncountable degree and $G - v$ is connected. Hence the strong version of the star-comb lemma (Lemma 11.1.5) applied to the neighbourhood $N(v)$ of v in G returns an infinite star attached to $N(v)$ and it follows that G is not finitely separable. Therefore, the blocks of any connected finitely separable graph G are all countable. Now to show that any connected finitely separable graph G is normally spanned, let us root the block graph of G arbitrarily (having in mind that the block graph is a tree). The block that is the root does have a normal spanning tree because it is countable (cf. Corollary 11.2.3), and we fix an arbitrary normal spanning tree. Then we consider the blocks of height one. Each block B of height one intersects the root block in precisely one vertex x , and we fix any normal spanning tree of B that is rooted at x (Jung has shown that prescribing the root x is possible, see Corollary 11.2.3). Proceeding in this fashion we fix for every block of G a normal spanning tree, and the way we choose their roots

ensures that the union of all these normal trees forms a normal spanning tree of G . So G is normally spanned, and hence Theorem 13.2.2 yields a spanning tree that reflects the undominated ends of G . By the backward implication of (i), all the fundamental cuts of this spanning tree are finite. \square

13.3. Duality theorems for undominated combs

In this section we prove our two duality theorems for undominated combs in full generality. The first theorem is phrased in terms of star-decompositions:

Theorem 13.7. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains an undominated comb attached to U ;
- (ii) G has a star-decomposition with finite separators such that U is contained in the central part and all undominated ends of G live in the leaves' parts.

Moreover, we may assume that the separators of the tree-decomposition in (ii) are pairwise disjoint and connected.

Proof. Clearly, at most one of (i) and (ii) can hold.

To establish that at least one of (i) and (ii) holds, we show $\neg(\text{i}) \rightarrow (\text{ii})$. By Theorem 12.2.8 we find a tree-decomposition (T, \mathcal{V}) of G with pairwise disjoint finite connected separators that displays the undominated ends of G . We let $W \subseteq V(T)$ consist of those nodes $t \in T$ whose parts V_t meet U . Then we root T arbitrarily and let T' be the subtree $\lceil W \rceil$ of T . Since U does not have any undominated end of G in its closure, it follows that T' must be rayless. We obtain the star S from T by contracting T' and all of the components of $T - T'$. Then we let (T, α) be the S_{\aleph_0} -tree corresponding to (T, \mathcal{V}) , so $(S, \alpha \upharpoonright \vec{E}(S))$ is an S_{\aleph_0} -tree that induces the desired star-decomposition which even satisfies the 'moreover' part. \square

The central part of the star-decomposition in Theorem 13.7 (ii) induces a subgraph of G that seems to carry the information that there is no undominated comb attached to U . Our second duality theorem for undominated combs confirms this suspicion:

Theorem 13.8. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains an undominated comb attached to U ;
- (ii) G has a connected subgraph that contains U and all whose rays are dominated in it.

Moreover, the subgraph H in (ii) can be chosen so as to reflect the ends in the closure of H .

Proof. To see that at most one of (i) and (ii) holds, consider any connected subgraph $H \subseteq G$ containing U such that every ray of H is dominated in H . We show that H obstructs the existence of an undominated comb in G attached to U . Assume for a contradiction that such a comb exists. Then the undominated end $\omega \in \Omega(G)$ of that comb's spine lies in the closure of U , and so applying the star-comb lemma in H to the attachment set $U' \subseteq U$ of that comb must yield another comb attached to U' . But this latter comb is dominated in H by assumption, and at the same time its spine is equivalent in G to the first comb's spine, contradicting that ω is undominated in G .

To establish that at least one of (i) and (ii) holds, we show $\neg(\text{i}) \rightarrow (\text{ii})$. Let (T, \mathcal{V}) be the star-decomposition from Theorem 13.7 (ii) also satisfying the 'moreover' part of the theorem. We claim that the graph $H = G[V_c]$ that is induced by the central part V_c of (T, \mathcal{V}) is as desired. Clearly, H contains U . And H is connected because the separators of (T, \mathcal{V}) are connected. Now if R is any ray in H , it is dominated in G by some vertex $v \in V_c$. This vertex v also dominates R in H because every infinite v – $(R - v)$ fan in G can be greedily turned into an infinite v – $(R - v)$ fan in H by employing the connectedness of the finite separators of the star-decomposition.

Finally, let us prove that H is as in the 'moreover' part of the theorem, i.e., let us show that H reflects $\partial_\Omega H$. For this let $\varphi: \Omega(H) \rightarrow \Omega(G)$ be the natural map satisfying $\eta \subseteq \varphi(\eta)$. We have to show that φ is injective with $\text{im}(\varphi) = \partial_\Omega H$.

To see that φ is injective, consider any distinct two ends η and η' of H and let $X \subseteq V(H)$ be a finite vertex set separating them in H . Since the separators of (T, \mathcal{V}) are pairwise disjoint and finite, we may assume that X includes all the separators that it meets. We claim that X separates $\varphi(\eta)$ and $\varphi(\eta')$ in G . Indeed,

otherwise some component of $G - X$, namely $C(X, \varphi(\eta)) = C(X, \varphi(\eta'))$, includes rays $R \in \eta$ and $R' \in \eta'$ together with a path connecting them. As R and R' are rays in H , the path has both its endvertices in H . But then this R - R' path can be turned into an R - R' path in $H - X$ by replacing some of its path segments with paths inside the connected separators that it meets (here we use that every separator meeting the path must avoid X).

It remains to verify $\text{im}(\varphi) = \partial_\Omega H$. The forward inclusion is immediate, we show the backward inclusion. Every ray in any end ω of G in the closure of H intersects H infinitely because the separators of the star-decomposition (T, \mathcal{V}) are all finite. Again we can employ the pairwise disjoint finite connected separators of the star-decomposition (T, \mathcal{V}) to turn the ray into a ray in H that intersects the original ray infinitely often. Then the new ray's end in H is included in ω . \square

14. Undominating stars

In Chapter 11 we found structures whose existence is complementary to the existence of a star or a comb attached to a given set U of vertices, and two types of these structures turned out to be relevant for both stars and combs: normal trees and tree-decompositions. A *comb* is the union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R . The last vertices of those paths are the *teeth* of this comb. Given a vertex set U , a *comb attached to U* is a comb with all its teeth in U , and a *star attached to U* is a subdivided infinite star with all its leaves in U . Then the set of teeth is the *attachment set* of the comb, and the set of leaves is the *attachment set* of the star. Given a graph G , a rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T , cf. [20]. For the definition of tree-decompositions see [20].

As stars and combs can interact with each other, this is not the end of the story. For example, a given vertex set U might be connected in a graph G by both a star and a comb, even with infinitely intersecting sets of leaves and teeth. To formalise this, let us say that a subdivided star S *dominates* a comb C if infinitely many of the leaves of S are also teeth of C . A *dominating star* in a graph G then is a subdivided star $S \subseteq G$ that dominates some comb $C \subseteq G$; and a *dominated comb* in G is a comb $C \subseteq G$ that is dominated by some subdivided star $S \subseteq G$. Thus, a star $S \subseteq G$ is undominating in G if it is not dominating in G ; and a comb $C \subseteq G$ is undominated in G if it is not dominated in G .

In Chapter 12 we determined structures whose existence is complementary to the existence of dominating stars or dominated combs. Like for arbitrary stars and combs, our duality theorems for dominating stars and dominated combs are phrased in terms of normal trees and tree-decompositions.

In Chapter 13 we determined structures whose existence is complementary to the existence of undominated combs. Our investigations showed that the types of complementary structures for undominated combs are quite different compared

to those for stars, combs, dominating stars and dominated combs. On the one hand, normal trees are too strong to serve as complementary structures, which is why we considered more general subgraphs instead. Tree-decompositions on the other hand are dynamic enough to allow for duality theorems, even in terms of star-decompositions—which are too strong to serve as complementary structures for stars, combs, dominating stars or dominated combs.

Among all the combinations of stars and combs, there is only one combination that we have yet to consider: undominating stars. Here, in this chapter, we determine structures whose existence is complementary to the existence of undominating stars. The types of complementary structures for undominating stars differ from those for stars, combs, dominating stars and dominated combs—surprisingly in the same way the types of complementary structures for undominated combs differ from them.

To begin, normal trees are too strong to serve as complementary structures for undominating stars: if G is an uncountable complete graph and $U = V(G)$, then G contains no undominating star attached to U but G has no normal spanning tree. However, if G contains no undominating star attached to U and U happens to be contained in a normal tree $T \subseteq G$, then the down-closure of U in T forms a locally finite subtree H . In this situation H witnesses that U is *tough* in G in that only finitely many components meet U whenever finitely many vertices are deleted from G . This property gives a candidate for a subgraph that might serve as a complementary structure, even when U is not contained in a normal tree. Call a graph G *tough* if its vertex set is tough in G , i.e., if deleting finitely many vertices from G always results in only finitely many components. It is well known that the tough graphs are precisely the graphs that are compactified by their ends, cf. [24]. Our first duality theorem for undominating stars is formulated in terms of tough subgraphs:

Theorem 14.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains an *undominating star* attached to U ;
- (ii) there is a *tough subgraph* $H \subseteq G$ that contains U .

As our second duality theorem for undominating stars, we also find star-decompositions that are complementary to undominating stars:

Theorem 14.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains an *undominating star* attached to U ;
- (ii) G has a tame *star-decomposition* such that U is contained in the central part and every critical vertex set of G lives in a leaf's part.

Here, a finite vertex set $X \subseteq V(G)$ is *critical* if infinitely many of the components of $G - X$ have their neighbourhood precisely equal to X . Critical vertex sets were introduced in [48]. As tangle-distinguishing separators, they have a surprising background involving the Stone-Ćech compactification of G , Robertson and Seymour's tangles from their graph-minor series, and Diestel's tangle compactification, cf. [21, 47, 63]. For the definitions of 'tame' and 'live', see Section 14.2. Tame tree-decompositions have finite adhesion sets.

While the wordings of our two duality theorems for undominating stars are similar to those of the duality theorems for undominated combs, their proofs are not. In fact, a whole new strategy is needed to prove these two theorems. The starting point of our strategy will be a very recent generalisation [31] of Robertson and Seymour's tree-of-tangles theorem from their graph-minor series [63].

This chapter is organised as follows. Section 14.1 establishes our duality theorem for undominating stars in terms of end-compactified subgraphs. Section 14.2 proves our duality theorem for undominating stars in terms of star-decompositions. In Section 14.3 we summarise the duality theorems of the first four chapters of this part.

We assume familiarity with the tools and terminology described in Chapter 11. For definitions and basic properties regarding separation systems refer to [22].

14.1. Tough subgraphs

In this section, we prove our duality theorem for undominating stars in terms of tough subgraphs:

Theorem 14.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) G contains an undominating star attached to U ;
- (ii) G has a tough subgraph that contains U .

We remark that the tough graphs are precisely the graphs that are compactified by their ends, see [24].

We prove that (i) and (ii) are complementary by proving that both \neg (i) and (ii) are equivalent to the assertion that U is tough in G . That \neg (i) is equivalent to U being tough in G will be shown in Lemma 14.1.1, and that (ii) is equivalent to U being tough in G will be shown in Theorem 14.1.2. It will be convenient to make this detour because U being tough in G is easier to work with than G not containing an undominating star attached to U .

Lemma 14.1.1. *A set U of vertices of a connected graph G is tough in G if and only if G contains no undominating star attached to U .*

Theorem 14.1.2. *A set U of vertices of a graph G is tough in G if and only if G has a tough subgraph that contains U .*

Proof of Theorem 14.1. Combine Lemma 14.1.1 and Theorem 14.1.2 above. \square

While the proof of Theorem 14.1.2 takes the rest of this section, that of Lemma 14.1.1 is easy and we shall provide it straight away. Recall that a finite set X of vertices of an infinite graph G is *critical* if the collection

$$\mathcal{C}_X^\infty := \{ C \in \mathcal{C}_X \mid N(C) = X \}$$

is infinite, where \mathcal{C}_X is the collection of all components of $G - X$. A critical vertex set X of G lies *in the closure* of M , where M is either a subgraph of G or a set of vertices of G , if infinitely many components in \mathcal{C}_X^∞ meet M .

Proof of Lemma 14.1.1. If U is tough in G then no critical vertex set of G lies in the closure of U . We know by Lemma 11.1.9 that every infinite set of vertices in a connected graph has an end or a critical vertex set in its closure. Therefore, every infinite subset $U' \subseteq U$ has an end of G in its closure and, in particular, there is always a comb in G attached to U' . Thus, every star in G attached to U must be dominating.

Conversely, if U is not tough in G , then there is a finite vertex set $X \subseteq V(G)$ such that some infinitely many components of $G-X$ meet U . Then infinitely many of these components send an edge to the same vertex $x \in X$ by the pigeonhole principle. This allows us to make x the centre of a star S attached to U by taking $x-U$ paths in $G[x+C]$, one for each of the infinitely many components C that meet U and have x in their neighbourhood. Now X obstructs the existence of a comb that has infinitely many teeth that are also leaves of S , and so S must be undominating. \square

Before we turn to the proof of Theorem 14.1.2, we summarise a few elementary properties that are complementary to containing an undominating star attached to a given vertex set U :

Lemma 14.1.3. *Let G be any connected graph, let $U \subseteq V(G)$ be any vertex set and let $(*)$ be the statement that G contains an undominating star attached to U . Then the following assertions are complementary to $(*)$:*

- (i) U is tough in G ;
- (ii) G has no critical vertex set that lies in the closure of U ;
- (iii) U is compactified by the ends of G that lie in the closure of U .

If U is normally spanned in G , then the following assertion is complementary to $()$ as well:*

- (iv) G contains a locally finite normal tree that contains U cofinally.

Proof. By Lemma 14.1.1 we have that (i) is complementary to $(*)$. The assertions (i) and (ii) are equivalent by the pigeonhole principle, and hence (ii) is complementary to $(*)$ as well. Property (iii) is in turn equivalent to (ii) because every graph is compactified by its ends and critical vertex sets in a compactification $|G|_{\Gamma} = G \cup \Omega(G) \cup \text{crit}(G)$ (see [48] for definitions): For (ii) \rightarrow (iii) note that

the closure $\overline{U} = U \cup \partial_\Omega U$ of U in $|G|_\Gamma$ is the desired compactification, and for $\neg(\text{ii}) \rightarrow \neg(\text{iii})$ note that for every critical vertex set X in the closure of U the infinitely many components of $G - X$ meeting U give rise to an open cover of $U \cup \partial_\Omega U$ in $|G|_\Gamma$ that has no finite subcover. That (iv) is complementary to $(*)$ has already been discussed in the introduction. \square

Now we turn to the proof of Theorem 14.1.2. If a graph G has a tough subgraph containing some vertex set U , then clearly U is tough in G . The reverse implication, which states that for every vertex set U that is tough in G the graph G contains a tough subgraph containing U , is harder to show and needs some preparation.

If U is tough in G , then no critical vertex set of G lies in the closure of U , that is, for every critical vertex set X of G only finitely many components in $\check{\mathcal{C}}_X$ meet U . The collection $\mathcal{C}(X)$ of these finitely many components gives rise to a separation $(\check{\mathcal{C}}_X \setminus \mathcal{C}(X), X) = (A_X, B_X)$ that we think of as pointing towards B_X . As $U \subseteq B_X$ for all critical vertex sets X , all the separations (A_X, B_X) point towards the tough vertex set U . Hence we have a candidate for a tough subgraph: the intersection $\bigcap \{G[B_X] \mid X \in \text{crit}(G)\}$. This candidate contains U because U is contained in all $G[B_X]$, but it can happen that our candidate is a non-tough induced $\overline{K^{\aleph_0}} \subseteq G$ with vertex set U , as the following example shows.

