

**Ends and tangles,
stars and combs,
minors and the Farey graph**

Dissertation

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

vorgelegt

im Fachbereich Mathematik

von

Jan Kurkofka

Hamburg

2020

Erstgutachter : Professor Reinhard Diestel (Betreuer)
Zweitgutachter : Professor Nathan Bowler
Drittgutachter : Professor Imre Leader

Datum der Disputation : 17. August 2020

Contents

1. Introduction	1
1.1. Part I: Ends and tangles	1
1.2. Part II: Stars and combs	2
1.3. Part III: End spaces	4
1.4. Part IV: The Farey graph	5
2. Tools and terminology	8
2.1. The star-comb lemma	8
2.2. Inverse limits	8
2.3. Separation systems, S -trees and tree sets	8
2.3.1. Separations of sets and abstract separation systems	9
2.3.2. Orientations	9
2.3.3. S -trees	10
2.3.4. Nested sets of separations and tree sets	10
2.3.5. Separations of graphs	11
2.4. Ends of graphs	11
2.4.1. Definition and notation	11
2.4.2. End spaces	12
2.4.3. Directions	12
I. Ends and tangles	14
3. Tangles and the Stone-Čech compactification of infinite graphs	15
3.1. Introduction	15
3.2. Reviewing Diestel's tangle compactification	16
3.2.1. The 1-complex of a graph	16
3.2.2. Tangles	16
3.2.3. Ends and Tangles	17
3.3. Tangles as inverse limit of finite spaces	18
3.4. Background on the Stone-Čech compactification of an infinite graph	19
3.4.1. Stone-Čech compactification of 1-complexes	19
3.4.2. Three examples	23
3.5. Comparing the Stone-Čech remainder with the tangle space	26
3.5.1. The Stone-Čech remainder of the vertex set	26
3.5.2. An auxiliary remainder	26
3.5.3. The components of the remainder can be distinguished by finite separators	27
3.5.4. Comparing \mathcal{P} with G^\times	31
3.6. Obtaining the tangle compactification from the Stone-Čech compactification	33
3.7. Three observations about the Stone-Čech compactification	34

4. A tree-of-tangles theorem for infinite-order tangles	37
4.1. Introduction	37
4.2. Ends, tangles and tree sets	39
4.2.1. Definition of \aleph_0 -tangles	39
4.2.2. Properties of \aleph_0 -tangles	40
4.2.3. Tree sets	41
4.3. Example section	43
4.3.1. Ultrafilters and tree sets	43
4.3.2. The problem case	46
4.4. The overall proof strategy	47
4.5. From principal collections of separators to tree sets	48
4.6. Applications	54
4.6.1. A structural connectivity result for infinite graphs	55
4.6.2. Collectionwise normality of the ultrafilter tangle space	55
4.7. Consistent orientation and lifting from torsos	58
4.8. Extending the tree set of the principal vertex sets	60
4.8.1. Modified torsos, proxies of corridors and lifting from modified torsos	62
4.8.2. Proxies of \aleph_0 -tangles	65
4.8.3. Efficiently distinguishing the proxies	67
4.8.4. Proof of the main result	69
II. Stars and combs	71
5. Duality theorems for stars and combs I: Arbitrary stars and combs	72
5.1. Introduction	72
5.2. Tools and terminology	76
5.2.1. The star-comb lemma	76
5.2.2. Ends of graphs	78
5.2.3. Critical vertex sets	79
5.2.4. Normal trees	80
5.2.5. Containing vertex sets cofinally	81
5.2.6. Tree-decompositions and S -trees	82
5.2.7. Tree-decompositions and S -trees displaying sets of ends	83
5.3. Combs	84
5.3.1. Normal trees	85
5.3.2. Tree-decompositions	86
5.3.3. Critical vertex sets	88
5.3.4. Rank	89
5.3.5. Combining the duality theorems	91
5.4. Stars	97
5.4.1. Normal trees	97
5.4.2. Tree-decompositions	98

6. Duality theorems for stars and combs II: Dominated combs	101
6.1. Introduction	101
6.2. Dominated combs and normal trees	103
6.3. Dominated combs and tree-decompositions	109
6.3.1. A duality theorem related to combs	110
6.3.2. A duality theorem related to stars	112
6.3.3. A duality theorem related to stars and combs	118
7. Duality theorems for stars and combs III: Undominated combs	125
7.1. Introduction	125
7.2. Undominated combs and rayless trees	129
7.2.1. Proof for normally spanned graphs	130
7.2.2. Deducing our duality theorem in terms of rayless trees	133
7.3. Spanning trees reflecting the undominated ends	137
7.4. Duality theorems for undominated combs	143
8. Duality theorems for stars and combs IV: Undominating stars	146
8.1. Introduction	146
8.2. Tough subgraphs	148
8.3. Star-decompositions	153
8.4. Overview of all duality results	160
9. End-faithful spanning trees in graphs without normal spanning trees	163
9.1. Introduction	163
9.2. Tools and terminology	165
9.3. Ranking T_κ -free graphs	166
9.4. Normally traceable graphs	167
9.5. End-faithful spanning trees	169
9.6. Trees reflecting the undominated ends	172
 III. End spaces	 174
10. Approximating infinite graphs by normal trees	175
10.1. Introduction	175
10.2. End spaces of graphs: a reminder	176
10.3. Proof of the main result	178
10.4. Consequences of the approximation result	181
10.5. Paracompactness in subspaces of end spaces	183
11. Uniquely determined ends and graphs	185
11.1. Introduction	185
11.2. Preliminaries	187
11.2.1. Normal trees	187
11.2.2. Tree-decompositions, \mathcal{S} -trees and ends	188

11.3. Countably determined directions and the first axiom of countability	189
11.4. Countably determined graphs and the second axiom of countability	194
IV. The Farey graph	201
12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$	202
12.1. Introduction	202
12.2. Preliminaries	204
12.3. Uniqueness and topological minors	204
12.3.1. Uniqueness	204
12.3.2. Minor versus topological minor	206
12.4. Overall proof strategy	206
12.5. Robustness	207
12.5.1. Finitely separating spanning trees	207
12.5.2. Analysing the components	209
12.5.3. Football minors	211
12.6. Proof of the main result	215
12.7. Outlook	218
13. The Farey graph is uniquely determined by its connectivity	219
13.1. Introduction	219
13.2. Preliminaries	220
13.3. Atypical Π -graphs and variations of the main result	220
13.3.1. Atypical Π -graphs	220
13.3.2. Variations of the main result	220
13.4. Overall proof strategy	221
13.5. Grain lines	222
13.6. Proof of the main result	226
14. Ubiquity and the Farey graph	228
14.1. Introduction	228
14.2. Proof of Theorem 1	230
14.3. Finding the Farey graph in the whirl graph	232
Appendix	237
15. English summary	237
16. Deutsche Zusammenfassung	240
17. Publications related to this dissertation	244
18. Declaration on my contributions	245
Acknowledgement	247

Contents

Bibliography	248
Eidesstattliche Versicherung	253

1. Introduction

“It has been reported that Dénes König, the author of the classic *Theorie der endlichen und unendlichen Graphen* (Leipzig, 1936), expressed a special liking for infinite graphs, which certainly receive substantial attention in his book. Nevertheless, the majority of combinatorialists seem to have concentrated on finite combinatorics, to the extent that it has almost seemed an eccentricity to think that graphs and other combinatorial structures can be either finite or infinite. However, there seems to be no logical reason why combinatorial structures should ‘usually’ be finite, and indeed this would preclude many fascinating avenues of exploration.” (C.St.J.A. Nash-Williams).

In this dissertation we will explore four exciting branches of structural infinite graph theory. Each of these offers individual problems, some of which we will solve. The techniques that we will use in our proofs span the whole breadth of non-set-theoretic infinite graph theory and include tools from general topology. Some of these techniques are new and we shall develop them in this dissertation. The four branches that we will explore are:

- (i) Ends and tangles
- (ii) Stars and combs
- (iii) End spaces
- (iv) The Farey graph

This dissertation consists of four parts, one for each branch, and all parts consist of several chapters, one for every problem in the branch that we solve. Here, then is a brief overview of the parts and their chapters, using the terminology from [26]. We remark, however, that every chapter will feature its own independent and more comprehensive introduction.

1.1. Part I: Ends and tangles

1.1.1. Chapter 3: Tangles and the Stone-Čech compactification of infinite graphs

Every locally finite connected graph can be naturally compactified by its ends to form its well-known end compactification, see e.g. [26, §8.6]. For graphs that are not locally finite, however, adding their ends no longer suffices to compactify them, and it has been a longstanding quest to decide what other ‘points at infinity’ besides the ends should be added to obtain a compactification, see e.g. Cartwright, Soardi and Woess [21] and Polat [71].

Recently, Diestel [25] proposed a solution to this problem: he generalised the end compactification from locally finite connected graphs to arbitrary graphs by generalising ends to infinite-order tangles, yielding the *tangle compactification*. Diestel then asked how the tangle compactification of an infinite graph relates to its Stone-Čech compactification [25, §6]. Indeed, it is well-known that the end

1. Introduction

compactification of a locally finite connected graph G can be described naturally in terms of its Stone-Čech compactification, namely, it is the quotient obtained by collapsing each connected component of the Stone-Čech remainder to a single point, see e.g. [1, §VI.3]. As our main result in this chapter, we show that this correspondence extends to all graphs when ends are generalised to tangles; see Theorem 3.1. Hence, even though Diestel’s reasoning and motivation behind the tangle compactification was purely combinatorial, it naturally happens to generalise the end compactification also in this second, more topological aspect.

1.1.2. Chapter 4: A tree-of-tangles theorem for infinite-order tangles

The *tree-of-tangles theorem*, one of the cornerstones of Robertson and Seymour’s proof of their graph-minor theorem, says:

Theorem. *Every finite graph G has a nested set of separations which efficiently distinguishes all the finite-order tangles in G that can be distinguished.*

Recently, Carmesin [19] extended the tree-of-tangles theorem to the infinite-order tangles of infinite graphs that are locally finite. Notably, Carmesin used his result to prove a conjecture of Halin [47] (in amended form) that had remained open for 50 years, and a conjecture of Diestel [28] (also in amended form) that had remained open for 20 years.

As our main result in this chapter, we extend Robertson and Seymour’s tree-of-tangles theorem to the infinite-order tangles of infinite graphs (and thus, we extend Carmesin’s result from locally finite graphs to all graphs); see Theorem 4.1. For our proof we develop a new concept called ‘corridors’, which we will use once more in Chapter 8 of the next part.

Our result has four applications: one connectivity result, one topological result about tangles, one application in the work of Elbracht, Kneip and Teegen [34], and the final application is the foundation of Chapter 8; see the comprehensive introduction of this chapter.

1.2. Part II: Stars and combs

1.2.1. Chapters 5–8: Duality theorems for stars and combs

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. The *star-comb lemma*, a standard tool in infinite graph theory [26], says:

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 .*

1. Introduction

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. These two possibilities apply separately to all their infinite sets U of vertices, and clearly, the smaller U the stronger the assertion.

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). Hence it is natural to ask for structures, more specific than combs and stars attached to U , whose existence is complementary to that of stars and combs attached to U , respectively.

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 the remaining three chapters of this series 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 four chapters in our series, each naming the substructure for which duality theorems are proved in its title:

- I: ARBITRARY STARS AND COMBS (Chapter 5)
- II: DOMINATING STARS AND DOMINATED COMBS (Chapter 6)
- III: UNDOMINATED COMBS (Chapter 7)
- IV: UNDOMINATING STARS (Chapter 8)

Our duality theorems will usually be phrased in terms of normal trees and tree-decompositions. We remark that the vast number of techniques that we will use in this series of four chapters already spans the whole breadth of non-set-theoretic infinite graph theory.

1.2.2. Chapter 9: End-faithful spanning trees in graphs without normal spanning trees

Schmidt [26, 78] characterised the class of rayless graphs by an ordinal rank function, which makes it possible to prove statements about rayless graphs by transfinite induction. At the turn of the millennium, Halin [44] 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

1. Introduction

answer Halin’s question in the affirmative: we characterise two important classes of graphs by an ordinal rank function.

As our first main result in this chapter, we characterise for every uncountable cardinal κ the class of graphs without a T_κ minor by an ordinal rank function (recall that T_κ denotes the κ -branching tree); see Theorem 9.1. This extends Seymour and Thomas’ characterisations [77].

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 natural 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 [47] conjectured that every connected graph has an end-faithful spanning tree. However, Seymour and Thomas [76] and Thomassen [83] constructed counterexamples. Ever since, it has been an open problem to characterise the class of graphs that admit an end-faithful spanning tree. A well-studied subclass is formed by the graphs with a normal spanning tree. In this chapter, 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; see Theorem 9.2. This subclass includes all other subclasses that have been structurally characterised. Our proof of Theorem 9.2 relies on a characterisation of the class of normally traceable graphs by an ordinal rank function that we provide; see Theorem 9.5.

1.3. Part III: End spaces

1.3.1. Chapter 10: Approximating infinite graphs by normal trees

Normal spanning trees are perhaps the most useful structural tool in infinite graph theory. Their importance arises from the fact that they capture the separation properties of the graph they span, and so in many situations it suffices to deal with the much simpler tree structure instead of the whole graph. For example, the end space of G coincides, even topologically, with the end space of any normal spanning tree of G . However, not every connected graph has a normal spanning tree, and the structure of graphs without normal spanning trees is still not completely understood [11, 32].

In order to harness and transfer the power of normal spanning trees to arbitrary connected graphs, we show that every connected graph can be approximated by a normal tree, up to some arbitrarily small error phrased in terms of neighbourhoods around its ends; see Theorem 10.1. The existence of such approximate normal trees has consequences of both combinatorial and topological nature.

On the combinatorial side, we show that a graph has a normal spanning tree as soon as it has normal spanning trees locally at each end; i.e., the only obstruction for a graph to having a normal spanning tree is an end for which none of its neighbourhoods has a normal spanning tree.

On the topological side, we show that end spaces of graphs are always paracompact, which gives unified and short proofs for a number of results by Diestel [24],

Sprüssel [80] and Polat [68], and answers an open question about metrizable of end spaces by Polat; see the comprehensive introduction.

1.3.2. Chapter 11: Countably determined ends and graphs

The *directions* of an infinite graph G are a tangle-like description of its ends: they are choice functions that choose compatibly for all finite vertex sets $X \subseteq V(G)$ a component of $G - X$.

Although every direction is induced by a ray, there exist directions of graphs that are not uniquely determined by any countable subset of their choices. We characterise these directions and their countably determined counterparts in terms of star-like substructures or rays of the graph; see Theorems 11.1 and 11.2.

Curiously, there exist graphs whose directions are all countably determined but which cannot be distinguished all at once by countably many choices. We structurally characterise the graphs whose directions can be distinguished all at once by countably many choices, and we structurally characterise the graphs which admit no such countably many choices; see Theorems 11.3 and 11.4. Our characterisations are phrased in terms of normal trees and tree-decompositions.

Our four (sub)structural characterisations imply combinatorial characterisations of the four classes of infinite graphs that are defined by the first and second axiom of countability applied to their end spaces: the two classes of graphs whose end spaces are first countable or second countable, respectively, and the complements of these two classes.

1.4. Part IV: The Farey graph




Figure 1.4.1.: The Farey graph




Figure 1.4.2.: The graph $T_{\aleph_0} * t$

The Farey graph, shown in Figure 1.4.1 and surveyed in [22, 49], plays a role in a number of mathematical fields ranging from group theory and number theory to geometry and dynamics [22]. Curiously, graph theory is not among these.

1.4.1. Chapter 12: Every infinitely edge-connected graph contains the Farey graph or $T_{\aleph_0} * t$ as a minor

In this chapter we show that the Farey graph plays a central role in graph theory too: it is one of two infinitely edge-connected graphs that must occur as a minor in every infinitely edge-connected graph; see Theorem 12.1. The other graph is $T_{\aleph_0} * t$, the graph obtained from the infinitely-branching tree T_{\aleph_0} by joining an additional vertex t to all its vertices; see Figure 1.4.2. Previously it was not known that there was any set of graphs determining infinite edge-connectivity by forming a minor-minimal list in this way, let alone a finite set.

Since both the Farey graph and $T_{\aleph_0} * t$ are planar, our result implies that every infinitely edge-connected graph contains a planar infinitely edge-connected graph as a minor. Thus, in this sense, infinite edge-connectivity is an inherently planar property.

1.4.2. Chapter 13: The Farey graph is uniquely determined by its connectivity

In the previous chapter we showed that the Farey graph is one of two infinitely edge-connected graphs that must occur as a minor in every infinitely edge-connected graph. Infinite edge-connectivity, however, is only one aspect of the connectivity of the Farey graph, and it contrasts with a second aspect: the Farey graph does not contain infinitely many independent paths between any two of its vertices. In this chapter we show that the Farey graph is uniquely determined by these two contrasting aspects of its connectivity: up to minor-equivalence, the Farey graph is the unique minor-minimal graph that is infinitely edge-connected but such that every two vertices can be finitely separated; see Theorem 13.1. This is the first graph-theoretic characterisation of the Farey graph.