For every $n \in \mathbb{N}$ let A_n be some countably infinite set, such that A_n is disjoint from every A_m with $m \neq n$ and also disjoint from \mathbb{N} . Let G be the graph on $\mathbb{N} \cup \bigcup_{n \in \mathbb{N}} A_n$ where every vertex in A_n is joined completely to $\{0, \dots, n\}$. Then the critical vertex sets are precisely the vertex sets of the form $\{0, \dots, n\}$. For every critical vertex set $X = \{0, \dots, n\}$ the collection of components $\check{\mathcal{C}}_X$ consists of the singletons in A_n and the component of $G - X$ that contains $\mathbb{N} \setminus X$. Therefore, if we set $U = \mathbb{N}$, then $G[B_X] = G - A_n$, and our candidate $\bigcap_X G[B_X]$ turns out to be $G[\mathbb{N}] = \overline{K^{\aleph_0}}$.

Although our approach in its naive form fails, this is not the end of it. We will stick to the idea but perform the construction in a more sophisticated way. For this we shall need the following notation and two structural results from [31] for critical vertex sets in graphs, Theorems 14.1.6 and 14.1.7 below. Essentially, these two theorems together will reveal that the separations (A_X, B_X) with X critical in G can be slightly modified to form a tree set.

A *tree set* is a nested separation system that has neither trivial elements nor degenerate elements, cf. [23]. When $(\vec{S}, \leq, *)$ is a tree set, we also call \vec{S} and S tree sets. In our setting, we shall not have to worry about trivial or degenerate separations too much. Indeed, usually our nested sets of separations will consist of separations (A, B) of a graph with neither $A \setminus B$ nor $B \setminus A$ empty, and these sets are known to form *regular* tree sets: tree sets that do not contain small elements.

Let S be any tree-set consisting of finite-order separations of G . A *part* of S is a vertex set of the form $\bigcap \{ B \mid (A, B) \in O \}$ where O is a consistent orientation of S . Thus, if O is any consistent orientation of S , then it defines a part, which in turn induces a subgraph of G . The graph obtained from this subgraph by adding an edge xy whenever x and y are two vertices of the part that lie together in the separator of some separation in O is called the *torso* of O (or of the part, if O is clear from context). Thus, torsos usually will not be subgraphs of G . We need the following standard lemma:

Lemma 14.1.4 ([31, Corollary 2.11]). *Let G be any graph and let $W \subseteq V(G)$ be any connected vertex set. If B is a part of a tree set of separations of G , then $W \cap B$ is connected in the torso of B .*

Given a collection \mathcal{Y} of (in this chapter usually finite) vertex sets of G we say that a vertex set X of G is \mathcal{Y} -*principal* if X meets for every $Y \in \mathcal{Y}$ at most one component of $G - Y$. And we say that \mathcal{Y} is *principal* if all its elements are \mathcal{Y} -principal.

If $X \subseteq V(G)$ meets precisely one component of $G - Y$ for some $Y \subseteq V(G)$, then we denote this component by $C_Y(X)$.

Every critical vertex set of a graph is \mathcal{X} -principal: since every two vertices in a critical vertex set X are linked by infinitely many independent paths (these exist as \mathcal{C}_X is infinite), no two vertices in X are separated by a finite vertex set.

Definition 14.1.5 ([31, Definition 5.9]). Suppose that \mathcal{Y} is a principal collection of vertex sets of a graph G . A function that assigns to every $X \in \mathcal{Y}$ a subset $\mathcal{H}(X) \subseteq \mathcal{C}_X$ is called *admissable* for \mathcal{Y} if for every two $X, Y \in \mathcal{Y}$ that are incomparable as sets we have either $C_X(Y) \notin \mathcal{H}(X)$ or $C_Y(X) \notin \mathcal{H}(Y)$. If additionally $|\mathcal{C}_X \setminus \mathcal{H}(X)| \leq 1$ for all $X \in \mathcal{Y}$, then \mathcal{H} is *strongly* admissable for \mathcal{Y} .

Theorem 14.1.6 ([31, Theorem 5.10]). *For every principal collection of vertex sets of a connected graph there is a strongly admissible function.*

Theorem 14.1.7 ([31, Theorem 5.11]). *Let G be any connected graph, let \mathcal{Y} be any principal collection of vertex sets of G and let \mathcal{K} be any admissible function for \mathcal{Y} . Then for every distinct two $X, Y \in \mathcal{Y}$, after possibly swapping X and Y ,*

either $(\mathcal{K}(X), X) \leq (Y, \mathcal{K}(Y))$ or $(\mathcal{K}(X), X) \leq (C_Y(X), Y) \leq (\mathcal{K}(Y), Y)$.

In particular, if $\emptyset \subsetneq \mathcal{K}(X) \subsetneq \mathcal{C}_X$ for all $X \in \mathcal{Y}$, then the separations $\{X, \mathcal{K}(X)\}$ form a regular tree set for which the separations $(\mathcal{K}(X), X)$ form a consistent orientation.

Suppose now that \mathcal{Y} is a principal collection of vertex sets of a graph G and that \mathcal{K} is an admissible function for \mathcal{Y} satisfying $\emptyset \subsetneq \mathcal{K}(X) \subsetneq \mathcal{C}_X$ for all $X \in \mathcal{Y}$. If T is the regular tree set $\{\{X, \mathcal{K}(X)\} \mid X \in \mathcal{Y}\}$ provided by Theorem 14.1.7, then we call T a *principal tree set* of G . By a slight abuse of notation, we also call the triple $(T, \mathcal{Y}, \mathcal{K})$ a *principal tree set*. In this context, we write $O_{\mathcal{K}}$ for the consistent orientation $\{(\mathcal{K}(X), X) \mid X \in \mathcal{Y}\}$ of T .

Corollary 14.1.8. *Let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. If U is tough in G , then there is a principal tree set $(T, \text{crit}(G), \mathcal{K})$ of G satisfying the following two conditions:*

- (i) *no element of $\mathcal{K}(X)$ meets U for any critical vertex set X ;*
- (ii) *$\mathcal{K}(X)$ is a cofinite subset of $\check{\mathcal{C}}_X$ for every critical vertex set X .*

Proof. As U is tough in G , for every critical vertex set X of G only finitely many components in $\check{\mathcal{C}}_X$ meet U ; we write \mathcal{F}_X for this finite collection. Theorem 14.1.6 yields a strongly admissible function \mathcal{K} for the collection $\text{crit}(G)$ of all the critical vertex sets of G . We alter this function by removing \mathcal{F}_X from $\mathcal{K}(X)$ for all X . Then \mathcal{K} is still admissible for $\text{crit}(G)$, and $\mathcal{K}(X)$ is a cofinite subcollection of $\check{\mathcal{C}}_X \setminus \mathcal{F}_X$ for all X . Now Theorem 14.1.7 says that the separations $\{X, \mathcal{K}(X)\}$ with X critical form a tree set, and that the oriented separations $(\mathcal{K}(X), X)$ form a consistent orientation of this tree set. \square

Proof of Theorem 14.1.2. If H is a tough subgraph of G covering U , then U is tough in H ; in particular, U is tough in G . Conversely, we need to show that for

every vertex set $U \subseteq V(G)$ that is tough in G there is a tough subgraph of G containing U . By Corollary 14.1.8 we find a principal tree set $(T, \text{crit}(G), \mathcal{K})$ so that no element of $\mathcal{K}(X)$ meets U for any critical vertex set X . We write B for the part of T that is defined by $O_{\mathcal{K}}$. Note that U is included in B .

First we claim that the torso of the part B is tough. To see this, consider any finite vertex set $X \subseteq B$. Only finitely many components of $G - X$ meet B : indeed, if infinitely many components of $G - X$ meet B , then by the pigeonhole principle we deduce that a subset X' of X is critical in G with infinitely many components in $\mathcal{C}_{X'}$ meeting B . But then $\bigcup \mathcal{K}(X')$ must meet B , contradicting that B is the part of T that is defined by $O_{\mathcal{K}} = \{(\mathcal{K}(X), X) \mid X \in \text{crit}(G)\}$. Thus $G - X$ has only finitely many components meeting B . By Lemma 14.1.4 each of these components induces a component of the torso minus X , and so deleting X from the torso results in at most finitely many components.

The tough torso of the part B , however, usually is not a subgraph of G . And the part B usually will not induce a tough subgraph of G . That is why as our next step, we construct a subgraph H of G that imitates the torso of B to inherit its toughness. More precisely, we obtain H from $G[B]$ by adding a subgraph L of G that has the following three properties:

- (L1) Every vertex of $L - B$ has finite degree in L .
- (L2) For every finite $X \subseteq B$ only finitely many components of $L - X$ avoid B .
- (L3) If x and y are distinct vertices in B that lie together in a critical vertex set of G , then L contains a B -path between x and y .

Before we begin the construction of L , let us verify that any L satisfying these three properties really gives rise to a tough subgraph $H = G[B] \cup L$. For this, consider any finite vertex set $X \subseteq V(H)$. By (L1) every vertex of $H - B$ has finite degree in H , and hence deleting it produces only finitely many new components. Therefore we may assume that X is included in B entirely. Every component of $H - X$ avoiding B is a component of $L - X$ avoiding B , and there are only finitely many such components by (L2). Hence it remains to show that there are only finitely many components of $H - X$ that meet B . We already know that the torso of B is tough, so deleting X from it results in at most finitely many components. Then property (L3) ensures that each of these finitely many components has its

vertex set included in a component of $H - X$. And hence there can only be finitely many components of $H - X$ that meet B .

Finally, we construct a subgraph $L \subseteq G$ satisfying the three properties (L1), (L2) and (L3). Choose $(\{x_\alpha, y_\alpha\})_{\alpha < \kappa}$ to be a transfinite enumeration of the collection of all unordered pairs $\{x, y\}$ where x and y are distinct vertices in B that lie together in a critical vertex set of G . Then we recursively construct L as a union $L = \bigcup_{\alpha < \kappa} P_\alpha$ where at step α we choose P_α from among all B -paths P in G between x_α and y_α so as to minimize the number $|E(P) \setminus E(\bigcup_{\xi < \alpha} P_\xi)|$ of new edges. (There is a B -path in G between x_α and y_α since x_α and y_α lie together in some critical vertex set X of G and $\mathcal{K}(X) \subseteq \mathcal{C}_X^\checkmark$ is non-empty.)

We verify that our construction yields an L satisfying (L1), (L2) and (L3).

(L1). For this, fix any vertex $\ell \in L - B$. It suffices to show that the edges of L at ℓ simultaneously extend to an ℓ - B fan in L . To see that this really suffices, use that ℓ is not contained in B to find some critical vertex set X of G with $\ell \in \bigcup \mathcal{K}(X)$. Then the ℓ - B fan at ℓ extending the edges of L at ℓ must have all its ℓ - B paths pass through the finite X , and so there can be only finitely many such paths, meaning that ℓ has finite degree in L .

Now to find the ℓ - B fan we proceed as follows. For every edge e of L at ℓ we write $\alpha(e)$ for the minimal ordinal α with $e \in E(P_\alpha)$. Then we write P_e for $P_{\alpha(e)}$, and we write Q_e for the ℓ - B subpath of P_e containing e . The paths Q_e form an ℓ - B fan, as we verify now. For this, we show that, if $e \neq e'$ are two distinct edges of L at ℓ , then Q_e and $Q_{e'}$ meet precisely in ℓ . Let e and e' be given. We abbreviate $\alpha(e) = \alpha$ and $\alpha(e') = \alpha'$. If $\alpha = \alpha'$ then $Q_e \cup Q_{e'} = P_\alpha$ and we are done. Otherwise $\alpha < \alpha'$, say. Then we assume for a contradiction that $\overset{\circ}{\ell}Q_{e'}$ does meet $\overset{\circ}{\ell}Q_e$. Without loss of generality we may assume that $Q_{e'}$ starts in ℓ and ends in $y_{\alpha'}$. We let t be the last vertex of $Q_{e'}$ in $\overset{\circ}{\ell}Q_e$. But then the graph $x_{\alpha'}P_{e'}\ell \cup \overset{\circ}{\ell}Q_e t P_{e'} y_{\alpha'}$ is connected and meets B precisely in the two vertices $x_{\alpha'}$ and $y_{\alpha'}$. Consequently, it contains a B -path P between $x_{\alpha'}$ and $y_{\alpha'}$. But then P avoids the edge e' , so the inclusion $E(P) \setminus E(\bigcup_{\xi < \alpha'} P_\xi) \subseteq E(P_{e'}) \setminus E(\bigcup_{\xi < \alpha'} P_\xi)$ must be proper. Therefore, P contradicts the choice of $P_{\alpha'}$ as desired.

(L2). For this, fix any finite vertex set $X \subseteq B$. Let \mathcal{C} be the set consisting of all the components of $L - X$ that avoid B . And let F consist of all the edges inside components from \mathcal{C} and all the edges of L between components from \mathcal{C} and X , i.e., $F = E(\bigcup \mathcal{C}) \cup E_L(\bigcup \mathcal{C}, X)$. As every component from \mathcal{C} meets some

edge from F it suffices to show that F is finite, a fact that we verify as follows. Every edge in F lies on a path P_α , and since P_α is a B -path between x_α and y_α we deduce $\{x_\alpha, y_\alpha\} \in [X]^2$. Thus the finite edge sets of the paths P_α with $\{x_\alpha, y_\alpha\} \in [X]^2$ cover F . Since X is finite so is $[X]^2$, and hence there are only finitely many such paths, meaning that F is finite.

(L3). This property holds by construction.

As (L1), (L2) and (L3) are now verified we conclude that L is as desired, which completes the proof of our first main result. \square

14.2. Star-decompositions

In this section we prove our second main result, a duality theorem for undominating stars in terms of star-decompositions, Theorem 14.2 below.

Before we state the theorem, let us recall the following definitions from the first chapter of this part, Chapter 11. A finite-order separation $\{X, \mathcal{C}\}$ of a graph G is *tame* if for no $Y \subseteq X$ both \mathcal{C} and $\mathcal{C}_X \setminus \mathcal{C}$ contain infinitely many components whose neighbourhoods are precisely equal to Y . The tame separations of G are precisely the finite-order separations of G that respect the critical vertex sets:

Lemma 14.2.1. *11.2.15 A finite-order separation $\{A, B\}$ of a graph G is tame if and only if every critical vertex set X of G together with all but finitely many components from \mathcal{C}_X is contained in one side of $\{A, B\}$.*

An S_{\aleph_0} -tree (T, α) is *tame* if all the separations in the image of α are tame. As a consequence of Lemma 14.2.1, if X is a critical vertex set of G and (T, α) is a tame S_{\aleph_0} -tree, then X induces a consistent orientation of the image of α by orienting every tame finite-order separation $\{A, B\}$ towards the side that contains X and all but finitely many of the components from \mathcal{C}_X . This consistent orientation, via α , also induces a consistent orientation of $\vec{E}(T)$. Then, just like for ends, the critical vertex set X either *lives* at a unique node $t \in T$ or *corresponds* to a unique end of T . As usual, these definitions for S_{\aleph_0} -trees carry over to tree-decompositions.

Theorem 14.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be infinite. Then the following assertions are complementary:*

- (i) *G contains an undominating star attached to U ;*
- (ii) *G has a tame star-decomposition such that U is contained in the central part and every critical vertex set of G lives in a leaf's part.*

The proof of this theorem is organised as follows. First, we state without proof a technical theorem, Theorem 14.2.2 below, and then show how it implies our main result, Theorem 14.2 above. In a last step we prepare and provide the proof of the technical theorem.

Note that the part of a star σ of separations of a graph G is $\bigcap \{ B \mid (A, B) \in \sigma \}$. Given two oriented separations \vec{s}_1, \vec{s}_2 of G we write $\vec{s}_1 \lesssim \vec{s}_2$ if either $\vec{s}_1 \leq \vec{s}_2$ or there is a component $C \in \mathcal{C}$ for $(\mathcal{C}, X) = \vec{s}_1$ such that $(\mathcal{C} \setminus \{C\}, X) \leq \vec{s}_2$. Here is the technical theorem:

Theorem 14.2.2. *Let G be any graph, and let $(T, \mathcal{Y}, \mathcal{X})$ be any principal tree set so that $O_{\mathcal{X}}$ defines an infinite part. Then G admits a star σ of finite-order separations such that the following two conditions hold:*

- (i) *the part defined by $O_{\mathcal{X}}$ is included in the part of σ ;*
- (ii) *for every $\vec{s} \in O_{\mathcal{X}}$ there is some $\vec{r} \in \sigma$ with $\vec{s} \lesssim \vec{r}$.*

The technical theorem implies our second main result, Theorem 14.2:

Proof of Theorem 14.2. First, we show that at most one of (i) and (ii) holds. By Lemma 14.1.3 we know that if G contains an undominating star attached to U , then G has a critical vertex X that lies in the closure of U . But then X lives in a leaf's part of the star-decomposition provided by (ii), and it follows that this part does contain infinitely many vertices from U , contradicting that U is contained in the central part and that the separations of the star-decomposition are finite.

Now, to show that at least one of (i) and (ii) holds, we show $\neg(\text{i}) \rightarrow (\text{ii})$. By Lemma 14.1.3 we know that U is tough in G . Then, by Corollary 14.1.8, we find a principal tree set $(T, \text{crit}(G), \mathcal{X})$ such that, for every critical vertex set X , no element of $\mathcal{X}(X)$ meets U and the inclusion $\mathcal{X}(X) \subseteq \check{\mathcal{C}}_X$ is cofinite. We claim that the star provided by Theorem 14.2.2 gives a star-decomposition of G meeting the requirements of (ii), a fact that can be verified as follows: First, the separations

of the form $(\mathcal{H}(X), X)$ with X critical and $\mathcal{H}(X)$ a cofinite subset of $\check{\mathcal{C}}_X$ are tame and thus our star-decomposition is tame. Next, by Theorem 14.2.2 (i), we have that U is contained in the central part of the star-decomposition. Finally, by Theorem 14.2.2 (ii), every critical vertex set of G lives in a leaf's part. \square

Next, we prepare the proof of our technical theorem, Theorem 14.2.2. First, we will need the following theorem by Kneip. A chain \mathcal{C} in a given poset is said to have *order-type* α for an ordinal α if \mathcal{C} with the induced linear order is order-isomorphic to α . The chain \mathcal{C} is then said to be an α -*chain*.

Theorem 14.2.3 ([44, Theorem 1]). *A tree set is isomorphic to the edge tree set of a tree if and only if it is regular and contains no $(\omega + 1)$ -chain.*

Besides this theorem, we will need the following concept of a corridor from [31]. Suppose that $(\vec{T}, \leq, *)$ is a tree set, and that O is a consistent orientation of \vec{T} . A *corridor* of O is an equivalence class of separations in O , where two separations $\vec{s}_1, \vec{s}_2 \in O$ are considered equivalent if there is $\vec{r} \in O$ with $\vec{s}_1, \vec{s}_2 \leq \vec{r}$, cf. [31, Lemma 7.1 and Definition 7.2]. As corridors are consistent partial orientations of tree sets on the one hand, and directed posets on the other hand, they come with a number of useful properties.

The *supremum* $\sup L$ of a set L of oriented separations of a graph is the oriented separation (A, B) with $A = \bigcup \{C \mid (C, D) \in L\}$ and $B = \bigcap \{D \mid (C, D) \in L\}$.

Lemma 14.2.4. *Let T be any regular tree set of separations of any graph G , let O be any consistent partial orientation of T and let γ be any corridor of O . Then the supremum of γ is nested with \vec{T} .*

Proof. Consider any unoriented separation $r \in T$. If there is a separation $\vec{s} \in \gamma$ such that r has an orientation \vec{r} with $\vec{r} \leq \vec{s}$, then $\vec{r} \leq \vec{s} \leq \sup \gamma$ as desired. As T is nested, r has for every separation $\vec{s} \in \gamma$ an orientation $\vec{r}(\vec{s})$ such that either $\vec{r}(\vec{s}) \leq \vec{s}$ or $\vec{s} \leq \vec{r}(\vec{s})$. By our first observation, we may assume that $\vec{s} \leq \vec{r}(\vec{s})$ for all $\vec{s} \in \gamma$. It suffices to show that $\vec{r}(\vec{s}_1) = \vec{r}(\vec{s}_2)$ for all $\vec{s}_1, \vec{s}_2 \in \gamma$, since then r has one orientation that lies above all elements of γ and, in particular, above the supremum of γ . Given $\vec{s}_1, \vec{s}_2 \in \gamma$ consider any $\vec{s}_3 \in \gamma$ with $\vec{s}_1, \vec{s}_2 \leq \vec{s}_3$. Then $\vec{s}_1, \vec{s}_2 \leq \vec{s}_3 \leq \vec{r}(\vec{s}_3)$. As T is regular, $\vec{r}(\vec{s}_3) = \vec{r}(\vec{s}_1) = \vec{r}(\vec{s}_2)$ follows. \square

Lemma 14.2.5. *Let T be any tree set of separations of any graph G and let O be any consistent orientation of T . Then the suprema of the corridors of O form a star.*

Proof. We have to show that for every two distinct corridors γ and δ of O the supremum (A, B) of γ and the supremum (C, D) of δ satisfy $(A, B) \leq (D, C)$. Let us write $\gamma = \{(A_i, B_i) \mid i \in I\}$ and $\delta = \{(C_j, D_j) \mid j \in J\}$. As γ is distinct from δ we have $(A_i, B_i) \leq (D_j, C_j)$ for all $i \in I$ and $j \in J$. Hence $(A, B) = (\bigcup_i A_i, \bigcap_i B_i) \leq (\bigcap_j D_j, \bigcup_j C_j) = (D, C)$. \square

Lemma 14.2.6. *Suppose that T is any tree set of separations of any graph G , that O is any consistent orientation of T , and that γ is any corridor of O . Then every finite subset of the separator of the supremum of γ is contained in the separator of some separation in γ .*

In particular, if the order of the separations in γ is bounded by some natural number n , then the supremum of γ has order at most n .

Proof. Let us write (A, B) for the supremum of γ and let Y be any finite subset of its separator $X := A \cap B$. For every vertex $y \in Y \subseteq A$ there is separation $(C_y, D_y) \in \gamma$ with $y \in C_y$. Since γ is a corridor we find a separation $(C, D) \in \gamma$ lying above all (C_y, D_y) . Then $Y \subseteq C$ as C includes all C_y , and $Y \subseteq D$ because $(C, D) \leq (A, B)$ gives $Y \subseteq X \subseteq B \subseteq D$. \square

Before we start with the proof of Theorem 14.2.2 we need two final ingredients: induced separation systems and parliaments. If $\vec{S} = (\vec{S}, \leq, *)$ is a separation system and $O \subseteq \vec{S}$ is any subset (usually a partial orientation of S), then O induces a separation system $O \cup O^*$ that is a subsystem of \vec{S} with the partial ordering and involution induced by \leq and $*$. We denote this subsystem by $\vec{S}[O]$.

Next, we define parliaments. Suppose that G is any graph, that $\vec{T} = (\vec{T}, \leq, *)$ is any regular tree set of finite-order separations of G , and that O is any consistent orientation of \vec{T} . For every number $n \in \mathbb{N}$ let $O_{\leq n}$ be the subset of O formed by the oriented separations in O whose separators have size at most n . Then, by Lemma 14.2.6, every corridor of $O_{\leq n}$ has a supremum of order at most n , and these suprema form a star for fixed n (cf. Lemma 14.2.5) which we denote by $\pi_n(O)$. The *parliament* of O , denoted by $\pi(O)$, is the union $\bigcup_{n \in \mathbb{N}} \pi_n(O)$. Notably, the parliament of O is a cofinal subset of $O \cup \pi(O)$. The parliament of O induces a

separation system $\vec{S}_{\aleph_0}[\pi(O)]$ that is a subsystem of \vec{S}_{\aleph_0} whose separations are all nested with each other. Furthermore, $\vec{S}_{\aleph_0}[\pi(O)]$ and \vec{T} are nested with each other in \vec{S}_{\aleph_0} by Lemma 14.2.4. Also, the parliament of O is a consistent orientation of $\vec{S}_{\aleph_0}[\pi(O)]$ where it defines the same part as O does for \vec{T} .

As one might expect, the inverses of corridors of parliaments have no ω -chains:

Lemma 14.2.7. *Let G be any graph, let \vec{T} be any regular tree set of finite-order separations of G , and let O be any consistent orientation of \vec{T} . Then the inverse γ^* of any corridor γ of \vec{T} has no ω -chain.*

Proof. Suppose for a contradiction that there is a sequence $\vec{s}_0 < \vec{s}_1 < \dots$ of separations $\vec{s}_n \in \gamma^*$. Note that $\vec{s} < \vec{r}$ with $\vec{s} \in \pi_m(O)$ and $\vec{r} \in \pi_n(O)$ implies $m < n$. Hence the function $g: \omega \rightarrow \omega$ assigning to each $n < \omega$ the least $k < \omega$ with $\vec{s}_n \in \pi_k(O)$ is strictly decreasing in that $g(m) > g(n)$ for all $m < n$, contradicting that there are only finitely many natural numbers $< g(0)$. \square

The corridors of a parliament usually stem from S_{\aleph_0} -trees:

Theorem 14.2.8. *Let G be any graph, let \vec{T} be any regular tree set of finite-order separations of G , and let O be any consistent orientation of \vec{T} such that $\vec{S}_{\aleph_0}[\pi(O)]$ is regular. Then for every corridor γ of the parliament of O the corresponding regular tree set $\vec{S}_{\aleph_0}[\gamma]$ is isomorphic to the edge tree set of a tree.*

Proof. Let γ be any corridor of the parliament of O . By Theorem 14.2.3, it suffices to show that $\vec{S}_{\aleph_0}[\gamma]$ has no $(\omega + 1)$ -chain. For this, suppose for a contradiction that $\vec{s}_0 < \vec{s}_1 < \dots < \vec{s}_\omega$ is an $(\omega + 1)$ -chain in $\vec{S}_{\aleph_0}[\gamma]$.

If \vec{s}_ω lies in γ , then so do all the other \vec{s}_n as γ is consistent. Note that $\vec{s} < \vec{r}$ with $\vec{s} \in \pi_m(O)$ and $\vec{r} \in \pi_n(O)$ implies $m < n$. Hence the function $f: \omega + 1 \rightarrow \omega$ assigning to each $\alpha \leq \omega$ the least $n < \omega$ with $\vec{s}_\alpha \in \pi_n(O)$ is strictly increasing in that $f(\alpha) < f(\beta)$ for all $\alpha < \beta$, contradicting $f(\omega) < \omega$.

Otherwise \vec{s}_ω lies in γ^* . If there is an $N < \omega$ with $\vec{s}_n \in \gamma^*$ for all $n \geq N$, without loss of generality $N = 0$, then γ^* has an ω -chain contradicting Lemma 14.2.7.

Therefore, we may assume that $\vec{s}_n \in \gamma$ for infinitely many $n < \omega$. Since γ is consistent, $\vec{s}_n \in \gamma$ for all $n < \omega$ follows. Using that γ is a corridor we find a separation $\vec{r} \in \gamma$ with $\vec{s}_\omega \leq \vec{r}$ and $\vec{s}_0 \leq \vec{r}$. For every $n < \omega$, either $\vec{s}_n \leq \vec{r}$ or $\vec{s}_n \leq \vec{r}$ or $\vec{s}_n \leq \vec{r}$ or $\vec{s}_n \leq \vec{r}$. We cannot have $\vec{s}_n \leq \vec{r}$ for any n , since this would imply $\vec{s}_0 < \vec{s}_n \leq \vec{r} \leq \vec{s}_0$ contradicting that $\vec{S}_{\aleph_0}[\pi(O)]$ is regular. We cannot have

$\vec{s}_n \leq \vec{r}$ for any n because γ is consistent. And we cannot have $\vec{s}_n \leq \vec{r}$, because then $\vec{s}_\omega \leq \vec{r} \leq \vec{s}_n < \vec{s}_\omega$ contradicts that $\vec{S}_{\aleph_0}[\pi(O)]$ is regular. Hence $\vec{s}_n \leq \vec{r}$ for all n . As γ contains no $(\omega + 1)$ -chains by the first case, there must be an $\ell < \omega$ with $\vec{s}_\ell = \vec{r}$. But this then contradicts $\vec{r} = \vec{s}_\ell < \vec{s}_{\ell+1} \leq \vec{r}$, completing the proof that $\vec{S}_{\aleph_0}[\gamma]$ has no $(\omega + 1)$ -chains. \square

Finally, we prove our technical theorem:

Proof of Theorem 14.2.2. Let $(T_{\mathcal{X}}, \mathcal{Y}, \mathcal{X})$ be any principal tree set of a connected graph G so that $O_{\mathcal{X}}$ defines an infinite part. We let O be the parliament of $O_{\mathcal{X}}$. Then the tree set $\vec{S}_{\aleph_0}[O]$ is regular: for every $n \in \mathbb{N}$ and every $(A, B) \in \pi_n(O_{\mathcal{X}}) \subseteq O$ we have that $A \setminus B$ contains the non-empty vertex set of the graph $\bigcup \mathcal{X}(X)$ for some $X \in \mathcal{Y}$, and $B \setminus A$ contains all but at most $|A \cap B| \leq n$ of the infinitely many vertices of the infinite part defined by O . Therefore, by Theorem 14.2.8 we find for every corridor γ of O an S_{\aleph_0} -tree $(T_\gamma, \alpha_\gamma)$ such that α_γ is an isomorphism between the edge tree set $\vec{E}(T_\gamma)$ of T_γ and $\vec{S}_{\aleph_0}[\gamma]$.

In a first step, we will use the S_{\aleph_0} -trees $(T_\gamma, \alpha_\gamma)$ to define stars σ_γ , one for every corridor γ of O , such that their union $\sigma = \bigcup_\gamma \sigma_\gamma$ is a candidate for the star that we seek. Then, in a second step, we will verify that σ is indeed as desired, completing the proof.

First step. We define stars σ_γ , one for each corridor γ of O , such that their union $\sigma := \bigcup_\gamma \sigma_\gamma$ is a candidate for the star that we seek. For this, consider any corridor γ of O . Then γ , as it orients the image of α_γ consistently, defines either a node or an end of T_γ (see Chapter 11).