1.4.3. Chapter 14: Ubiquity and the Farey graph




Figure 1.4.3.: The whirl graph, coloured

Let us call two $u-v$ paths *order-compatible* if they traverse their common vertices in the same order. In this chapter we construct the whirl graph shown in Figure 1.4.3

1. Introduction

and show that, for all pairs of vertices u and v , the whirl graph contains k edge-disjoint order-compatible u - v paths for every integer k , but not infinitely many; see Theorem 14.1.

But what does this have to do with the Farey graph? Everything! We shall use the Cantor set to show that the whirl graph contains the Farey graph as a minor with branch sets of size two, but that it contains neither the Farey graph nor $T_{\aleph_0} * t$ as a topological minor. This property makes the whirl graph the ideal example to show that the main results of the two previous chapters are both best possible; see the comprehensive introduction of this chapter.

2. Tools and terminology

Any graph-theoretic notation not explained here can be found in Diestel's textbook [26]. A non-trivial path P is an A -path for a set A of vertices if P has its endvertices but no inner vertex in A .

2.1. The star-comb lemma

The star-comb lemma [26, Lemma 8.2.2] is a standard tool in infinite graph theory. 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 .*

2.2. Inverse limits

A partially ordered set (I, \leq) is said to be *directed* if for every two $i, j \in I$ there is some $k \in I$ with $k \geq i, j$. Let $(X_i \mid i \in I)$ be a family of topological spaces indexed by some directed poset (I, \leq) . Furthermore, suppose that we have a family $(\varphi_{ji}: X_j \rightarrow X_i)_{i \leq j \in I}$ of continuous mappings which are the identity on X_i in case of $i = j$ and which are *compatible* in that $\varphi_{ki} = \varphi_{ji} \circ \varphi_{kj}$ for all $i \leq j \leq k$. Then both families together are said to form an *inverse system*, and the maps φ_{ji} are called its *bonding maps*. We denote such an inverse system by $\{X_i, \varphi_{ji}, I\}$ or $\{X_i, \varphi_{ji}\}$ for short if I is clear from context. Its *inverse limit* $\varprojlim X_i = \varprojlim (X_i \mid i \in I)$ is the topological space

$$\varprojlim X_i = \{ (x_i)_{i \in I} \mid \varphi_{ji}(x_j) = x_i \text{ for all } i \leq j \} \subseteq \prod_{i \in I} X_i.$$

Whenever we define an inverse system without specifying a topology for the spaces X_i first, we tacitly assume them to carry the discrete topology. If each X_i is (non-empty) compact Hausdorff, then so is $\varprojlim X_i$. In particular, $\varprojlim X_i$ is non-empty if all X_i are non-empty and finite. See [36] or [72] for more.

2.3. Separation systems, S -trees and tree sets

Separation systems, S -trees and tree sets are standard tools in graph minor theory. In this section we briefly recall the definitions from [23, 26, 29] that we need, without detailed explanations: for these we refer to the citations. Tangles will be recalled in the respective chapters; however, we remark that in this dissertation we will only work with infinite-order tangles.

2.3.1. Separations of sets and abstract separation systems

A *separation of a set* V is an unordered pair $\{A, B\}$ such that $A \cup B = V$. The ordered pairs (A, B) and (B, A) are its *orientations*. Then the *oriented separations* of V are the orientations of its separations. The map that sends every oriented separation (A, B) to its *inverse* (B, A) is an involution that reverses the partial ordering

$$(A, B) \leq (C, D) \quad :\Leftrightarrow \quad A \subseteq C \text{ and } B \supseteq D$$

since $(A, B) \leq (C, D)$ is equivalent to $(D, C) \leq (B, A)$.

More generally, a *separation system* is a triple $(\vec{S}, \leq, *)$ where (\vec{S}, \leq) is a partially ordered set and $*$: $\vec{S} \rightarrow \vec{S}$ is an order-reversing involution. We refer to the elements of \vec{S} as *oriented separations*. If an oriented separation is denoted by \vec{s} , then we denote its *inverse* \vec{s}^* as \overleftarrow{s} , and vice versa. That $*$ is *order-reversing* means $\vec{r} \leq \vec{s} \Leftrightarrow \overleftarrow{r} \geq \overleftarrow{s}$ for all $\vec{r}, \vec{s} \in \vec{S}$.

A *separation* is an unordered pair of the form $\{\vec{s}, \overleftarrow{s}\}$, and then denoted by s . Its elements \vec{s} and \overleftarrow{s} are the *orientations* of s . The set of all separations $\{\vec{s}, \overleftarrow{s}\} \subseteq \vec{S}$ is denoted by S . When a separation is introduced as s without specifying its elements first, we use \vec{s} and \overleftarrow{s} (arbitrarily) to refer to these elements. Every subset $S' \subseteq S$ defines a separation system $\vec{S}' := \bigcup S' \subseteq \vec{S}$ with the ordering and involution induced by \vec{S} .

Separations of sets, and their orientations, are an instance of this abstract setup if we identify $\{A, B\}$ with $\{(A, B), (B, A)\}$. Here is another example: The set $\vec{E}(T) := \{(x, y) \mid xy \in E(T)\}$ of all *orientations* (x, y) of the edges $xy = \{x, y\}$ of a tree T forms a separation system with the involution $(x, y) \mapsto (y, x)$ and the natural partial ordering on $\vec{E}(T)$ in which $(x, y) < (u, v)$ if and only if $xy \neq uv$ and the unique $\{x, y\}$ – $\{u, v\}$ path in T is $\overset{x}{\dot{}}yT\overset{u}{\dot{}}v = yTu$.

In the context of a given separation system $(\vec{S}, \leq, *)$, a *star (of separations)* is a subset $\sigma \subseteq \vec{S}$ such that $\vec{r} \leq \overleftarrow{s}$ for all distinct $\vec{r}, \overleftarrow{s} \in \sigma$; see [26, Fig. 12.5.1] for an illustration.¹ If t is a node of a tree T , then the set

$$\vec{F}_t := \{(x, t) \mid xt \in E(T)\}$$

is a star in $\vec{E}(T)$.

2.3.2. Orientations

An *orientation* of a separation system \vec{S} , or of a set S of separations, is a subset $O \subseteq \vec{S}$ such that $|O \cap \{\vec{s}, \overleftarrow{s}\}| = 1$ for all $s \in S$. A *partial orientation* of S is an orientation of a subset of S . A subset $O \subseteq \vec{S}$ is *consistent* if there are no two distinct separations $r, s \in S$ with orientations $\vec{r} < \overleftarrow{s}$ and $\overleftarrow{r}, \vec{s} \in O$.

¹Officially, in [23] a star σ is additionally required to consist only of oriented separations \vec{s} satisfying $\vec{s} \neq \overleftarrow{s}$. In this paper, however, all separations considered will satisfy this condition, which is why we will hide it for the convenience of the reader.

2.3.3. S -trees

An S -tree is a pair (T, α) such that T is a tree and $\alpha: \vec{E}(T) \rightarrow \vec{S}$ propagates the ordering on $\vec{E}(T)$ and commutes with inversion: that $\alpha(\vec{e}) \leq \alpha(\vec{f})$ if $\vec{e} \leq \vec{f} \in \vec{E}(T)$ and $(\alpha(\vec{e}))^* = \alpha(\vec{e})$ for all $\vec{e} \in \vec{E}(T)$; see [26, Fig. 12.5.2] for an illustration. Thus, every node $t \in T$ is associated with a star \vec{F}_t in $\vec{E}(T)$ which α sends to a star $\alpha[\vec{F}_t]$ in \vec{S} . A tree-decomposition (T, \mathcal{V}) , for example, makes T into an S -tree for the set of separations it induces [26, §12.5]. For oriented edges $(x, y) \in \vec{E}(T)$ we will write $\alpha(x, y)$ instead of $\alpha((x, y))$. Note that S -trees are ‘closed under taking minors’: if (T, α) is an S -tree and $T' \preceq T$, then $(T', \alpha \upharpoonright \vec{E}(T'))$ is again an S -tree when we view $E(T')$ as a subset of $E(T)$.

2.3.4. Nested sets of separations and tree sets

Two separations are *nested* if they have comparable orientations. Two oriented separations \vec{r}, \vec{s} are *nested* if r and s are nested. A set, either of separations or of oriented separations, is *nested* if every two of its elements are nested. For example, if T is a tree, then both $E(T)$ and $\vec{E}(T)$ are nested.

To state the definition of a tree set, we need the following definitions. An oriented separation $\vec{r} \in \vec{S}$ is

- (i) *degenerate* if $\vec{r} = \vec{r}$,
- (ii) *trivial* if there is a separation $s \in S$ such that both $\vec{r} < \vec{s}$ and $\vec{r} < \vec{s}$, and
- (iii) *small* if $\vec{r} \leq \vec{r}$.

The only degenerate separation of a set V is (V, V) ; its small separations are precisely the ones of the form (A, V) with $A \subseteq V$. All degenerate and trivial separations are small.

A separation system is

- (i) *essential* if it contains neither degenerate nor trivial elements, and
- (ii) *regular* if it contains no small elements.

If $(\vec{S}, \leq, *)$ is essential or regular, then we also call \vec{S} and S *essential* or *regular*, respectively. Regular implies essential.

A *tree set* is a nested essential separation system. If $(\vec{S}, \leq, *)$ is a tree set, then we also call \vec{S} and S *tree sets*. If T is a tree, then $\vec{E}(T)$ is a tree set, the *edge tree set* of T .

In this dissertation, separations usually will not be small, and hence separation systems usually will be regular. This means that when we define a candidate for a tree set and have to verify that it really is a tree set, it will usually suffice to verify nestedness.

Gollin and Kneip [53] characterised the tree sets that are isomorphic to the edge tree set of a tree. An *isomorphism* between two separation systems is a bijection between their underlying sets that respects both their partial orderings and their involutions. 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*.

2. Tools and terminology

Theorem 2.3.1 ([53, 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.*

2.3.5. Separations of graphs

A *separation of a graph G* is a separation $\{A, B\}$ of the set $V(G)$ (meaning $A \cup B = V(G)$) such that G has no edge ‘jumping’ the *separator* $A \cap B$, meaning that G contains no edge between $A \setminus B$ and $B \setminus A$. Thus, (oriented) separations of graphs are an instance of (oriented) separations of sets. The *order* of $\{A, B\}$ is the cardinal $|A \cap B|$. The set of all finite-order separations of a graph G is denoted by $S_{\aleph_0} = S_{\aleph_0}(G)$.

If (A, B) and (C, D) are two separations of G , then

- (i) $(A, B) \vee (C, D) := (A \cup C, B \cap D)$ is their *supremum*, and
- (ii) $(A, B) \wedge (C, D) := (A \cap C, B \cup D)$ is their *infimum*.

Supremum and infimum satisfy De Morgan’s law: $(\vec{r} \vee \vec{s})^* = \vec{r} \wedge \vec{s}$.

The following non-standard notation often will be useful as an alternative perspective on separations of graphs. 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 with $A, B \subseteq V(G)$ can be written in this way. For the orientations of $\{X, \mathcal{C}\}$ we write

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

If C is a component of $G - X$ we write $\{X, C\}$ instead of $\{X, \{C\}\}$. Similarly, we write (C, X) and (X, C) instead of $(\{C\}, X)$ and $(X, \{C\})$, respectively.

2.4. Ends of graphs

2.4.1. Definition and notation

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 [47], 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

2. Tools and terminology

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 $G[B \setminus A]$.

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 . If M is a vertex set and $\partial_\Omega M$ is empty, then M is *dispersed*.

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 [26] 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).

2.4.2. End spaces

If $X \in \mathcal{X}$ is any finite set of vertices of a graph G and C is any component of $G - X$, then we write $\Omega(X, C)$ for the set of ends ω of G with $C(X, \omega) = C$, and abbreviate $\Omega(X, \omega) := \Omega(X, C(X, \omega))$. And if \mathcal{C} is any collection of components of $G - X$, we write $\Omega(X, \mathcal{C}) := \bigcup \{ \Omega(X, C) \mid C \in \mathcal{C} \}$.

The collection of sets $\Omega(X, C)$ with $X \in \mathcal{X}$ and C a component of $G - X$ form a basis for a topology on Ω .

2.4.3. Directions

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 [30] (Theorem 2.4.1 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 limit description of the directions: note that \mathcal{X} is directed 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 $c_{X', X}: \mathcal{C}_{X'} \rightarrow \mathcal{C}_X$ for $X' \supseteq X$ send each component of $G - X'$ to the component of $G - X$ containing it; then $\{ \mathcal{C}_X, c_{X', X}, \mathcal{X} \}$ is an inverse system whose inverse limit, by construction, consists of the directions.

2. Tools and terminology

Theorem 2.4.1 ([30, 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. $\Omega = \varprojlim \mathcal{C}_X$.*

Note that the bijection $\omega \mapsto f_\omega$ is in fact a homeomorphism between the end space Ω and the inverse limit $\varprojlim \mathcal{C}_X$.

Part I.

Ends and tangles

3. Tangles and the Stone-Čech compactification of infinite graphs

3.1. Introduction

Every locally finite connected graph can be naturally compactified by its ends to form its well-known end compactification, see e.g. [26, §8.6]. For graphs that are not locally finite, however, adding their ends no longer suffices to compactify them, and it has been a longstanding quest to decide what other ‘points at infinity’ besides the ends should be added to obtain a compactification, see e.g. Cartwright, Soardi and Woess [21] and Polat [71].

Recently, Diestel [25] proposed a solution to this problem employing Robertson and Seymour’s notion of a tangle [73], which naturally generalises the end compactification (using the terminology from [26, §12.5]): First, he observed that an end ω of a graph G orients every finite-order separation $\{A, B\}$ of G towards the side that contains a tail from every ray in ω ; and since these orientations for distinct separations are consistent in a number of ways, every end naturally induces an infinite-order tangle of G in this way. Diestel then proceeded to show that, conversely, every infinite-order tangle of a locally finite connected graph G is defined by an end in this way. Thus, if G is locally finite and connected, there is a canonical bijection between its infinite-order tangles and its ends.

Finally, Diestel showed that *every* graph, in particular also the non-locally finite ones, is compactified by its infinite-order tangles in much the same way as the ends of a locally finite connected graph compactify it in its end-compactification. The arising *tangle compactification* coincides with the end compactification if G is locally finite and connected. Hence, for the tangle compactification, it is precisely those infinite-order tangles not corresponding to an end which need to be added as points at infinity besides the ends in order to compactify the graph.

Diestel concludes his paper with the question of how the tangle compactification of an infinite graph relates to its Stone-Čech compactification [25, §6]. Indeed, it is well-known that the end compactification of a locally finite connected graph G can be described naturally in terms of its Stone-Čech compactification, namely, it is the quotient obtained by collapsing each connected component of the Stone-Čech remainder to a single point, see e.g. [1, §VI.3]. As our main result, we show that this correspondence extends to all graphs when ends are generalised to tangles. Hence, even though Diestel’s reasoning and motivation behind the tangle compactification was purely combinatorial, it naturally happens to generalise the end compactification also in this second, more topological aspect.

Theorem 3.1. *The tangle compactification of any graph G is obtained from its Stone-Čech compactification βG by first declaring G to be open¹ in βG and then collapsing each connected component of the Stone-Čech remainder to a single point.*

¹When G is locally compact, it is automatically open in βG , and so this step is redundant for locally finite graphs.

3. Tangles and the Stone-Ćech compactification of infinite graphs

This chapter is organised as follows: First, in Section 3.2 we recall graph-theoretic background and provide a brief summary of Diestel’s tangle compactification of an infinite graph. In Section 3.3, we describe the remainder of the tangle compactification as an inverse limit of finite discrete spaces. In Section 3.4, we provide the necessary background on the Stone-Ćech compactification, and explain how the quotient relation defining the 1-complex G can be used to describe the Stone-Ćech compactification of an infinite graph as a ‘fake 1-complex’ on standard intervals and *non-standard* intervals (where the non-standard intervals are the standard subcontinua of the remainder of the positive half-line). Sections 3.5 and 3.6 contain the proof of our main theorem. We conclude this chapter in Section 3.7 with three additional observations about the tangle compactification that might be of independent interest. In particular, we show that no compactification of a non-locally finite graph can both be Hausdorff and have a totally disconnected remainder.

3.2. Reviewing Diestel’s tangle compactification

From now on, we fix an arbitrary connected simple infinite graph $G = (V, E)$.