If γ defines a node t of T_γ , then t has precisely one neighbour in T_γ . Indeed, γ is the down-closure in $\vec{S}_{\aleph_0}[\gamma]$ of the star $\alpha_\gamma(\vec{F}_t)$ where

$$\vec{F}_t = \{ (e, s, t) \in \vec{E}(T_\gamma) \mid e = st \in T_\gamma \}.$$

Note that all separations in $\alpha_\gamma(\vec{F}_t)$ are maximal in γ . Hence, if t has two distinct neighbours k_1 and k_2 in T_γ , then γ contains a separation \vec{r} that lies above both $\alpha_\gamma(k_1, t)$ and $\alpha_\gamma(k_2, t)$, contradicting the maximality in the corridor γ of at least one of these two separations (here we also use that $\alpha_\gamma(k_1, t)$ and $\alpha_\gamma(k_2, t)$ are distinct for distinct neighbours k_1 and k_2 of t because α_γ is injective). Therefore, t is a leaf of T_γ . Call its neighbour k . Then $\alpha_\gamma(k, t)$ is the maximal element of the corridor γ , and we let $\sigma_\gamma := \{ \alpha_\gamma(k, t) \}$.

Otherwise γ defines an end of T_γ from which we pick a ray $R_\gamma = v_\gamma^0 v_\gamma^1 \dots$ all whose edges are oriented forward by γ in that $\vec{s}_\gamma^n := \alpha_\gamma(v_\gamma^n, v_\gamma^{n+1})$ lies in γ for all $n \in \mathbb{N}$. Then we let

$$\sigma_\gamma := \{ \vec{s}_\gamma^0 \} \cup \{ \vec{s}_\gamma^n \wedge \overleftarrow{s}_\gamma^{n-1} : n \geq 1 \}. \quad (14.2.1)$$

(See Figure 14.2.1.)




Figure 14.2.1.: The light grey area depicts $B \setminus A$, the grey area depicts $A \setminus B$ and the dark grey area depicts $A \cap B$ of the separation $(A, B) := \vec{s}_\gamma^n \wedge \overleftarrow{s}_\gamma^{n-1}$ from the proof of Theorem 14.2.2.

Let us check that σ_γ really is a star. On the one hand, it follows from $\vec{s}_\gamma^0 \leq \vec{s}_\gamma^{n-1}$ that $\vec{s}_\gamma^0 \leq \overleftarrow{s}_\gamma^n \vee \vec{s}_\gamma^{n-1} = (\vec{s}_\gamma^n \wedge \overleftarrow{s}_\gamma^{n-1})^*$ for all $n \geq 1$. And on the other hand, for $1 \leq n < m$, we infer from $\vec{s}_\gamma^{n-1} \leq \vec{s}_\gamma^n \leq \vec{s}_\gamma^{m-1} \leq \vec{s}_\gamma^m$ that

$$\vec{s}_\gamma^m \wedge \overleftarrow{s}_\gamma^{m-1} \leq \overleftarrow{s}_\gamma^{m-1} \leq \overleftarrow{s}_\gamma^n \leq \overleftarrow{s}_\gamma^n \vee \vec{s}_\gamma^{n-1} = (\vec{s}_\gamma^n \wedge \overleftarrow{s}_\gamma^{n-1})^*.$$

Since all \vec{s}_γ^n have finite order, so do the infima of which σ_γ is composed. This technique of turning a ray into a star of separations has been introduced by Carmesin [15] in his ‘Proof that Lemma 6.8 implies Lemma 6.7’.

Second step. We prove that σ is as desired. First, we show condition (i), which states that the part defined by $O_{\mathcal{H}}$ is included in the part of σ . For every separation $\vec{s} \in \sigma$ there is some separation $\vec{r} \in O$ satisfying $\vec{s} \leq \vec{r}$. Hence the part

of σ includes the part of O , which in turn includes the part of $O_{\mathcal{H}}$ because O is the parliament of $O_{\mathcal{H}}$.

It remains to verify condition (ii), which states that for every $(\mathcal{H}(X), X) \in O_{\mathcal{H}}$ there is some $\vec{s} \in \sigma$ with $(\mathcal{H}(X), X) \lesssim \vec{s}$. For this, let any vertex set $X \in \mathcal{Y}$ be given. As O is cofinal in $O_{\mathcal{H}} \cup O$, there is a separation $\vec{s}_X \in O$ above $(\mathcal{H}(X), X)$. Let γ be the corridor of O containing \vec{s}_X . We check the following two cases.

In the first case, σ_γ is a singleton, formed by the maximal element \vec{s} of γ , giving

$$(\mathcal{H}(X), X) \leq \vec{s}_X \leq \vec{s} \in \sigma.$$

In the second case, σ_γ is of the form (14.2.1). Then, as O is nested with $T_{\mathcal{H}}$, the separation $(\mathcal{H}(X), X)$ induces a consistent orientation of the image of α_γ , as follows. The orientation consists of all $\vec{r} \in \vec{S}_{\aleph_0}[\gamma]$ that satisfy either $\vec{r} \leq (\mathcal{H}(X), X)$ or $(\mathcal{H}(X), X) < \vec{r}$. Now this consistent orientation defines either a node or an end of T_γ . Since $\vec{s}_X \in \gamma$ lies above $(\mathcal{H}(X), X)$ and since γ^* contains no ω -chains by Lemma 14.2.7, it must be a node t of T_γ . Let $P = t_0 \dots t_k$ be the t - R_γ path in T_γ and let $n \in \mathbb{N}$ be the number with $v_\gamma^n = t_k$, see Figure 14.2.2 (the ray $R_\gamma = v_\gamma^0 v_\gamma^1 \dots$ was defined right above (14.2.1)).




Figure 14.2.2.: The orientation of the image $\vec{S}_{\aleph_0}[\gamma]$ of α_γ and the path P in the second step of the proof of Theorem 14.2.2.

We claim that we may assume $n \neq 0$. For this, it suffices to show that we may assume that \vec{s}_γ^0 lies in the orientation that defines t . So let us consider the

case that $\overleftarrow{s}_\gamma^0$ instead of $\overrightarrow{s}_\gamma^0$ lies in the orientation that defines t . In this case we have either $\overleftarrow{s}_\gamma^0 \leq (\mathcal{H}(X), X)$ or $(\mathcal{H}(X), X) < \overleftarrow{s}_\gamma^0$. But actually, we cannot have $\overleftarrow{s}_\gamma^0 \leq (\mathcal{H}(X), X)$ because otherwise $(\mathcal{H}(X), X) \leq \overrightarrow{s}_X$ would imply that $\overleftarrow{s}_\gamma^0 \leq \overrightarrow{s}_X$ meaning that $\overleftarrow{s}_\gamma^0$ and \overrightarrow{s}_X violate the consistency of γ . Therefore, we must have $(\mathcal{H}(X), X) < \overleftarrow{s}_\gamma^0$, and then we are done because $\overrightarrow{s}_\gamma^0$ is an element of σ_γ . Thus, we may assume $n > 0$.

If the path P is non-trivial, i.e., if $t_0 = t$ is distinct from $t_k = v_\gamma^n$, then we consider the separation $\vec{r}_P = \alpha_\gamma(t_{k-1}, t_k) \in \gamma$ associated with the last edge $t_{k-1}t_k$ of P . By the definition of P , the separation \vec{r}_P satisfies either $\vec{r}_P \leq (\mathcal{H}(X), X)$ or $(\mathcal{H}(X), X) < \vec{r}_P$. The former inequality would violate the consistency of γ as $\vec{r}_P \leq (\mathcal{H}(X), X) \leq \overrightarrow{s}_X$ would follow (here we use that $\vec{S}_{\mathfrak{N}_0}[\gamma] \subseteq \vec{S}_{\mathfrak{N}_0}[O]$ is regular to ensure $\vec{r}_P \neq \overrightarrow{s}_X$). Hence $(\mathcal{H}(X), X) < \vec{r}_P$. As t_{k-1} is distinct from v_γ^{n-1} , and both vertices have v_γ^n as a neighbour in T_γ , we obtain the inequalities $\vec{r}_P \leq \overrightarrow{s}_\gamma^n$ and $\vec{r}_P \leq \overleftarrow{s}_\gamma^{n-1}$. Thus,

$$(\mathcal{H}(X), X) \leq \vec{r}_P \leq \overrightarrow{s}_\gamma^n \wedge \overleftarrow{s}_\gamma^{n-1} \in \sigma.$$

Otherwise the path P is trivial, i.e., $t_0 = t_k$ where $t_0 = t$ and $t_k = v_\gamma^n$. By the definition of t we have either $\overrightarrow{s}_\gamma^{n-1} \leq (\mathcal{H}(X), X)$ or $(\mathcal{H}(X), X) < \overrightarrow{s}_\gamma^{n-1}$, and we have either $\overleftarrow{s}_\gamma^n \leq (\mathcal{H}(X), X)$ or $(\mathcal{H}(X), X) < \overleftarrow{s}_\gamma^n$. The case $\overleftarrow{s}_\gamma^n \leq (\mathcal{H}(X), X)$ is impossible since otherwise $(\mathcal{H}(X), X) \leq \overrightarrow{s}_X \in \gamma$ would imply that $\overleftarrow{s}_\gamma^n \leq \overrightarrow{s}_X$ meaning that $\overleftarrow{s}_\gamma^n$ and \overrightarrow{s}_X violate the consistency of γ . Therefore, we have either $(\mathcal{H}(X), X) \leq \overrightarrow{s}_\gamma^{n-1} \wedge \overleftarrow{s}_\gamma^{n-1} \in \sigma$ as desired, or we have $\overrightarrow{s}_\gamma^{n-1} \leq (\mathcal{H}(X), X) < \overleftarrow{s}_\gamma^n$. For this latter case, we show that there is a component $C \in \mathcal{H}(X)$ such that $\overrightarrow{s}_\gamma^{n-1} \leq (C, X)$ holds. This suffices to complete the proof, because then the inequalities $(\mathcal{H}(X) \setminus \{C\}, X) \leq (X, C) \leq \overleftarrow{s}_\gamma^{n-1}$ and $(\mathcal{H}(X) \setminus \{C\}, X) \leq (\mathcal{H}(X), X) < \overrightarrow{s}_\gamma^n$ give

$$(\mathcal{H}(X) \setminus \{C\}, X) \leq \overrightarrow{s}_\gamma^{n-1} \wedge \overleftarrow{s}_\gamma^{n-1} \in \sigma.$$

The separation $\overrightarrow{s}_\gamma^{n-1} \in O$ is, by definition, the supremum of some corridor δ of $\{(A, B) \in O_{\mathcal{H}} : |A \cap B| \leq \ell\}$ for some number $\ell \in \mathbb{N}$. Then every separation $(\mathcal{H}(Y), Y) \in \delta$ satisfies $(\mathcal{H}(Y), Y) \leq \overrightarrow{s}_\gamma^{n-1} \leq (\mathcal{H}(X), X)$. In particular, as the principal tree set $T_{\mathcal{H}}$ satisfies the conclusions of Theorem 14.1.7, every separation $(\mathcal{H}(Y), Y) \in \delta$ satisfies $(\mathcal{H}(Y), Y) \leq (C_X(Y), X)$. Hence in order to show that $\overrightarrow{s}_\gamma^{n-1} \leq (C, X)$ for some component $C \in \mathcal{H}(X)$, it suffices to show that

$C_X(Y) = C_X(Y')$ for every two separations $(\mathcal{K}(Y), Y)$ and $(\mathcal{K}(Y'), Y')$ in δ . Given $(\mathcal{K}(Y), Y)$ and $(\mathcal{K}(Y'), Y')$, consider any separation $(\mathcal{K}(Z), Z) \in \delta$ above the two. Then $(\mathcal{K}(Z), Z) \leq (C_X(Z), X)$ implies that both $C_X(Y)$ and $C_X(Y')$ are contained in $C_X(Z)$, giving $C_X(Y) = C_X(Y')$ as desired. \square

14.3. Overview of all duality results

In this section we summarise all duality theorems of this part. A very brief overview of the complementary structures is given by the following table:

	normal tree	tree-decomposition	other
combs	✓	✓	✓
stars	✓	✓	
dominated combs	✓	✓	
dominating stars	✓	✓	
undominated comb	✗	✓	✓
undominating star	✗	✓	✓

Here, a check mark means, for example, that we proved a duality theorem for combs in terms of normal trees, whereas the two crosses mean that normal trees cannot serve as complementary structures for undominated combs or undominating stars.

Finally, we summarise our duality theorem for combs, stars and combinations of the two explicitly in five theorems:

Theorem (Combs). *Let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. Then the following assertions are equivalent:*

- (i) G does not contain a *comb* attached to U ;
- (ii) there is a *rayless normal tree* $T \subseteq G$ that contains U (moreover, T can be chosen such that it contains U cofinally);
- (iii) G has a *rayless tree-decomposition* into parts each containing at most finitely many vertices from U and whose parts at non-leaves of the decomposition tree are all finite (moreover, the tree-decomposition displays $\partial_\Omega U$ and may be chosen with connected separators);
- (iv) for every infinite $U' \subseteq U$ there is a *critical vertex set* $X \subseteq V(G)$ such that infinitely many of the components in \mathcal{C}_X meet U' ;
- (v) G has a *U -rank*;
- (vi) G has a rooted tame *tree-decomposition* (T, \mathcal{V}) that covers U cofinally and satisfies the following four assertions:
 - (T, \mathcal{V}) is the squeezed expansion of a normal tree of G that contains the vertex set U cofinally;
 - every part of (T, \mathcal{V}) meets U finitely and parts at non-leaves are finite;
 - (T, \mathcal{V}) displays $\partial_T U \subseteq \text{crit}(G)$;
 - the rank of T is equal to the U -rank of G .

Theorem (Stars). *Let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. Then the following assertions are equivalent:*

- (i) G does not contain a *star* attached to U ;
- (ii) there is a *locally finite normal tree* $T \subseteq G$ that contains U and all whose rays are undominated in G (moreover, T can be chosen such that it contains U cofinally and every component of $G - T$ has finite neighbourhood);
- (iii) G has a *locally finite tree-decomposition* with finite and pairwise disjoint separators such that each part contains at most finitely many vertices of U (moreover, the tree-decomposition can be chosen with connected separators and such that it displays $\partial_T U \subseteq \Omega(G)$);

Theorem (Dominating stars and dominated comb). *Let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. Then the following assertions are equivalent:*

- (i) G does not contain a *dominating star* attached to U ;
- (ii) G does not contain a *dominated comb* attached to U ;
- (iii) there is a *normal tree* $T \subseteq G$ that contains U and all whose rays are undominated in G (moreover, the normal tree T can be chosen such that it contains U cofinally and every component of $G - T$ has finite neighbourhood);
- (iv) G has a *tree-decomposition* (T, \mathcal{V}) such that
 - each part contains at most finitely many vertices from U ;
 - all parts at non-leaves of T are finite;
 - (T, \mathcal{V}) has essentially disjoint connected separators;
 - (T, \mathcal{V}) displays $\partial_\Omega U$.

Theorem (Undominated combs). *Let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. Then the following assertions are equivalent:*

- (i) G does not contain an *undominated comb* attached to U ;
- (ii) G has a *star-decomposition* with finite separators such that U is contained in the central part and all undominated ends of G live in the leaves' parts (moreover, the star-decomposition can be chosen with pairwise disjoint and connected separators);
- (iii) G has a connected *subgraph* that contains U and all whose rays are dominated in it (moreover, the subgraph can be chosen so as to reflect the ends in its closure).

Moreover, if U is normally spanned in G , we may add

- (iv) there is a *rayless tree* $T \subseteq G$ that contains U .