3.2.1. The 1-complex of a graph

In the *1-complex* of G which we denote also by G , every edge $e = xy$ is a homeomorphic copy $[x, y] := \{x\} \sqcup \mathring{e} \sqcup \{y\}$ of $\mathbb{I} = [0, 1]$ with \mathring{e} corresponding to $(0, 1)$ and points in \mathring{e} being called *inner edge points*. The space $[x, y]$ is called a *topological edge*, but we refer to it simply as *edge* and denote it by e as well. For each subcollection $F \subseteq E$ we write \mathring{F} for the set $\bigsqcup_{e \in F} \mathring{e}$ of inner edge points of edges in F . By $E(v)$ we denote the set of edges incident with a vertex v . The point set of G is $V \sqcup \mathring{E}$, and an open neighbourhood basis of a vertex v of G is given by the unions $\bigcup_{e \in E(v)} [v, i_e)$ of half open intervals with each i_e some inner edge point of e . Note that the 1-complex of G is (locally) compact if and only if the graph G is (locally) finite, and also that the 1-complex fails to be first-countable at vertices of infinite degree. Note that if the graph G has no isolated vertices, then its 1-complex can be obtained from the disjoint sum $\bigoplus_{e \in E} \mathbb{I}_e$ of copies \mathbb{I}_e of the unit interval by taking the quotient with respect to a suitable equivalence relation on $\bigoplus_{e \in E} \{0, 1\}$.

3.2.2. Tangles

Next, we formally introduce a definition of \aleph_0 -tangles provided by Diestel [25] which, as he proved, is equivalent to the original one due to Robertson and Seymour [73]. In the next subsection, however, we explain a third, equivalent viewpoint for tangles (due to Diestel), which describes \aleph_0 -tangles as the elements of the compact Hausdorff inverse limit $\varprojlim \beta(\mathcal{C}_X)$ and which we take as our point of reference for the remainder of this chapter.

The *interior* of a star $\sigma = \{(A_i, B_i) \mid i \in I\} \subseteq \vec{S}_{\aleph_0}$ is the intersection $\bigcap_{i \in I} B_i$.

3. Tangles and the Stone-Ćech compactification of infinite graphs

Definition 3.2.1. An \aleph_0 -tangle (of G) is a consistent orientation of S_{\aleph_0} that contains no finite star of finite interior as a subset. We write Θ for the set of all \aleph_0 -tangles.

3.2.3. Ends and Tangles

If ω is an end of G , then letting

$$\tau_\omega := \{ (A, B) \in \vec{S}_{\aleph_0} \mid C(A \cap B, \omega) \subseteq G[B \setminus A] \}$$

defines an injection $\Omega \hookrightarrow \Theta$, $\omega \mapsto \tau_\omega$ from the ends of G into the \aleph_0 -tangles. Therefore, we call the tangles of the form τ_ω the *end tangles* of G . By abuse of notation we write Ω for the collection of all end tangles of G , so we have $\Omega \subseteq \Theta$.

In order to understand the \aleph_0 -tangles that are not ends, Diestel studied an inverse limit description of Θ . If τ is an \aleph_0 -tangle of the graph, then for each $X \in \mathcal{X}$ it also chooses one *big side* from each bipartition $\{\mathcal{C}, \mathcal{C}'\}$ of \mathcal{C}_X , namely the $\mathcal{H} \in \{\mathcal{C}, \mathcal{C}'\}$ with $(X, \mathcal{H}) \in \tau$. Since it chooses these sides consistently, it induces an ultrafilter $U(\tau, X)$ on \mathcal{C}_X , one for every $X \in \mathcal{X}$, which is given by

$$U(\tau, X) = \{ \mathcal{C} \subseteq \mathcal{C}_X \mid (X, \mathcal{C}) \in \tau \},$$

and these ultrafilters are compatible in that they form a limit of the inverse system $\{ \beta(\mathcal{C}_X), \beta(c_{X',X}), \mathcal{X} \}$. Here, each set \mathcal{C}_X is endowed with the discrete topology and $\beta(\mathcal{C}_X)$ denotes its Stone-Ćech compactification. Every bonding map $\beta(c_{X',X})$ is the unique continuous extension of $c_{X',X}$ that is provided by the Stone-Ćech property (see Theorem 3.4.1 (ii)). As one of his main results, Diestel showed that the map

$$\tau \mapsto (U(\tau, X) \mid X \in \mathcal{X})$$

defines a bijection between the tangle set Θ and the inverse limit $\varprojlim \beta(\mathcal{C}_X)$. From now on, we view the tangle space Θ as the compact Hausdorff space $\varprojlim \beta(\mathcal{C}_X)$.

In his paper, Diestel moreover showed that the ends of G are precisely those \aleph_0 -tangles whose induced ultrafilters are all principal. For every \aleph_0 -tangle τ we write \mathcal{X}_τ for the collection of all $X \in \mathcal{X}$ for which the induced ultrafilter $U(\tau, X)$ is free. The set \mathcal{X}_τ is empty if and only if τ is an end tangle; an \aleph_0 -tangle τ with \mathcal{X}_τ non-empty is called an *ultrafilter tangle*. For every ultrafilter tangle τ the set \mathcal{X}_τ has a least element X_τ of which it is the up-closure. We characterised the sets of the form X_τ combinatorially in [62, Theorem 4.10]: they are precisely the *critical* vertex sets of G , finite sets $X \subseteq V$ whose deletion leaves some infinitely many components each with neighbourhood precisely equal to X , and they can be used together with the ends to compactify the graph, [62, Theorem 4.11].

We conclude our summary of ‘Ends and tangles’ with the formal construction of the tangle compactification. To obtain the tangle compactification $|G|_\Theta$ of a graph G we extend the 1-complex of G to a topological space $G \sqcup \Theta$ by declaring as open in addition to the open sets of G , for all $X \in \mathcal{X}$ and all $\mathcal{C} \subseteq \mathcal{C}_X$, the sets

$$\mathcal{O}_{|G|_\Theta}(X, \mathcal{C}) := \bigcup \mathcal{C} \cup \mathring{E}(X, \bigcup \mathcal{C}) \cup \{ \tau \in \Theta \mid \mathcal{C} \in U(\tau, X) \}$$

and taking the topology this generates. Notably, $|G|_\Theta$ contains Θ as a subspace.

3. Tangles and the Stone-Čech compactification of infinite graphs

Theorem 3.2.2 ([25, Theorem 1]). *Let G be any graph, possibly disconnected.*

- (i) $|G|_{\Theta}$ is a compactification of G with totally disconnected remainder.
- (ii) If G is locally finite and connected, then $|G|_{\Theta}$ coincides with the Freudenthal compactification of G .

The tangle compactification is Hausdorff if and only if G is locally finite. However, the subspace $|G|_{\Theta} \setminus \mathring{E}$ is compact Hausdorff. Teegen [81] generalised the tangle compactification to topological spaces.

3.3. Tangles as inverse limit of finite spaces

The Stone-Čech compactification of a discrete space can be viewed as the inverse limit of all its finite partitions, where each finite partition carries the discrete topology. In this section, we extend this fact to the tangle space.

We start by choosing the point set for our directed poset:

$$\Gamma := \{ (X, P) \mid X \in \mathcal{X} \text{ and } P \text{ is a finite partition of } \mathcal{C}_X \}.$$

Notation. If an element of Γ is introduced just as γ , then we write $X(\gamma)$ and $P(\gamma)$ for the sets satisfying $(X(\gamma), P(\gamma)) = \gamma$. Given $X \subseteq X' \in \mathcal{X}$ and a finite partition P of \mathcal{C}_X we write $P \upharpoonright X'$ for the finite partition

$$\{ c_{X',X}^{-1}(\mathcal{C}) \mid \mathcal{C} \in P \} \setminus \{\emptyset\}$$

that P induces on $\mathcal{C}_{X'}$.

Letting $(X, P) \leq (Y, Q)$ whenever $X \subseteq Y$ and Q refines $P \upharpoonright Y$ defines a directed partial ordering on Γ :

Lemma 3.3.1. (Γ, \leq) is a directed poset.

Proof. Checking the poset properties is straightforward; we verify that it is directed: Given any two elements (X, P) and (Y, Q) of Γ let R be the coarsest refinement of $P \upharpoonright (X \cup Y)$ and $Q \upharpoonright (X \cup Y)$. Then $(X, P), (Y, Q) \leq (X \cup Y, R) \in \Gamma$. \square

For a reason that will become clear in the proof of our next theorem, we consider a cofinal subset of Γ , namely

$$\Gamma' := \{ \gamma \in \Gamma \mid \forall \mathcal{C} \in P(\gamma): V[\mathcal{C}] \text{ is infinite} \}.$$

Lemma 3.3.2. Γ' is cofinal in Γ .

Proof. Given $(X, P) \in \Gamma$ we put

$$X' = X \cup \bigcup \{ V[\mathcal{C}] \mid \mathcal{C} \in P \text{ with } V[\mathcal{C}] \text{ finite} \}.$$

Then $(X, P) \leq (X', P \upharpoonright X') \in \Gamma'$ as desired. \square

3. Tangles and the Stone-Čech compactification of infinite graphs

We aim to describe the tangle space as an inverse limit of finite Hausdorff spaces. For this, we choose Γ as our directed poset, and for each $\gamma \in \Gamma$ we let \mathcal{P}_γ be the set $P(\gamma)$ endowed with the discrete topology. Our bonding maps $f_{\gamma',\gamma}: \mathcal{P}_{\gamma'} \rightarrow \mathcal{P}_\gamma$ send each $\mathcal{C}' \in \mathcal{P}_{\gamma'}$ to the unique $\mathcal{C} \in \mathcal{P}_\gamma$ with $\mathcal{C}' \upharpoonright X(\gamma) \subseteq \mathcal{C}$. Since the spaces \mathcal{P}_γ are compact Hausdorff, so is their inverse limit

$$\mathcal{P} := \varprojlim (\mathcal{P}_\gamma \mid \gamma \in \Gamma).$$

By [36, Corollary 2.5.11] we may replace Γ with its cofinal subset Γ' without changing the inverse limit \mathcal{P} , so we assume without loss of generality that $\Gamma = \Gamma'$.

Notation. If τ is an \aleph_0 -tangle and $\gamma = (X, P) \in \Gamma$ is given, then we write $\mathcal{C}(\tau, \gamma)$ for the unique partition class of P that is contained in the ultrafilter $U(\tau, X)$.

Theorem 3.3.3. *For any graph G , its tangle space is homeomorphic to the inverse limit \mathcal{P} , i.e. $\Theta \cong \mathcal{P}$.*

Proof. Letting $\varphi_\gamma: \Theta \rightarrow \mathcal{P}_\gamma$ assign $\mathcal{C}(\tau, \gamma)$ to each tangle $\tau \in \Theta$ defines a collection of maps that are compatible as tangles are consistent. To see that our maps are continuous, it suffices to note that for all $\gamma \in \Gamma$ and $\mathcal{C} \in \mathcal{P}_\gamma$ we have

$$\varphi_\gamma^{-1}(\mathcal{C}) = \{ \tau \in \Theta \mid \mathcal{C} \in U(\tau, X(\gamma)) \}.$$

The set $V[\mathcal{C}]$ is infinite due to $\Gamma = \Gamma'$, so Diestel's [25, Lemma 3.7] ensures that the preimage $\varphi_\gamma^{-1}(\mathcal{C})$ is non-empty, i.e. that our maps are surjective. Since the tangle space Θ is compact and the inverse limit \mathcal{P} is Hausdorff, the maps φ_γ combine into a continuous surjection $\varphi: \Theta \rightarrow \mathcal{P}$ (cf. [36, Corollary 3.2.16]). Moreover, φ is injective, so it follows from compactness that φ is a homeomorphism. \square

3.4. Background on the Stone-Čech compactification of an infinite graph

3.4.1. Stone-Čech compactification of 1-complexes

The following characterisation of the Stone-Čech compactification is well-known:

Theorem 3.4.1 (Cf. [36],[39]). *Let X be a Tychonoff space. The following are equivalent for a Hausdorff compactification $\gamma X \supseteq X$:*

- (i) $\gamma X = \beta X$,
- (ii) every continuous function $f: X \rightarrow T$ to a compact Hausdorff space T has a continuous extension $\hat{f}: \gamma X \rightarrow T$ with $\hat{f} \upharpoonright X = f$,
- (iii) every continuous function $f: X \rightarrow \mathbb{I}$ has a continuous extension $\hat{f}: \gamma X \rightarrow \mathbb{I}$ with $\hat{f} \upharpoonright X = f$.

Moreover, if X is normal², then we may add

²In this chapter, the property *normal* always includes *Hausdorff*.

3. Tangles and the Stone-Čech compactification of infinite graphs

- (iv) any two closed disjoint sets $Z_1, Z_2 \subseteq X$ have disjoint closures in γX ,
(v) for any two closed sets $Z_1, Z_2 \subseteq X$ we have

$$\overline{Z_1 \cap Z_2}^{\gamma G} = \overline{Z_1}^{\gamma G} \cap \overline{Z_2}^{\gamma G}.$$

Remarkably, from (iv) it follows that whenever X is normal and $Y \subseteq X$ is closed, then $\overline{Y}^{\beta X} = \beta Y$. (Also cf. [36, Corollary 3.6.8].) In particular $\overline{V}^{\beta G} = \beta V$.

Ultrafilter limits

Consider a compact Hausdorff space X . If $x = (x_i \mid i \in I)$ is a family of points $x_i \in X$ and U is an ultrafilter on the index set I , then there is a unique point $x_U \in \overline{\{x_i \mid i \in I\}} \subseteq X$ defined by

$$\{x_U\} = \bigcap_{J \in U} \overline{\{x_i \mid i \in J\}}.$$

Indeed, since U is a filter, the collection $\{\{x_i \mid i \in J\} \mid J \in U\}$ has the finite intersection property, and so by compactness of X , the intersection over their closures is non-empty; and it follows from Hausdorffness of X that the intersection can contain at most one point. We also write

$$x_U = U\text{-lim } x = U\text{-lim } (x_i \mid i \in I)$$

and call x_U the limit of $(x_i \mid i \in I)$ along U , or U -limit of x . Note that if U is the principal ultrafilter generated by $i \in I$, then $x_U = x_i$.

For an alternative description, put $T = \overline{\{x_i \mid i \in I\}} \subseteq X$ and view I as a discrete space, so that the index function

$$\tilde{x}: I \rightarrow \{x_i \mid i \in I\} \subseteq T, \quad i \mapsto x_i$$

is continuous and βI is given by the space of ultrafilters on I . Then the Stone-Čech extension $\beta \tilde{x}: \beta I \rightarrow T$ of the index function \tilde{x} maps each ultrafilter $U \in \beta I$ to x_U .

More generally, if $(X_i \mid i \in I)$ is a family of subsets of a compact Hausdorff space X and U is an ultrafilter on the index set I , then we write

$$X_U = U\text{-lim } (X_i \mid i \in I) := \bigcap_{J \in U} \bigcup_{i \in J} X_i \subseteq X$$

and call X_U the U -limit of $(X_i \mid i \in I)$. Regarding ultrafilter limits, we have the following well-known lemma.

Lemma 3.4.2. *Suppose that $f: X \rightarrow D$ is a continuous surjection with X normal and D discrete. Then the fibres of $\beta f: \beta X \rightarrow \beta D$ are precisely the sets $U\text{-lim } (f^{-1}(d) \mid d \in D)$ with U an ultrafilter on D .*

Proof. First, for an arbitrary subset $J \subseteq D$ the preimages $f^{-1}(J)$ and $f^{-1}(D \setminus J)$ partition X into closed subsets, and hence induce a partition of βX into closed

3. Tangles and the Stone-Čech compactification of infinite graphs

subsets $\overline{f^{-1}(J)}$ and $\overline{f^{-1}(D \setminus J)}$. Since also $(\beta f)^{-1}(\overline{J})$ and $(\beta f)^{-1}(\overline{D \setminus J})$ partition βX , it follows from $f^{-1}(J) \subseteq (\beta f)^{-1}(\overline{J})$ that $f^{-1}(J) = (\beta f)^{-1}(\overline{J})$ for all $J \subseteq D$.

Therefore, for an arbitrary ultrafilter $U \in \beta D$ we have

$$(\beta f)^{-1}(U) = (\beta f)^{-1}\left(\bigcap_{J \in U} \overline{J^{\beta D}}\right) = \bigcap_{J \in U} \overline{f^{-1}(J)}^{\beta X} = U\text{-lim}(f^{-1}(d) \mid d \in D),$$

which is the assertion of the lemma. \square

Two facts about continua

We shall need the following two simple lemmas about continua. Recall that a continuum is a non-empty compact connected Hausdorff space.

Lemma 3.4.3. *Let X be a compact Hausdorff space, and $C \subseteq X$ a connected subspace. Then $\overline{C} \subseteq X$ is a continuum.* \square

A family $(C_i \mid i \in I)$ of subcontinua of some topological space is said to be *directed* if for any $i, j \in I$ there exists a $k \in I$ such that $C_k \subseteq C_i \cap C_j$.