Theorem (Undominating stars). *Let G be any connected graph and let $U \subseteq V(G)$ be any vertex set. Then the following assertions are equivalent:*

- (i) *G does not contain an undominating star attached to U ;*
- (ii) *there is a tough subgraph $H \subseteq G$ that contains U ;*
- (iii) *G has a tame star-decomposition such that U is contained in the central part and every critical vertex set of G lives in a leaf's part.*

Moreover, if U is normally spanned, we may add

- (iv) *there is a locally finite normal tree $T \subseteq G$ that contains U .*

15. End-faithful spanning trees

Schmidt [20, 64] characterised the class of rayless graphs by an ordinal rank function, which makes it possible to prove statements about rayless graphs by transfinite induction. For example, Bruhn, Diestel, Georgakopoulos and Sprüssel [3, 20] proved the unfriendly partition conjecture for the class of rayless graphs in this way.

At the turn of the millennium, Halin [41] asked in his legacy collection of problems whether Schmidt's rank can be generalised to characterise other important classes of graphs besides the class of rayless graphs. In this chapter we answer Halin's question in the affirmative: we characterise two important classes of graphs by an ordinal rank function.

As our first main result, we characterise for every uncountable cardinal κ the class of graphs without a T_κ minor by an ordinal rank function that we call the κ -rank (recall that T_κ denotes the κ -branching tree):

Theorem 15.1. *For every graph G and every uncountable cardinal κ the following assertions are equivalent:*

- (i) G contains no T_κ minor;
- (ii) G has a κ -rank.

This extends Seymour and Thomas' characterisations [66]. We remark that, for regular uncountable cardinals κ , they also showed that a graph contains a T_κ minor if and only if it contains a subdivision of T_κ .

Our second main result addresses another largely open problem raised by Halin. Call a spanning tree T of a graph G *end-faithful* if the map $\varphi: \Omega(T) \rightarrow \Omega(G)$ satisfying $\omega \subseteq \varphi(\omega)$ is bijective. Here, $\Omega(T)$ and $\Omega(G)$ denote the set of ends of T and of G , respectively. Halin [37] conjectured that every connected graph has an end-faithful spanning tree. However, Seymour and Thomas [65] and Thomassen [70] constructed uncountable counterexamples; for instance, there

exists a connected graph that has precisely one end but all whose spanning trees must contain a subdivision of T_{\aleph_1} . Ever since, it has been an open problem to characterise the class of graphs that admit an end-faithful spanning tree.

Normal spanning trees are important examples of end-faithful spanning trees. Given a graph G , a rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T , cf. [20]. Call a set U of vertices of a graph G *normally spanned* in G if U is contained in a tree $T \subseteq G$ that is normal in G . The graph G is *normally spanned* if $V(G)$ is normally spanned in G , i.e., if G has a normal spanning tree. Thus, every normally spanned graph has an end-faithful spanning tree.

A second existence result for end-faithful spanning trees is due to Polat [60] and directly addresses the counterexamples by Seymour and Thomas and by Thomassen: every connected graph that does not contain a subdivision of T_{\aleph_1} has an end-faithful spanning tree.

As our second main result, we determine a new subclass of the class of graphs with an end-faithful spanning tree. Call a connected graph G *normally traceable* if it has a rayless tree-decomposition into parts that are normally spanned in G . For the definition of tree-decompositions see [20].

Theorem 15.2. *Every normally traceable graph has an end-faithful spanning tree.*

Our theorem easily extends the two known existence results for end-faithful spanning trees: On the one hand, every normally spanned graph has a trivial tree-decomposition into one normally spanned part. On the other hand, every connected graph without a subdivision of T_{\aleph_1} has a rayless tree-decomposition into countable parts by the characterisation of Seymour and Thomas [66], and countable vertex sets are normally spanned.

In both cases, the extension is proper: The \aleph_1 -*branching trees with tops* are the graphs obtained from the rooted T_{\aleph_1} by selecting uncountably many rooted rays and adding for every selected ray R a new vertex, its *top*, and joining it to infinitely many vertices of R [29]. Every T_{\aleph_1} with tops has a star-decomposition into normally spanned parts where T_{\aleph_1} forms the central part and each top plus its neighbours forms a leaf's part. However, not every T_{\aleph_1} with tops has a normal spanning tree [29, 54], and every T_{\aleph_1} with tops contains T_{\aleph_1} as a subgraph.

Carmesin [15] has amended Halin's conjecture about end-faithful spanning trees:

He showed that every connected graph G has a spanning tree T that is *end-faithful* for its undominated ends in that every undominated end ω of G is uniquely represented by an end η of T with $\eta \subseteq \omega$. Recall that a vertex v of a graph G *dominates* a ray $R \subseteq G$ if there is an infinite v - R fan in G . Rays not dominated by any vertex are *undominated*. An end of G is *dominated* or *undominated* if one (equivalently: each) of its rays is dominated or undominated, respectively, see [20].

Carmesin pointed out that his result becomes false when one replaces ‘is end-faithful for’ with ‘reflects’ in its wording. Here, a spanning tree T of a graph G *reflects* the undominated ends of G if it is end-faithful for the undominated ends of G and every end η of T represents an undominated end ω of G with $\eta \subseteq \omega$. In Chapter 13, we proved that normally spanned graphs have spanning trees reflecting their undominated ends. As our third main result, we extend this to the class of normally traceable graphs:

Theorem 15.3. *Every normally traceable graph has a spanning tree that reflects its undominated ends.*

Our theorem extends two existence results on rayless spanning trees. For a connected graph G , having a rayless spanning tree is equivalent to all the ends of G being dominated if G is normally spanned [7] or if G does not contain a subdivision of T_{\aleph_1} [60]. The following corollary extends these results, and any T_{\aleph_1} with all tops witnesses that this extension is proper.

Corollary 15.4. *For every normally traceable graph G , having a rayless spanning tree is equivalent to all the ends of G being dominated.* \square

Finally, as our fifth main result we characterise the class of normally traceable graphs by an ordinal rank function that we call the normal rank:

Theorem 15.5. *For every graph G the following assertions are equivalent:*

- (i) G is normally traceable;
- (ii) G has a normal rank.

We use this in the proofs of all our results on normally traceable graphs.

This chapter is organised as follows. In Section 15.1 we introduce the κ -rank and prove Theorem 15.1. Then, in Section 15.2 we introduce the normal rank and prove Theorem 15.5. We prove Theorem 15.2 in Section 15.3 and we prove Theorem 15.3 in Section 15.4.

15.1. Ranking T_κ -free graphs

In this section we characterise for every uncountable cardinal κ the class of graphs without a T_κ minor by an ordinal rank function that we call the κ -rank.

Suppose that κ is any infinite cardinal. Let us assign κ -rank 0 to all the graphs of order less than κ . Given an ordinal $\alpha > 0$, we assign κ -rank α to every graph G that does not already have a κ -rank $< \alpha$ and which has a set X of less than κ many vertices such that every component of $G - X$ has some κ -rank $< \alpha$. Note that the \aleph_0 -rank is Schmidt's rank [20, 64].

The κ -rank behaves quite similar to Schmidt's rank [20, p. 243]: When disjoint graphs G_i have κ -ranks $\xi_i < \alpha$, their union clearly has a κ -rank of at most α ; if the union is finite, it has κ -rank $\max_i \xi_i$. Induction on α shows that subgraphs of graphs of κ -rank α also have a κ -rank of at most α . Conversely, joining less than κ many new vertices to a graph, no matter how, will not change its κ -rank.

Not every graph has a κ -rank. Indeed, an inflated κ -branching tree cannot have a κ -rank, since deleting less than κ many of its vertices always leaves a component that contains another inflated κ -branching tree. As subgraphs of graphs with a κ -rank also have a κ -rank, this means that only graphs without a T_κ minor can have a κ -rank. But all these do:

Theorem 15.1. *For every graph G and every uncountable cardinal κ the following assertions are equivalent:*

- (i) G contains no T_κ minor;
- (ii) G has a κ -rank.

Hence the κ -rank characterises the class of graphs without a T_κ minor.

Our proof relies upon a theorem by Seymour and Thomas [66] that we recall here. For every set M we denote by $[M]^{<\kappa}$ the set of all subsets of M of cardinality $< \kappa$. Now, given a graph G , we write \mathcal{C}_X for the set of components of $G - X$ for every set $X \subseteq V(G)$ of vertices. An *escape* of order κ in G is a function σ which assigns to each $X \in [V(G)]^{<\kappa}$ the vertex set $V[\mathcal{C}] := \bigcup \{ V(C) \mid C \in \mathcal{C} \}$ of a subset $\mathcal{C} \subseteq \mathcal{C}_X$ in such a way that:

- (i) if $X \subseteq Y$, then $\sigma(Y) \subseteq \sigma(X)$,
- (ii) if $X \subseteq Y$, then for $\sigma(X) = V[\mathcal{C}]$ every component $C \in \mathcal{C}$ intersects $\sigma(Y)$,
and
- (iii) $\sigma(\emptyset) \neq \emptyset$.

We speak of (i), (ii) and (iii) as the first, second and third *escape axioms*. We remark that Seymour and Thomas' escapes can in fact be seen as more general predecessors of directions which describe the ends of a graph by a theorem of Diestel and Kühn [27].

Theorem 15.1.1 ([66, Theorem 1.3]). *For every graph G and every uncountable cardinal κ the following assertions are equivalent:*

- (i) G contains a T_κ minor;
- (ii) G has an escape of order κ .

We are now ready to prove Theorem 15.1:

Proof of Theorem 15.1. We show the equivalence $\neg(\text{i}) \leftrightarrow \neg(\text{ii})$. The forward implication has already been pointed out above. For the backward implication suppose that G has no κ -rank; we show that G must contain a T_κ minor. By Theorem 15.1.1 it suffices to find an escape of order κ in G . We define a candidate σ for such an escape as follows. Given any vertex set $X \in [V(G)]^{<\kappa}$ we call a component C of $G - X$ *bad* if it has no κ -rank, and we let $\sigma(X) := V[\mathcal{C}]$ for the collection \mathcal{C} of all the bad components of $G - X$. It remains to show that σ satisfies all three escape axioms.

Having no κ -rank is closed under taking supergraphs, so the first axiom holds. For the second axiom, let any two vertex sets $X \subseteq Y \in [V(G)]^{<\kappa}$ be given, and consider any component $C \in \mathcal{C}$ for $\sigma(X) = V[\mathcal{C}]$. Then $C - Y$ must have a component that has no κ -rank, and this component then is bad as desired. Finally, the third axiom holds because the graph G must have a bad component. \square

15.2. Normally traceable graphs

In this section we characterise the class of normally traceable graphs by an ordinal rank function that we call the normal rank.

Let G be any connected graph. A connected subgraph $H \subseteq G$ has *normal rank* 0 in G if the vertex set of H is normally spanned in G . Given an ordinal $\alpha > 0$, a connected subgraph $H \subseteq G$ has *normal rank* α in G if it does not already have a normal rank $< \alpha$ in G and if there is a vertex set $X \subseteq V(H)$ that is normally spanned in G such that every component of $H - X$ has some normal rank $< \alpha$ in G .

The graph G has *normal rank* α for an ordinal α if G has normal rank α in G .

Theorem 15.5. *For every connected graph G the following assertions are equivalent:*

- (i) G is normally traceable;
- (ii) G has a normal rank.

Moreover, if G has a tree-decomposition witnessing that G is normally traceable, then G has normal rank at most the rank of the decomposition tree. Conversely, if G has a normal rank, then G is normally traceable and this is witnessed by a tree-decomposition whose decomposition tree has as rank the normal rank of G .

Before we prove this theorem, we point out a few properties of the normal rank.

Lemma 15.2.1. *Let G be any connected graph.*

- (i) *If G has \aleph_1 -rank α , then G has some normal rank $\leq \alpha$.*
- (ii) *There are graphs that have a normal rank but that have neither an \aleph_1 -rank nor a normal spanning tree.*

Proof. (i) We show that every connected subgraph $H \subseteq G$ of \aleph_1 -rank α has normal rank $\leq \alpha$ in G , by induction on α ; for $H = G$ this establishes (i). Any connected countable subgraph of G is normally spanned in G by Jung's Theorem 11.2.5, so the base case holds. For the induction step suppose that $\alpha > 0$. We find a countable vertex set $X \subseteq V(H)$ so that every component of $H - X$ has some \aleph_1 -rank $< \alpha$. As X is countable it is also normally spanned in G . By the induction

hypothesis every component of $H - X$ has normal rank $< \alpha$ in G . Hence X witnesses that H has normal rank $\leq \alpha$ in G .

(ii) Let G be any T_{\aleph_1} with tops. Then G has normal rank 1 because $G - T_{\aleph_1}$ consists only of isolated vertices. However, G has no \aleph_1 -rank by Theorem 15.1, and G has no normal spanning tree as pointed out by Diestel and Leader [29]. \square

Lemma 15.2.2. *Let $H \subseteq H' \subseteq G$ be any three connected graphs.*

(i) *If H' has normal rank α in G , then H has normal rank $\leq \alpha$ in G .*

(ii) *If H has normal rank α in G , then H has normal rank $\leq \alpha$ in H' .*

In particular, if H has normal rank α in G , then H has normal rank $\leq \alpha$.

Proof. (i) Induction on α . If $\alpha = 0$, then the vertex set of H' is normally spanned in G ; in particular, the vertex set of $H \subseteq H'$ is normally spanned in G .

Otherwise $\alpha > 0$. Then there exists a vertex set $X \subseteq V(H')$ that is normally spanned in G such that every component of $H' - X$ has normal rank $< \alpha$ in G . Every component of $H - X$ is contained in a component of $H' - X$ and hence has normal rank $< \alpha$ in G by the induction hypothesis. Thus, H has normal rank $\leq \alpha$ in G .

(ii) Induction on α . If $\alpha = 0$, then the vertex set of H is normally spanned in G . In particular, by Jung's Theorem 11.2.5, the vertex set of H is normally spanned in $H' \subseteq G$, so H has normal rank 0 in H' as desired.

Otherwise $\alpha > 0$. Then there exists a vertex set $X \subseteq V(H)$ that is normally spanned in G such that every component of $H - X$ has normal rank $< \alpha$ in G . Note that X is also normally spanned in $H' \subseteq G$ by Jung's Theorem 11.2.5. By the induction hypothesis, every component of $H - X$ has normal rank $< \alpha$ in H' . Thus, H has normal rank $\leq \alpha$ in H' . \square

Proof of Theorem 15.5. Let G be any connected graph. To show the equivalence (i) \leftrightarrow (ii) together with the 'moreover' part of the theorem, it suffices to show the following two assertions:

- (1) If G has a tree-decomposition witnessing that G is normally traceable, then G has a normal rank which is at most the rank of the decomposition tree.
- (2) If G has a normal rank, then G is normally traceable and this is witnessed by a tree-decomposition whose decomposition tree has rank at most the normal rank of G .

(1) We show that every connected subgraph $H \subseteq G$ that has a rayless tree-decomposition (T, \mathcal{V}) into parts that are normally spanned in G does have normal rank $\leq \alpha$ in G for α the rank of T . We prove this by induction on α ; for $H = G$ and α equal to the rank of the decomposition tree of some tree-decomposition of G witnessing that G is normally traceable we obtain (1). If H and (T, \mathcal{V}) are such that $\alpha = 0$, then T is finite, and hence the union of all the parts in \mathcal{V} is normally spanned in G by Jung's Theorem 11.2.5; in particular, $V(H)$ is normally spanned in G and hence has normal rank 0 in G .