Lemma 3.4.4 ([36, Theorem 6.1.18]). *The intersection of any directed family of continua is again a continuum.* \square

The Stone-Čech compactification of a disjoint sum of intervals

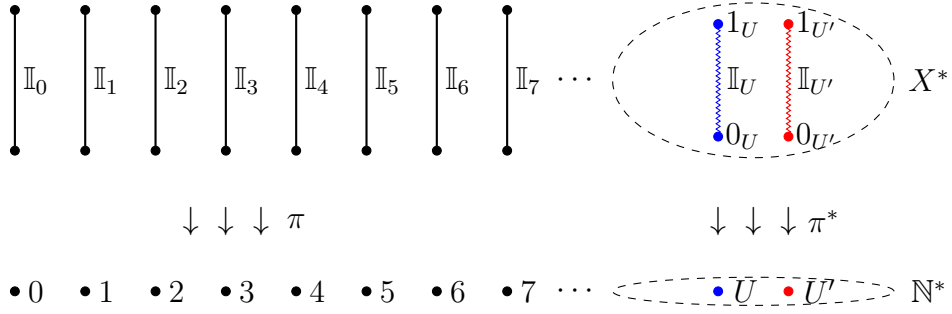
Recall that the 1-complex of a connected graph G can be obtained from the topological sum of disjoint unit intervals (one for each edge) by identifying suitable endpoints, and using the quotient topology. To formalise this, consider the topological space $\mathbb{M}_E = \mathbb{I} \times E$ where $E = E(G)$ carries the discrete topology. Then $G = \mathbb{M}_E / \sim$ for some suitable equivalence relation identifying endpoints. Write \mathbb{I}_e for $\mathbb{I} \times \{e\} \subseteq \mathbb{M}_E$, and x_e for $(x, e) \in \mathbb{I}_e$, so $\mathbb{M}_E = \bigoplus_{e \in E} \mathbb{I}_e$.

Our next results, and in particular Theorem 3.4.9, say that the Stone-Čech compactification of a 1-complex G (which to our knowledge hasn't been studied at all) can be understood through the Stone-Čech compactification $\beta \mathbb{M}_E$ of \mathbb{M}_E (which has been studied extensively over the past decades, see e.g. the survey [48]).

Lemma 3.4.5 ([48, Corollary 2.2]). *Let $X = \bigoplus_{i \in I} K_i$ be a topological sum of continua, and view I as a discrete space. Consider the continuous projection $\pi: X \rightarrow I$, sending K_i to $i \in I$. The components of βX are the fibres of the map $\beta\pi: \beta X \rightarrow \beta I$.*

Suppose for a moment that $X = \bigoplus_{i \in I} K_i$ has only countably many components, i.e. that $I = \mathbb{N}$. Write $X^* = \beta X \setminus X$ for the Stone-Čech remainder. In the lemma, $\beta\pi$ denotes the Stone-Čech extension of π , where we interpret π as a continuous map from X into the compact Hausdorff space $\beta \mathbb{N} \supseteq \mathbb{N}$. And since π has compact fibres (also called *perfect map*), the extension $\beta\pi$ restricts to a continuous map $\pi^* = \beta\pi \upharpoonright X^*: X^* \rightarrow \mathbb{N}^*$, i.e. it maps the remainder of βX to the remainder of $\beta \mathbb{N}$, [36, Theorem 3.7.16]. The figure below illustrates this for $X = \mathbb{M}_{\mathbb{N}}$:

3. Tangles and the Stone-Čech compactification of infinite graphs



Now, for every ultrafilter $U \in \beta\mathbb{N}$ the fibre $\beta\pi^{-1}(U)$ is a connected component of βX , which is also denoted by K_U . By Lemma 3.4.2 we have

$$\beta\pi^{-1}(U) = K_U = U\text{-lim} (K_i \mid i \in I) = \bigcap_{J \in U} \overline{\bigcup_{i \in J} K_i}^{\beta X}.$$

Also, if $(x_i \mid i \in I)$ is a family of points with $x_i \in K_i$, then x_U is the unique point of $K_U \cap \overline{\{x_i \mid i \in I\}}^{\beta X}$. If the spaces K_i are homeomorphic copies of a single space and the points $x_i \in K_i$ correspond to the same point ξ of the original space, then we write ξ_U for x_U . For example, if each K_i is a copy of the unit interval and x_i corresponds to 0 for all $i \in I$, then $x_U = 0_U$.

We shall also need the following lemma plus corollary:

Lemma 3.4.6 ([48, Lemma 2.3]). *For a family $(x_i \mid i \in I)$ of points $x_i \in K_i$, the point x_U is a cut-point of K_U if and only if $\{i \mid x_i \text{ is a cut-point of } K_i\} \in U$.*

Notation. In the context of $X = \mathbb{M}_E$ we write $\check{\mathbb{I}}_U$ for $\mathbb{I}_U \setminus \{0_U, 1_U\}$.

Corollary 3.4.7. *The spaces $\mathbb{I}_U \setminus \{0_U\}$, $\mathbb{I}_U \setminus \{1_U\}$ and $\check{\mathbb{I}}_U$ are connected.*

Proof. The non-standard interval $[0_U, (\frac{1}{2})_U]$ is homeomorphic to \mathbb{I}_U (cf. [48, Proposition 2.8]). Thus $(0_U, (\frac{1}{2})_U]$ is connected by Lemma 3.4.6. So is $[(\frac{1}{2})_U, 1_U)$. Since both meet in $(\frac{1}{2})_U$, so is their union $\check{\mathbb{I}}_U$. \square

Quotients

As we are interested in 1-complexes, i.e. in quotients of \mathbb{M}_E , we provide a theorem how the quotient operation relates to the Stone-Čech functor. We need the following lemma, which is easily verified (alternatively see Theorems 2.4.13 and 1.5.20 from [36]).

Lemma 3.4.8. *Let V be a closed discrete subset of a normal space X , and suppose that \sim is an equivalence relation on V . Then X/\sim is again normal.* \square

3. Tangles and the Stone-Čech compactification of infinite graphs

Theorem 3.4.9. *Let V be a closed discrete subset of a normal space X , and suppose that \sim is an equivalence relation on V . Let $\{V_i \mid i \in I\}$ be the collection of all \sim -classes. Consider the equivalence relation \sim_β on $\overline{V}^{\beta X}$ into equivalence classes of the form*

$$V_U = U\text{-lim}(V_i \mid i \in I) = \bigcap_{J \in U} \bigcup_{i \in J} V_i^{\beta X},$$

one for each ultrafilter U on I , and singletons. Then X/\sim is again normal and

$$\beta(X/\sim) = (\beta X)/\sim_\beta.$$

Proof. Let us write $V/\sim = I$ where I is endowed with the discrete topology. The quotient X/\sim is normal by Lemma 3.4.8, so its Stone-Čech compactification exists. Also, the quotient map $q: X \rightarrow X/\sim$ is a continuous closed map ([36, Proposition 2.4.3]), and so $q[V] = I$ is closed in X/\sim .

Now by the Stone-Čech property in Theorem 3.4.1 (ii), the map $q: X \rightarrow X/\sim \subseteq \beta(X/\sim)$ extends to a continuous surjection $\beta q: \beta X \rightarrow \beta(X/\sim)$.

We claim that each non-trivial fibre of βq is of the form V_U for each ultrafilter U on I . Since every continuous surjection $f: Z \rightarrow Y$ from a compact space Z onto a Hausdorff space Y gives rise to a homeomorphism between the quotient $Z/\{f^{-1}(y) \mid y \in Y\}$ over the fibres of f and Y , this implies the desired result.

First, note that βq maps $\overline{V}^{\beta X}$ onto $\overline{I}^{\beta(X/\sim)}$, and restricts to a bijection on the respective complements, as V is closed. Moreover, as $I \subseteq X/\sim$ is closed and discrete, we have $\overline{I}^{\beta(X/\sim)} = \beta I = \{U \mid U \text{ is an ultrafilter on } I\}$. Hence, by Lemma 3.4.2 the fibres of βq are just $(\beta q)^{-1}(U) = V_U$, one for each ultrafilter U on I . \square

Corollary 3.4.10. *Let X be a normal space and $V \subseteq X$ a closed discrete subset. Then X/V is again normal and*

$$\beta(X/V) = \beta X / (\overline{V}^{\beta X}). \quad \square$$

Corollary 3.4.11. *Let X and Y be two disjoint normal spaces, and suppose that $A = \{a_i \mid i \in I\} \subseteq X$ and $B = \{b_i \mid i \in I\} \subseteq Y$ are infinite closed discrete subspaces. Consider the quotient $Z = (X \oplus Y)/\sim$ where we identify pairs $\{a_i, b_i\}$ for all $i \in I$. Then*

$$\beta Z = (\beta X \oplus \beta Y)/\sim_\beta$$

where we identify pairs $\{a_U, b_U\}$ for all ultrafilters U on I . \square

3.4.2. Three examples

Before turning towards the proof of our main result, we illustrate the above topological lemmas by three representative examples: We discuss the Stone-Čech compactification of the infinite ray R , the infinite star S_λ of degree λ , and the dominated ray D , and compare it side by side with the \aleph_0 tangles of these examples.