Otherwise H and (T, \mathcal{V}) are such that $\alpha > 0$. Let $W \subseteq V(T)$ be any finite vertex set such that every component of $T - W$ has rank $< \alpha$. Then the vertex set $X := \bigcup_{t \in W} V_t \subseteq V(H)$ is normally spanned in G by Jung's Theorem 11.2.5. Every component of $H - X$ is contained in $\bigcup_{t \in T'} G[V_t]$ for some component T' of $T - W$, so by the induction hypothesis every component of $H - X$ has normal rank $< \alpha$ in G . Thus, H has normal rank $\leq \alpha$ in G .

(2) Suppose that G is any connected graph that has a normal rank. We show that every connected subgraph $H \subseteq G$ of normal rank α in G has a rayless tree-decomposition (T, \mathcal{V}) into parts that are normally spanned in G such that T has rank $\leq \alpha$, by induction on the normal rank α of H in G ; for $H = G$ this establishes (2). If $\alpha = 0$, then $V(H)$ is normally spanned in G and the trivial tree-decomposition of H into the single part $V(H)$ is as desired.

Otherwise $\alpha > 0$. Then there exists a vertex set $X \subseteq V(H)$ that is normally spanned in G such that every component of $H - X$ has normal rank $< \alpha$ in G . By the induction hypothesis, every component C of $H - X$ has a rayless tree-decomposition (T_C, \mathcal{V}_C) with $\mathcal{V}_C = (V_C^t \mid t \in T_C)$ such that every part is normally spanned in G and the rank of T_C is $< \alpha$. Without loss of generality the trees T_C are pairwise disjoint. We choose from every tree T_C an arbitrary node $t_C \in T_C$. Then we let the tree T be obtained from the disjoint union $\bigcup_C T_C$ by adding a new vertex t_* that we join to all the chosen nodes t_C . We define the family $\mathcal{V} = (V^t \mid t \in T)$ by letting $V^t := V_C^t \cup X$ for all $t \in T_C \subseteq T$ and $V^{t_*} := X$. Then (T, \mathcal{V}) is a rayless tree-decomposition of H into parts that are normally spanned in G by Jung's Theorem 11.2.5, and the rank of T is $\leq \alpha$ because every component of $T - t_*$ has rank $< \alpha$. \square

15.3. End-faithful spanning trees

In this section we prove that every normally traceable graph has an end-faithful spanning tree. Our proof requires some preparation.

Suppose that H is any subgraph of G and $\varphi: \Omega(H) \rightarrow \Omega(G)$ is the natural map satisfying $\eta \subseteq \varphi(\eta)$ for every end η of H . Furthermore, suppose that a set $\Psi \subseteq \Omega(G)$ of ends of G is given. We say that H is *end-faithful* for Ψ if $\varphi \upharpoonright \varphi^{-1}(\Psi)$ is injective and $\text{im}(\varphi) \supseteq \Psi$. And H *reflects* Ψ if φ is injective with $\text{im}(\varphi) = \Psi$. A spanning tree of G that is end-faithful for all the ends of G is *end-faithful*. The following lemma is a reformulation of Lemma 11.1.11

Lemma 15.3.1. *If G is any graph and $T \subseteq G$ is any normal tree, then T reflects the ends of G in the closure of T .*

Lemma 15.3.2. *Let G be any graph and let $\Psi \subseteq \Omega(G)$ be any set of ends of G . If $H \subseteq G$ is a spanning forest that reflects Ψ and T is a component of H such that every other component of H has a neighbour in T , then G has a spanning tree that reflects Ψ .*

Proof. Fix for every component $T' \neq T$ of H an edge $e_{T'}$ between T' and T . It is straightforward to check that the spanning tree consisting of H plus all the edges $e_{T'}$ reflects the ends in Ψ . \square

Lemma 15.3.3. *Let G be any graph with a spanning tree $T \subseteq G$ that reflects a set $\Psi \subseteq \Omega(G)$ and let $R \subseteq G$ be a ray from some end in Ψ . Then there exists a spanning tree $T' \subseteq G$ that reflects Ψ and contains R .*

Moreover, T' can be chosen such that no end other than the end of R lies in the closure of the symmetric difference $E(T) \Delta E(T')$ (viewed as a subgraph of G).

The ‘moreover’ part of the lemma says that T and T' differ only locally. Note that there may also be no end in the closure of $E(T) \Delta E(T')$.

Proof. Given $T \subseteq G$, Ψ and R , we root T arbitrarily and write ω for the end of R in G . Furthermore, we write R_T for the unique rooted ray in T that is equivalent to R , and we pick a sequence P_0, P_1, \dots of pairwise disjoint R - R_T paths in G . We write C for the comb $C := R \cup \bigcup_n P_n$ consisting of R and all the paths P_n , and we write U for the vertex set of the subtree $[C]_T$ of T . Note that $R_T \subseteq [C]_T$ because the paths P_0, P_1, \dots meet R_T infinitely often. By standard arguments we

have $\partial_\Omega C = \{\omega\}$, and so $\partial_\Omega U = \{\omega\}$ follows by Lemma 11.1.13. Since T reflects Ψ and $\lceil C \rceil_T$ contains only rays from ω , we deduce that $\lceil C \rceil_T$ is either rayless or one-ended. As $\lceil C \rceil_T$ contains the ray R_T , it is one-ended.

Next, we define an edge set $F \subseteq E(\lceil C \rceil_T)$, as follows. If R has a tail in R_T , then we set $F = \emptyset$. Otherwise R has no tail in R_T . Then we select infinitely many pairwise edge-disjoint C -paths Q_0, Q_1, \dots in the ray R_T (these exist because R has no tail in R_T). We choose one edge of every path Q_n and we let F consist of all the chosen edges, completing the definition of F .

The graph $(\lceil C \rceil_T \cup C) - F$ is a connected subgraph of G and inside it, we extend C arbitrarily to a spanning tree T_R . Then T_R has vertex set U , and T_R reflects $\{\omega\}$: Every ray R' in T_R that is disjoint from R meets at most one component of $C - R$ because C and R' are contained in the tree T_R , and hence R' must have a tail in $\lceil C \rceil_T - C$. But $\lceil C \rceil_T$ contains just one rooted ray, namely the ray R_T , and either R_T contains a tail of R or F consists of infinitely many edges of R_T , contradicting the existence of R' in $T_R \subseteq (\lceil C \rceil_T \cup C) - F$. It remains to extend T_R to a spanning tree of G reflecting Ψ . For this, we consider the collection $\{T_i \mid i \in I\}$ of all the components of $T - U$. By the choice of U , every end ω' of G other than ω is still represented by an end of one of the trees T_i : Indeed, if ω' is an end of G other than ω , then it does not lie in the closure of U , and hence every ray in ω' has a tail that avoids U . In particular, every ray in T that lies in ω' has some tail that avoids U . Therefore, the union of T_R and all the trees T_i is a spanning forest of G reflecting Ψ .

We extend this spanning forest to a spanning tree T' by adding all the T_i - T_R edges of T for every $i \in I$ (note that T contains precisely one T_i - T_R edge for every $i \in I$ as $T \cap G[U] = \lceil C \rceil_T$ is connected). Then T' reflects Ψ again by Lemma 15.3.2. To see $\partial_\Omega(E(T) \Delta E(T')) \subseteq \{\omega\}$ recall $\partial_\Omega G[U] = \{\omega\}$ and note that the symmetric difference is contained in $G[U]$ entirely. \square

Lemma 15.3.4. *Let G be any graph and let $X \subseteq V(G)$ be any vertex set.*

- (i) *Every end of G is contained in the closure of X in G or in the closure of some component of $G - X$ in G .*
- (ii) *Every end of G that is contained in the closure of two distinct components of $G - X$ in G is also contained in the closure of X in G .*

Proof. (i) Let ω be any end of G and let $R \in \omega$ be any ray. Then either the vertex set of R intersects X infinitely, or R has a tail that is contained in some component C of $G - X$. In the first case, ω is contained in the closure of X , and in the second case it is contained in the closure of C in G .

(ii) Let C and C' be two distinct components of $G - X$ and suppose that ω is any end of G that is contained in the closure of both C and C' in G . If $S \subseteq V(G)$ is any finite vertex set, then the component $C(S, \omega)$ meets both C and C' . As X separates C and C' in G it follows that $C(S, \omega)$ meets X as well. We conclude that ω is contained in the closure of X in G . \square

Lemma 15.3.5. *Let G be any connected graph, let $X \subseteq V(G)$ be normally spanned in G and let C be any component of $G - X$ so that $G[C \cup X]$ is connected. If C has normal rank ξ in G , then $G[C \cup X]$ has normal rank $\leq \xi$.*

Proof. Suppose that C is a component of $G - X$ that has normal rank ξ in G . If $\xi = 0$, then $V(C)$ is normally spanned in G and C has a normal spanning tree by Jung's Theorem 11.2.5, so C has normal rank 0 as desired. Otherwise there is a vertex set $Y \subseteq V(C)$ that is normally spanned in G and satisfies that every component of $C - Y$ has normal rank $< \xi$ in G . Note that $X \cup Y$ is normally spanned in G by Jung's Theorem 11.2.5. Therefore $X \cup Y$ witnesses that $G[C \cup X]$ has normal rank $\leq \xi$ in G . Finally, Lemma 15.2.2 (ii) implies that $G[C \cup X]$ has normal rank $\leq \xi$. \square

Theorem 15.2. *Every normally traceable graph has an end-faithful spanning tree.*

Proof. By Theorem 15.5 we may prove the statement via induction on the normal rank of G . If G has normal rank 0, then it has a normal spanning tree, and normal spanning trees are end-faithful. For the induction step suppose that G has normal rank $\alpha > 0$, and let $X \subseteq V(G)$ be any vertex set that is normally spanned in G and satisfies that every component of $G - X$ has normal rank $< \alpha$ in G . By replacing X with the vertex set of any normal tree in G that contains X , we may assume that X is the vertex set of a normal tree $T_{\text{NT}} \subseteq G$; indeed, every component of $G - X$ still has normal rank $< \alpha$ in G by Lemma 15.2.2 (i). Note that, by Lemma 15.3.1, the tree T_{NT} reflects the ends of G in the closure of X .

By Lemma 15.3.4 (i), every end of G is contained in the closure of X in G or in the closure of some component of $G - X$. And by Lemma 15.3.4 (ii), every end

of G that is contained in the closure of two distinct components of $G - X$ in G is also contained in the closure of X in G . Thus, by Lemma 15.3.2 it suffices to find in each component C of $G - X$ a spanning forest H_C so that every component of H_C sends an edge in G to T_{NT} and so that H_C reflects $\partial_\Omega C \setminus \partial_\Omega X$.

For this, consider any component C of $G - X$. Let P be the (possibly one-way infinite) path in T_{NT} that is formed by the down-closure of $N(C)$ in T_{NT} . Then by Lemma 15.3.5 the graph $G[C \cup P]$ has normal rank $< \alpha$, and therefore satisfies the induction hypothesis. Hence we find an end-faithful spanning tree T_C of $G[C \cup P]$. By Lemma 15.3.3 we may assume that the path P is a subgraph of T_C if this path is a ray. It is straightforward to check that $H_C := T_C - X$ is as desired. \square

15.4. Trees reflecting the undominated ends

In this section we prove that every normally traceable graph has a spanning tree that reflects its undominated ends. Our proof requires the following theorem:

Theorem 15.4.1 ([7, Theorem 3.2]). *Let G be any graph and let $U \subseteq V(G)$ be normally spanned in G . Then there is a tree $T \subseteq G$ that contains U and reflects the undominated ends in the closure of U .*

Theorem 15.3. *Every normally traceable graph has a spanning tree that reflects its undominated ends.*

Proof. By Theorem 15.5 we may prove the statement via induction on the normal rank of G . If G has normal rank 0, then it is normally spanned. Thus, by Theorem 15.4.1, the graph G has a spanning tree that reflects its undominated ends. For the induction step suppose that G has normal rank $\alpha > 0$, and let $X \subseteq V(G)$ be any vertex set that is normally spanned in G and satisfies that every component of $G - X$ has normal rank $< \alpha$ in G . By replacing X with any normal tree in G that contains X , we may assume that X is the vertex set of a normal tree $T_{\text{NT}} \subseteq G$; indeed, every component of $G - X$ still has normal rank $< \alpha$ in G by Lemma 15.2.2 (i).

We claim that it suffices to find in every component C of $G - T_{\text{NT}}$ a spanning forest H_C such that every component of H_C sends an edge in G to T_{NT} and H_C reflects the undominated ends of G in $\partial_\Omega C \setminus \partial_\Omega T_{\text{NT}}$. This can be seen as follows. Suppose that we find such a spanning forest H_C in every component C of $G - X$.

By Theorem 15.4.1 we find a tree $T_{\text{UD}} \subseteq G$ that contains $X = V(T_{\text{NT}})$ and reflects the undominated ends of G in the closure of T_{NT} . Then we set $H'_D := H_C \cap D$ for every component D of $G - T_{\text{UD}}$ and the component C of $G - X$ containing it. Now consider the spanning forest H of G that is the union of all forests H'_D with the tree T_{UD} . We show that H reflects the undominated ends of G .

On the one hand, all the rays in H belong to undominated ends of G , and H contains no two disjoint rays from the same undominated end of G . On the other hand, let ω be any undominated end of G . If ω lies in the closure of T_{NT} , then $T_{\text{UD}} \subseteq H$ contains a ray from ω . Otherwise ω does not lie in the closure of T_{NT} . Then ω lies in the closure of a component C of $G - T_{\text{NT}}$ by Lemma 15.3.4 (i), so H_C contains a ray R from ω . Furthermore, ω does not lie in the closure of T_{UD} because by the star-comb lemma every tree in G contains a ray from every undominated end in its closure, and T_{UD} reflects only the undominated ends of G in the closure of T_{NT} ; in particular, R has a tail $R' \subseteq R$ that avoids T_{UD} . Then $R' \subseteq H'_D \subseteq H$ for the component D of $G - T_{\text{UD}}$ that contains R' , completing the proof that H reflects the undominated ends of G . It remains to show that G has a spanning tree that reflects the undominated ends of G ; such a tree arises from H by Lemma 15.3.2.

To complete the proof, we show that every component C of $G - T_{\text{NT}}$ has a spanning forest H_C such that every component of H_C sends an edge in G to T_{NT} and H_C reflects the undominated ends of G in $\partial_\Omega C \setminus \partial_\Omega T_{\text{NT}}$. So let C be any component of $G - X$ and let P be the (possibly one-way infinite) path in T_{NT} that is formed by the down-closure of $N(C)$ in T_{NT} . Then by Lemma 15.3.5 the graph $G[C \cup P]$ satisfies the induction hypothesis. Hence we find a spanning tree T_C of $G[C \cup P]$ reflecting the undominated ends of $G[C \cup P]$. By Lemma 15.3.3 we may assume that the path P is a subgraph of T_C if this path is an undominated ray in $G[C \cup P]$. It is straightforward to check that $H_C := T_C - X$ is as desired. \square

Appendix

English summary

In the following we give a brief summary of the results that we provide in the three parts of this dissertation.

I. Monochromatic generalised paths

Answering a conjecture of Soukup in the affirmative, we prove that every r -edge-coloured complete bipartite graph with bipartition classes of the same infinite cardinality admits a partition of its vertex set into $2r - 1$ monochromatic generalised paths (Theorem 2.2). This bound is best possible in the sense that for every $\kappa \geq \aleph_0$ there are r -edge-colourings of $K_{\kappa, \kappa}$ for which the graph cannot be partitioned into $2r - 2$ monochromatic paths.

Furthermore, our discussion leads to a conceptually simpler closing argument for Soukup's Theorem 2.1, stating that the vertex set of every k -edge-coloured complete graph of infinite cardinality can be partitioned into monochromatic generalised paths of different colours.

The key idea for our proofs is to combine Soukup's techniques from his original paper with our notion and construction method of X -robust paths.

II. Ends of digraphs

In this part we develop an end space theory for directed graphs. As for undirected graphs, the ends of a digraph are points at infinity to which its rays converge. Unlike for undirected graphs, some ends are joined by limit edges; these are crucial for obtaining the end space of a digraph as a natural (inverse) limit of its finite contraction minors.

As our main result in Chapter 8, we show that the notion of directions of an undirected graph, a tangle-like description of its ends, extends to digraphs: there is a one-to-one correspondence between the directions of a digraph and its ends and limit edges (Theorem 8.2 and Theorem 8.3).

In the course of this we develop a number of fundamental tools and techniques for the study of ends of digraphs, such as the necklace lemma and the directed star-comb lemma.

In Chapter 9, we introduce the topological space $|D|$ formed by a digraph D together with its ends and limit edges. We then show that the digraphs that are compactified by $|D|$ are precisely the solid ones (Theorem 9.1). Furthermore, we show that if $|D|$ is compact, it is the inverse limit of finite contraction minors of D (Theorem 9.2).

To illustrate the use of this, we extend to the space $|D|$ two statements about finite digraphs that do not generalise verbatim to infinite digraphs. The first statement is the characterisation of finite Eulerian digraphs by the condition that the in-degree of every vertex equals its out-degree (see Theorem 9.3). The second statement is the characterisation of strongly connected finite digraphs by the existence of a closed Hamilton walk (see Theorem 9.4).

In Chapter 10 we introduce a concept of depth-first search trees in infinite digraphs, which we call normal spanning arborescences.

We show that normal spanning arborescences are end-faithful: every end of a digraph is represented by exactly one ray in the normal spanning arborescence that starts from the root (Theorem 10.1). We further show that every normal spanning arborescence of a solid digraph reflects its horizon (Theorem 10.2). Finally, we prove a Jung-type criterion for the existence of normal spanning arborescences: Every digraph with a vertex that can reach all the other vertices and whose vertex set can be written as a countable union of dispersed set has a normal spanning arborescence (Theorem 10.3).

III. Stars and combs

In the first four chapters of this part we determine structures whose existence is complementary to the existence of two substructures that are particularly fundamental to the study of connectedness in infinite graphs: stars and combs.

In Chapter 11 we determine structures whose existence is complementary to the existence of arbitrary stars and combs. We offer several duality theorems: in terms of normal trees (Theorem 11.1 and Theorem 11.6), tree-decompositions (Theorem 11.2 and Theorem 11.7), critical vertex sets (Theorem 11.3) and the

U -rank (Theorem 11.4).

In Chapter 12 we determine structures whose existence is complementary to the existence of dominating stars and dominated combs. As dominating stars exist if and only if dominated combs do, the structures complementary to them coincide. Like for arbitrary stars and combs, our duality theorems for dominated combs (and dominating stars) are phrased in terms of normal trees (Theorem 12.1) or tree-decompositions (Theorem 12.2).

The complementary structures we provide for dominated combs unify those for stars and combs and allow us to derive our duality theorems for stars and combs from those for dominated combs. This is surprising given that our complementary structures for stars and combs are quite different: those for stars are locally finite whereas those for combs are rayless.

In Chapter 13 we determine structures whose existence is complementary to the existence of undominated combs. We describe their complementary structures in terms of rayless trees (Theorem 13.1) and of star-decompositions (Theorem 13.7).

Applications include a complete characterisation, in terms of normal spanning trees, of the graphs whose rays are dominated but which have no rayless spanning tree (Corollary 13.1.3). Only two such graphs had so far been constructed, by Seymour and Thomas [65] and by Thomassen [70]. As a corollary, we show that graphs with a normal spanning tree have a rayless spanning tree if and only if all their rays are dominated (Corollary 13.2).

Another application settles a problem left unsolved by Carmesin [15]: The graphs whose undominated ends are reflected by a suitable spanning tree can be characterised in terms of normal spanning trees (Theorem 13.3). In particular, we show that every graph that has a normal spanning tree does have a spanning tree reflecting its undominated ends.

In Chapter 14 we determine structures whose existence is complementary to the existence of undominating stars. This completes our solution to the problem of finding complementary structures for stars, combs, and their combinations. Our duality theorems are phrased in terms of end-compactified subgraphs (Theorem 14.1) and tree-decompositions (Theorem 14.2).

The last chapter of this part, Chapter 15, can be seen as a spin-off of its previous chapters: there we consider two problems raised by Halin and provide further applications of our duality theorems and the tools that we provide in their proofs.

Schmidt [64] characterised the class of rayless graphs by an ordinal rank function, which makes it possible to prove statements about rayless graphs by transfinite induction. Halin [41] asked in his legacy collection of problems whether Schmidt's rank function can be generalised to characterise other important classes of graphs. We answer Halin's question in the affirmative: we characterise two important classes of graphs by an ordinal rank function.

Seymour and Thomas have characterised for every uncountable cardinal κ the class of graphs without a T_κ minor. We extend their characterisations by an ordinal rank function, one for every uncountable cardinal κ (Theorem 15.1).

Another largely open problem raised by Halin asks for a characterisation of the class of graphs with an end-faithful spanning tree. A well-studied subclass is formed by the graphs with a normal spanning tree. We determine a larger subclass, the class of normally traceable graphs, which consists of the connected graphs with a rayless tree-decomposition into normally spanned parts (Theorem 15.2). Investigating the class of normally traceable graphs further, we prove that all its graphs have spanning trees reflecting their undominated ends (Theorem 15.3). Our proofs rely on a characterisation of the class of normally traceable graphs by an ordinal rank function that we provide (Theorem 15.5).

Deutsche Zusammenfassung

Im Folgenden fassen wir die Resultate dieser Dissertation zusammen.

I. Monochromatische verallgemeinerte Wege

Wir bestätigen Soukups Vermutung, dass es zu jeder Färbung der Kanten eines balancierten vollständig bipartiten Graphen mit r Farben eine Partition der Eckenmenge in $2r - 1$ monochromatische verallgemeinerte Wege gibt (Satz 2.2). Diese Schranke ist bestmöglich in dem Sinne, dass es für alle $\kappa \geq \aleph_0$ Färbungen der Kanten von $K_{\kappa, \kappa}$ mit r Farben gibt, sodass es keine Partition der Eckenmenge in $2r - 2$ monochromatische verallgemeinerte Wege gibt.

Unsere Techniken liefern zudem ein konzeptuell simpleres Schlussargument für Soukups Beweis von Satz 2.1.

Die Schlüsselidee für unsere Beweise besteht darin, Soukups Techniken mit unseren X -robusten Wegen zu verstärken.

II. Enden von Digraphen

In diesem Teil der Dissertation entwickeln wir eine Endenraumtheorie für Digraphen. Genau wie bei Enden von ungerichteten Graphen sind auch Enden von Digraphen Punkte im Unendlichen, gegen die ihre Strahlen konvergieren. Anders als bei Enden von ungerichteten Graphen sind manche Enden durch Limes Kanten verbunden; diese sind essenziell, um den Endenraum eines Digraphen als (inversen) Limes von endlichen Kontraktionsminoren zu erhalten.

Im Hauptresultat von Kapitel 8 zeigen wir, dass sich Richtungen von ungerichteten Graphen, knäuelartige Beschreibungen ihrer Enden, auf Digraphen übertragen lassen: Die Richtungen eines jeden Digraphen korrespondieren bijektiv zu seinen Enden und Limes Kanten (Satz 8.2 und Satz 8.3).

Im Zuge dessen entwickeln wir einige fundamentale Werkzeuge und Techniken für die Untersuchung von Enden in Graphen, wie etwa das Halskettenlemma oder

das gerichtete Stern-Kamm Lemma.

In Kapitel 9 führen wir den topologischen Raum $|D|$ ein, welcher aus D und den Enden und Limeskanten von D geformt wird. Danach charakterisieren wir die Digraphen, welche durch diesen Raum kompaktifiziert werden (Satz 9.1). Des Weiteren zeigen wir, dass der Raum $|D|$, sobald er kompakt ist, der inverse Limes von endlichen Kontraktionsminoren von D ist (Satz 9.2).

Um den typischen Nutzen hiervon zu illustrieren, erweitern wir zwei Aussagen über endliche Digraphen, die nicht wörtlich auf unendliche Digraphen erweitert werden können, auf den Raum $|D|$. Die erste Aussage ist die Charakterisierung von Eulerschen Digraphen durch die Bedingung, dass der in-Grad einer jeden Ecke ihrem aus-Grad entspricht (Satz 9.3). Die zweite Aussage ist die Charakterisierung von stark zusammenhängenden endlichen Digraphen durch die Existenz eines geschlossenen Hamilton Kantenzugs (Satz 9.4).

In Kapitel 10 führen wir ein Konzept von Tiefensuchbäumen in unendlichen Digraphen ein. Die entsprechenden Bäume nennen wir normale Spannarotheszenzen.

Wir zeigen, dass normale Spannarotheszenzen endentreu sind: jedes Ende wird durch genau einen in der Wurzel beginnenden Strahl der normalen Spannarotheszenzen repräsentiert (Satz 10.1). Weiter zeigen wir, dass sich die hierdurch ergebende Bijektion auf einen Homöomorphismus zwischen dem Horizont eines Digraphen und dem Horizont einer jeden normalen Spannarotheszenzen des Digraphen erweitern lässt (Satz 10.2). Zu guter Letzt beweisen wir ein Jung-ähnliches Kriterium für die Existenz von normalen Spannarotheszenzen (Satz 10.3).

III. Sterne und Kämmen

In den ersten vier Kapiteln dieses Teils der Dissertation ermitteln wir Strukturen, deren Existenz komplementär zu der Existenz von zwei Teilstrukturen ist, die besonders fundamental für die Untersuchung vom Zusammenhang unendlicher Graphen sind: Sterne und Kämmen.

In Kapitel 11 ermitteln wir Strukturen, deren Existenz komplementär zur Existenz von beliebigen Sternen und Kämmen ist. Wir liefern zahlreiche Dualitätsergebnisse: in Hinsicht auf normale Spannarotheszenzen (Satz 11.1 und Satz 11.6), Baumzerlegungen (Satz 11.2 und Satz 11.7), kritische Eckenmengen (Satz 11.3) und

dem U -Rang (Satz 11.4).

In Kapitel 12 ermitteln wir Strukturen, deren Existenz komplementär zur Existenz von dominierenden Sternen und dominierten Kämmen ist. Da dominierende Sterne genau dann existieren, wenn dominierte Kämmen existieren, stimmen ihre komplementären Strukturen überein. Wie bei beliebigen Sternen und Kämmen, sind unsere Dualitätssätze für dominierte Kämmen (und dominierende Sterne) hinsichtlich normaler Bäumen (Satz 12.1) und Baumzerlegungen (Satz 12.2.5) formuliert.

Die komplementären Strukturen, die wir für dominierte Kämmen bereitstellen, vereinen jene für Sterne und Kämmen und erlauben uns die Dualitätssätze für Sterne und Kämmen von jenen für dominierte Kämmen abzuleiten. Das ist überraschend, da unsere Strukturen für Sterne und Kämmen recht unterschiedlich sind: jene für Sterne sind lokal endlich, während jene für Kämmen strahlenlos sind.

In Kapitel 13 ermitteln wir Strukturen, deren Existenz komplementär zur Existenz von undominierten Kämmen ist. Wir beschreiben ihre komplementären Strukturen hinsichtlich strahlenloser Bäume (Satz 13.1) und Baumzerlegungen (Satz 13.7).

Anwendungen umfassen eine vollständige Charakterisierung (hinsichtlich normaler Spannbäume) der Graphen, deren Strahlen dominiert sind, aber die keinen strahlenlosen Spannbaum haben (Korollar 13.1.3). Nur zwei solche Graphen wurden bisher konstruiert (von Seymour und Thomas [65] und Thomassen [70]). Als Korollar zeigen wir, dass Graphen mit normalem Spannbaum genau dann einen strahlenlosen Spannbaum haben, wenn alle ihre Strahlen dominiert sind (Korollar 13.2).

Eine weitere Anwendung löst ein Problem, das Carmesin [15] offen ließ: Die Graphen, deren undominierte Enden durch einen geeigneten Spannbaum reflektiert werden, können hinsichtlich normaler Spannbäume charakterisiert werden (Theorem 13.3). Insbesondere zeigen wir, dass jeder Graph, der einen normalen Spannbaum hat, auch einen Spannbaum hat, der die undominierten Enden reflektiert.

In Kapitel 14 ermitteln wir Strukturen, deren Existenz dual zu der Existenz von undominierenden Sternen ist. Das schließt unsere Lösung des Problems komplementäre Strukturen von Sternen, Kämmen und deren Kombinationen zu finden ab. Unsere Sätze sind hinsichtlich Enden-kompaktifizierender Teilgraphen

(Satz 14.1) und Baumzerlegungen (Satz 14.2) formuliert.

Das letzte Kapitel dieses Teils, Kapitel 15, kann als Spinn-off der vorherigen Kapitel angesehen werden: wir betrachten zwei Probleme, die Halin stellte, und stellen weitere Anwendungen unserer Dualitätssätze und der Werkzeuge, die wir für diese entwickelt haben, vor.

Schmidt charakterisierte die Klasse der strahlenlosen Graphen durch eine ordinale Rangfunktion, welche es ermöglicht, Aussagen über strahlenlose Graphen mit transfiniten Induktion zu beweisen. Halin fragte, ob sich Schmidts Rangfunktion so verallgemeinern lässt, dass weitere wichtige Graphenklassen charakterisiert werden können. Wir beantworteten Halins Frage positiv: wir charakterisieren zwei wichtige Graphenklassen durch eine ordinale Rangfunktion.

Seymour und Thomas charakterisierten für jede überabzählbare Kardinalität κ die Klasse der Graphen ohne einen T_κ Minor. Wir erweitern ihr Resultat um eine Charakterisierung hinsichtlich einer ordinalen Rangfunktion (eine für jede überabzählbare Kardinalität κ , Satz 15.1).

Ein weiteres, weitestgehend offenes, Problem von Halin fragt nach einer Charakterisierung der Graphen mit endentremem Spannbaum. Eine gut untersuchte Teilklasse ist die Klasse der normal aufgespannten Graphen. Wir bestimmen eine größere Teilklasse, die Klasse der normal verfolgbar Graphen, welche aus den zusammenhängenden Graphen besteht, die eine strahlenlose Baumzerlegung in normal aufgespannte Teile haben. Die Klasse der normal verfolgbar Graphen untersuchen wir noch weiter und beweisen, dass alle ihre Graphen einen Spannbaum haben der die undomierten Enden reflektiert (Satz 15.2). Unsere Beweise basieren auf einer Charakterisierung der Klasse der normal verfolgbar Graphen durch eine ordinale Rangfunktion, die wir bereitstellen (Satz 15.5).

Publications related to this dissertation

The following articles are related to this dissertation:

Part I:

- (1) Part I is based on [13].

Part II:

- (2) Chapter 8 is based on [10].
- (3) Chapter 9 is based on [11].
- (4) Chapter 10 is based on [12].

Part III:

- (5) Chapter 11 is based on [5].
- (6) Chapter 12 is based on [6].
- (7) Chapter 13 is based on [7].
- (8) Chapter 14 is based on [8].
- (9) Chapter 15 is based on [9].

Declaration on my contributions

The research conducted in this thesis is based on collaborative work with three authors. The corresponding articles are listed on the previous page. In the following I will point out my contributions to this research.

I. Monochromatic generalised paths

This part is a generalisation of my master's thesis, where I proved the \aleph_1 case of Soukup's conjecture. However our proof strategy of the general case considerably differs.

Pitz—who also supervised my Master's thesis—contributed ideas and asked critical questions. I came up with the three main ingredients explained in Chapter 4.0.2 and figured out how to combine them to gain a proof of Soukup's conjecture. Pitz wrote the first draft of the introduction and of Chapter 7. I first drafted Chapters 4, 5 and 6.

II. Ends of digraphs

This part is based on a series of three papers [10–12] that I wrote together with Melcher. The research was conducted in close collaboration and we share an equal amount of work: both in developing the main ideas and in drafting the papers.

The project started with a new approach to ends of digraphs. I told Melcher about my ideas and invited him to join the project.

I first drafted Chapter 8 and Melcher wrote the first draft of Chapter 9. We drafted the introduction and preliminaries of Chapter 10 together. I first drafted Sections 10.2, 10.3 and 10.5 and Melcher first drafted Section 10.4.

III. Stars and combs

This part is based on five papers [5–9] that I wrote together with Kurkofka. The research was conducted in close collaboration and we share an equal amount of work: both in developing the main ideas and in drafting the papers.

Kurkofka started the project and partially drafted Chapter 11. He presented his draft to me and I had some suggestion to extend it—so he invited me to join the project. Some days later, I told Kurkofka about my ideas to consider combinations of stars and combs. In one and a half years of close collaboration we further developed our theory, including our results in the last chapter of this part (which can be seen as a spin-off of the previous ones).

Kurkofka first drafted Chapter 11 and Chapter 14 and I wrote the first draft of Chapter 12 and Chapter 13. We drafted the introduction and the preliminaries of Chapter 15 together. Kurkofka first drafted Section 15.1 and Section 15.2 and I wrote the first draft of Section 15.3 and of Section 15.4.

Acknowledgement

First, I would like to thank my supervisor Prof. Reinhard Diestel. You fascinated me for graph theory and inspired me a lot—not only in the last three years, but also during my bachelor and master. As my supervisor you gave me a lot of mathematical advice and you also stood up for me several times.

Second, I would like to thank Dr. Max Pitz. Your mind set and positive way of thinking had a great impact on my personal development. You always had a sympathetic ear and gave me a lot of constructive feedback. I really enjoyed working with you and learning from you.

Next, I would like to thank my colleagues and friends. Thank you Jan Kurkofka and Ruben Melcher. We had a lot of fun during our research and learned a lot from each other. I will never forget the great time that we spent together. Thank you Fabiana Bornholdt for the wonderful time outside of working hours. You gave me the necessary balance that I needed.

Finally, many thanks to you mum and dad for all your love. You have supported me all the way and were always there for me.

Bibliography

- [1] J. Bang-Jensen and G. Gutin, *Digraphs: theory, algorithms and applications*, Springer Science & Business Media, 2008. ↑33, 61
- [2] H. Bruhn and R. Diestel, *Duality in infinite graphs*, *Comb., Probab. & Comput.* **15** (2006), 75–90. ↑2, 32, 59, 182, 200, 202
- [3] H. Bruhn, R. Diestel, A. Georgakopoulos, and P. Sprüssel, *Every rayless graph has an unfriendly partition*, *Combinatorica* **30** (2010), no. 5, 521–532. ↑36, 111, 131, 229
- [4] C. Bürger, *Decomposing Edge-Coloured Infinite Graphs into Monochromatic Paths and Cycles*, Master’s Thesis, 2017. ↑8, 9
- [5] C. Bürger and J. Kurkofka, *Duality theorems for stars and combs I: Arbitrary stars and combs*, 2020. available at arXiv:2004.00594. ↑36, 44, 251, 253
- [6] ———, *Duality theorems for stars and combs II: Dominating stars and dominated combs*, 2020. arXiv:2004.00593. ↑126, 251, 253
- [7] ———, *Duality theorems for stars and combs III: Undominated combs*, 2020. arXiv:2004.00592. ↑231, 240, 251, 253
- [8] ———, *Duality theorems for stars and combs IV: Undominating stars*, 2020. arXiv:2004.00591. ↑251, 253
- [9] ———, *End-faithful spanning trees in graphs without normal spanning trees*, 2020. arXiv:2004.00591. ↑251, 253
- [10] C. Bürger and R. Melcher, *Ends of digraphs I: Basic theory*, 2020. In preparation. ↑69, 251, 252
- [11] ———, *Ends of digraphs II: The topological point of view*, 2020. In preparation. ↑98, 251, 252
- [12] ———, *Ends of digraphs III: Normal arborescences*, 2020. In preparation. ↑251, 252
- [13] C. Bürger and M. Pitz, *Partitioning edge-coloured infinite complete bipartite graphs into monochromatic paths*, *Israel Journal of Mathematics* (2020). <https://doi.org/10.1007/s11856-020-2030-z>. ↑251
- [14] C. Carathéodory, *Über die Begrenzung einfach zusammenhängender Gebiete*, *Math. Annalen* **73** (1913), no. 3, 323–370. ↑32

- [15] J. Carmesin, *All graphs have tree-decompositions displaying their topological ends*, *Combinatorica* **39** (2019), no. 3, 545–596. ↑117, 125, 142, 150, 160, 165, 181, 195, 222, 230, 245, 249
- [16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*, MIT press, 2009. ↑87
- [17] R. Diestel, *The end structure of a graph: recent results and open problems*, *Discrete Mathematics* **100** (1992), 313–327. ↑97, 125
- [18] ———, *Graph Theory*, 3rd ed., Springer, 2005. ↑2, 32, 59
- [19] ———, *Locally finite graphs with ends: a topological approach*, *Discrete Mathematics* **310–312** (2010), 2750–2765 (310); 1423–1447 (311); 21–29 (312). arXiv:0912.4213v3. ↑118
- [20] ———, *Graph Theory*, 5th ed., Springer, 2016. ↑2, 5, 14, 32, 33, 36, 42, 44, 50, 59, 61, 66, 80, 86, 87, 88, 90, 96, 107, 110, 112, 114, 117, 118, 119, 120, 123, 126, 131, 132, 135, 149, 157, 173, 174, 179, 181, 182, 206, 229, 230, 231, 232
- [21] ———, *Ends and Tangles*, *Abh. Math. Sem. Univ. Hamburg* **87** (2017), no. 2, 223–244. Special issue in memory of Rudolf Halin, arXiv:1510.04050v3. ↑111, 119, 208
- [22] ———, *Abstract Separation Systems*, *Order* **35** (2018), no. 1, 157–170. arXiv:1406.3797v6. ↑208
- [23] ———, *Tree Sets*, *Order* **35** (2018), no. 1, 171–192. arXiv:1512.03781v3. ↑212
- [24] ———, *End spaces and spanning trees*, *J. Combin. Theory (Series B)* **96** (2006), no. 6, 846–854. ↑60, 118, 149, 153, 157, 158, 207, 209
- [25] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, *J. Combin. Theory (Series B)* **87** (2003), 197–206. ↑33, 50
- [26] R. Diestel and D. Kühn, *On Infinite Cycles I*, *Combinatorica* **24** (2004), 68–89. ↑32, 59
- [27] ———, *Graph-theoretical versus topological ends of graphs*, *J. Combin. Theory (Series B)* **87** (2003), 197–206. ↑115, 117, 118, 145, 233
- [28] ———, *Topological Paths, Cycles and Spanning Trees in Infinite Graphs*, *Europ. J. Comb.* **25** (2004), 835–862. ↑182
- [29] R. Diestel and I. Leader, *Normal spanning trees, Aronszajn trees and excluded minors*, *J. London Math. Soc.* **63** (2001), 16–32. ↑127, 181, 190, 230, 235
- [30] R. Diestel, *Locally finite graphs with ends: A topological approach, I. Basic theory*, *Discrete Mathematics* **311** (2011), no. 15, 1423–1447. ↑59
- [31] A. K. Elm and J. Kurkofka, *A tree-of-tangles theorem for infinite-order tangles*, 2020. arXiv:2003.02535. ↑139, 208, 211, 212, 213, 218
- [32] R. Engelking, *General Topology*, second, Sigma Series in Pure Mathematics, vol. 6, Heldermann Verlag, Berlin, 1989. ↑66, 117, 118

- [33] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, Mathematische Zeitschrift **33** (1931), no. 1, 692–713. ↑32
- [34] J. Geelen and B. Joeris, *A generalization of the Grid Theorem*, 2016. arXiv:1609.09098v1. ↑114
- [35] A. Georgakopoulos, *Infinite Hamilton cycles in squares of locally finite graphs*, Advances in Mathematics **220** (2009), no. 3, 670–705. ↑2, 32, 59
- [36] J. P. Gollin and K. Heuer, *Characterising k -connected sets in infinite graphs*, 2018. arXiv:1811.06411. ↑114, 115
- [37] R. Halin, *Über unendliche Wege in Graphen*, Math. Annalen **157** (1964), 125–137. ↑4, 32, 116, 125, 195, 229
- [38] ———, *Charakterisierung der Graphen ohne unendliche Wege*, Arch. Math. **16** (1965), no. 1, 227–231. ↑111, 130
- [39] ———, *Graphen ohne unendliche Wege*, Math. Nachr. **31** (1966), no. 1–2, 111–123. ↑111
- [40] ———, *Simplicial decompositions of infinite graphs*, Advances in graph theory, annals of discrete mathematics, 1978. ↑114, 153, 202
- [41] ———, *Miscellaneous problems on infinite graphs*, J. Graph Theory **35** (2000), 128–151. ↑198, 229, 246
- [42] B. Joeris, *Connectivity, tree-decompositions and unavoidable-minors*, Ph.D. Thesis, 2015. ↑114
- [43] H. A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, Math. Nachr. **41** (1969), 1–22. ↑34, 90, 110, 126, 127, 153, 181
- [44] J. Kneip and P. Gollin, *Representations of infinite tree sets*, 2020. arXiv:1908.10327. ↑218
- [45] P. Komjáth, *Martin’s axiom and spanning trees of infinite graphs*, J. Combin. Theory (Series B) **56** (1992), no. 1, 141–144. ↑184
- [46] K. Kunen, *Set Theory*, Studies in logic, College Publications, 2011. ↑115
- [47] J. Kurkofka and M. Pitz, *Tangles and the Stone-Čech compactification of infinite graphs*, 2018. arXiv:1806.00220v1. ↑111, 119, 208
- [48] ———, *Ends, tangles and critical vertex sets*, Math. Nachr. **292** (2019), no. 9, 2072–2091. arXiv:1804.00588v1. ↑111, 119, 208, 210
- [49] A. G. L. Gerencsér, *On Ramsey-type problems*, Ann. Univ. Sci. Budapest. Eötvös Sect. Math **10** (1967), 167–170. ↑8
- [50] P. Z. L. Ribes, *Profinite Groups*, 2nd ed., Springer, 2010. ↑65, 66
- [51] L. S. M. Elekes D. Soukup and Z. Szentmiklóssy, *Decompositions of edge-coloured infinite complete graphs into monochromatic paths*, Discrete Mathematics **340** (2017), no. 8, 2053–2069. ↑7, 8

- [52] V. Neumann-Lara, *The dichromatic number of a digraph*, J. Combin. Theory, Series B **33** (1982), no. 3, 265–270. ↑47, 95
- [53] B. Oporowski, J. Oxley, and R. Thomas, *Typical subgraphs of 3- and 4-connected graphs*, J. Combin. Theory (Series B) **57** (1993), no. 2, 239–257. ↑114
- [54] M. Pitz, *A new obstruction for normal spanning trees*, 2020. arXiv:2005.04150. ↑230
- [55] A. Pokrovskiy, *Partitioning edge-coloured complete graphs into monochromatic cycles and paths*, J. Combin. Theory, Series B **106** (2014), 70–97. ↑8
- [56] N. Polat, *Développments terminaux des graphes infinis III. Arbres maximaux sans rayon, cardinalité maximum des ensembles disjoints de rayons*, Math. Nachr. **115** (1984), no. 1, 337–352. ↑180, 185
- [57] ———, *Topological aspects of infinite graphs*, Cycles and rays: basic structures in finite and infinite graphs, Proc. NATO Adv. Res. Workshop, Montreal/Can. 1987, NATO ASI Ser. (1990), 197–220. ↑180, 185
- [58] ———, *Ends and multi-endings, I*, J. Combin. Theory (Series B) **67** (1996), 86–110. ↑32
- [59] ———, *Ends and multi-endings, II*, J. Combin. Theory (Series B) **68** (1996), 56–86. ↑32
- [60] ———, *End-faithful spanning trees in T_{\aleph_1} -free graphs*, J. Graph Theory **26** (1997), no. 4, 175–181. ↑185, 230, 231
- [61] R. Rado, *Monochromatic paths in graphs*, Annals of discrete mathematics, 1978, pp. 191–194. ↑1, 7, 18
- [62] L. Ribes and P. Zalesskii, *Profinite Groups*, Springer, 2010. ↑117
- [63] N. Robertson and P. D. Seymour, *Graph Minors. X. Obstructions to Tree-Decomposition*, J. Combin. Theory (Series B) **52** (1991), 153–190. ↑111, 208
- [64] R. Schmidt, *Ein Ordnungsbegriff für Graphen ohne unendliche Wege mit einer Anwendung auf n -fach zusammenhängende Graphen*, Arch. Math. **40** (1983), no. 1, 283–288. ↑33, 36, 44, 107, 131, 229, 232, 246
- [65] P. Seymour and R. Thomas, *An end-faithful spanning tree counterexample*, Discrete Mathematics **95** (1991), no. 1, 321–330. ↑180, 184, 229, 245, 249
- [66] P. D. Seymour and R. Thomas, *Excluding infinite trees*, Trans. Amer. Math. Soc. **335** (1993), 597–630. ↑229, 230, 232, 233
- [67] D. Soukup, *Colouring problems of Erdős and Rado on infinite graphs*, Ph.D. Thesis, 2015. ↑8, 15
- [68] ———, *Decompositions of edge-coloured infinite complete graphs into monochromatic paths ii*, Israel Journal of Mathematics **221** (2017), no. 1, 235–273. ↑1, 8, 9, 11, 12, 13, 17, 19, 20, 27, 28, 29, 30

- [69] C. Thomassen, *Duality of Infinite Graphs*, J. Combin. Theory (Series B) **33** (1982), 137–160. ↑202
- [70] ———, *Infinite connected graphs with no end-preserving spanning trees*, J. Combin. Theory (Series B) **54** (1992), no. 2, 322–324. ↑180, 184, 229, 245, 249
- [71] H. Whitney, *Non-separable and planar graphs*, Trans. Am. Math. Soc. **34** (1932), 339–362. ↑182
- [72] J. Zuther, *Planar strips and an end concept for digraphs*, Ph.D. Thesis, 1996. ↑2, 32, 46
- [73] J. Širáň, *End-faithful forests and spanning trees in infinite graphs*, Discrete Mathematics **95** (1991), no. 1, 331–340. ↑180, 184, 185

Eidesstattliche Versicherung

Hiermit versichere ich an Eides statt, die vorliegende Dissertation selbst verfasst und keine anderen als die angegebenen Hilfsmittel benutzt zu haben. Darüber hinaus versichere ich, dass diese Dissertation nicht in einem früheren Promotionsverfahren eingereicht wurde.