3. Tangles and the Stone-Čech compactification of infinite graphs

The infinite ray

Consider the infinite ray R with vertex set $V = \{v_n \mid n \in \mathbb{N}\}$ and edge set $E = \{v_n v_{n+1} \mid n \in \mathbb{N}\}$. Since R is locally finite, the space of \aleph_0 -tangles consists solely of the single end of R , by Theorem 3.2.2 (ii). Moreover, the 1-complex R is homeomorphic to the positive half line $\mathbb{H} = [0, \infty)$, so they have the same Stone-Čech remainder $R^* = \mathbb{H}^*$. The space \mathbb{H}^* has been extensively investigated, see e.g. [48] for a survey. At this point, however, we are content to provide the standard argument showing that the Stone-Čech remainder of the infinite ray is indeed connected, confirming the connection between components in the remainder of the Stone-Čech compactification and the \aleph_0 -tangles.

Example 3.4.12. The infinite ray has a connected Stone-Čech remainder.

Proof. Deleting a vertex v_n from R leaves behind exactly one infinite component $C_n = R[v_{n+1}, v_{n+2}, \dots]$. Then $\bigcap_{n \in \mathbb{N}} \overline{C_n}^{\beta R}$ is a continuum by Lemmas 3.4.3 and 3.4.4. We claim that

$$R^* = \bigcap_{n \in \mathbb{N}} \overline{C_n}^{\beta R}.$$

Indeed, “ \supseteq ” holds as any vertex and edge of R is removed eventually by the intersection. For “ \subseteq ” note that for any $n \in \mathbb{N}$ we have $R = R[v_0, \dots, v_{n+1}] \cup C_n$, and hence

$$R^* \subseteq \overline{R[v_0, \dots, v_{n+1}]}^{\beta R} \cup \overline{C_n}^{\beta R},$$

since the closure operator distributes over finite unions. But $R[v_0, \dots, v_{n+1}]$ is compact, and hence closed in the Hausdorff space βR , implying

$$\overline{R[v_0, \dots, v_{n+1}]}^{\beta R} = R[v_0, \dots, v_{n+1}] \subseteq R.$$

It follows $R^* \subseteq \overline{C_n}^{\beta R}$ for all $n \in \mathbb{N}$ as desired. \square

The infinite star




Figure 3.4.1.: The Stone-Čech compactification of the countable infinite star

For any cardinal λ we denote by S_λ the star of degree λ . Clearly, this star has no end, so all \aleph_0 -tangles are ultrafilter tangles. As a consequence of [62, Theorem 4.10],

3. Tangles and the Stone-Čech compactification of infinite graphs

the ultrafilter tangles correspond precisely to the free ultrafilters on λ . The 1-complex of S_λ is obtained from \mathbb{M}_E (with E a discrete space of cardinality λ) via

$$S_\lambda = \mathbb{M}_E / \{0_e \mid e \in E\}.$$

Example 3.4.13. The Stone-Čech remainder of an infinite star S_λ is homeomorphic to $\mathbb{M}_E^* \setminus \{0_U \mid U \in E^*\}$. Each connected component of S_λ^* is homeomorphic to $\mathbb{I}_U \setminus \{0_U\}$ for some free ultrafilter $U \in E^*$.

Proof. Since $S_\lambda = \mathbb{M}_E / \{0_e \mid e \in E\}$, it follows immediately from Corollary 3.4.10 that $\beta S_\lambda = \beta \mathbb{M}_E / \overline{\{0_e \mid e \in E\}}^{\beta \mathbb{M}_E}$. Since the equivalence class $\overline{\{0_e \mid e \in E\}}^{\beta \mathbb{M}_E}$ corresponds to the center vertex of S_λ , it follows for the remainder of βS_λ that

$$S_\lambda^* = \mathbb{M}_E^* \setminus \overline{\{0_e \mid e \in E\}}^{\beta \mathbb{M}_E} = \mathbb{M}_E^* \setminus \{0_U \mid U \in E^*\}.$$

By Lemma 3.4.5 and Corollary 3.4.7, the connected components of the remainder $\mathbb{M}_E^* \setminus \{0_U \mid U \in E^*\}$ are given by $\mathbb{I}_U \setminus \{0_U\}$ for each free ultrafilter U on E . \square

The dominated ray

The dominated ray D is the quotient of an infinite star S_{\aleph_0} and a ray R where the leaves of S_{\aleph_0} , denoted as in the previous example by $\{1_n \mid n \in \mathbb{N}\}$, are identified pairwise with vertices of the ray, denoted by $\{v_n \mid n \in \mathbb{N}\}$ (see Fig. 3.4.2). Since deleting any finite set of vertices from D leaves only one infinite component, the sole end of D is the one and only \aleph_0 -tangle.




Figure 3.4.2.: The dominated ray with dominating vertex c

Example 3.4.14. The dominated ray D has a connected Stone-Čech remainder.

Proof. By Corollary 3.4.11, the Stone-Čech remainder of D is homeomorphic to the quotient $(S_{\aleph_0}^* \oplus R^*) / \sim_\beta$ where $1_U \sim_\beta v_U$ for every ultrafilter $U \in \mathbb{N}^*$ and $1_U \in \mathbb{I}_U$ and $v_U \in R^*$. It follows that every connected component $\mathbb{I}_U \setminus \{0_U\}$ of $S_{\aleph_0}^*$ (see Example 3.4.13) is, via the identified points $1_U \sim_\beta v_U$, attached to the connected remainder R^* (see Example 3.4.12) of βR , and so D^* is indeed connected. \square

3.5. Comparing the Stone-Čech remainder with the tangle space

3.5.1. The Stone-Čech remainder of the vertex set

Due to $\beta G = (\beta \mathbb{M}_E) / \sim_\beta$ for any representation \mathbb{M}_E / \sim of G (Theorem 3.4.9) we may view $\beta V = \overline{V}^{\beta G} \subseteq \beta G$ as the closure of $\{[0_e]_{\sim_\beta}, [1_e]_{\sim_\beta} \mid e \in E\}$ in the quotient $(\beta \mathbb{M}_E) / \sim_\beta$. In particular, the non-standard intervals \mathbb{I}_U (with $U \in E^*$) may interact with V or its Stone-Čech remainder V^* . In this subsection, we have a closer look at this interaction.

In the next lemma, we write $V^* = G^* \cap \overline{V}^{\beta G}$. Since $\beta V = \overline{V}^{\beta G}$, this potential double meaning does no harm.

Lemma 3.5.1. *Let \mathbb{M}_E / \sim be a representation of G , and let $U \in E^*$ be any free ultrafilter. Then at most one of the endpoints 0_U and 1_U of \mathbb{I}_U is contained in some \sim_β -class that belongs to V , and at least one of them is contained in some \sim_β -class that belongs to V^* .*

Proof. A vertex $x \in G$, viewed as \sim_β -class (Theorem 3.4.9), contains an endpoint of \mathbb{I}_U if and only if $E(x) \in U$. And since $|E(x) \cap E(y)| \leq 1$ for every distinct two vertices $x, y \in G$, at most one vertex $x \in G$ can satisfy $E(x)$. \square

Lemma 3.5.2. *Let G be a graph, and let C be a connected component of the Stone-Čech remainder G^* . Then $C \cap V^* \neq \emptyset$. In particular, the connected components of G^* induce a closed partition of V^* .*

Proof. Consider a representation $G = \mathbb{M}_E / \sim$ of G , and recall that by Corollary 3.4.7, every non-standard component \mathbb{I}_U of \mathbb{M}_E^* remains connected upon deleting one or both of the endpoints 0_U and 1_U .

Consider some connected component C of G^* . Then for some $\mathbb{I}_U \subseteq \mathbb{M}_E^*$ we have $\check{\mathbb{I}}_U \subseteq C$. Therefore, it suffices to show that for every free ultrafilter $U \in E^*$ at least one of $[0_U]_{\sim_\beta}$ and $[1_U]_{\sim_\beta}$ is in V^* . This is the content of Lemma 3.5.1. \square

3.5.2. An auxiliary remainder

The remainder G^* not being compact prevents us from using topological machinery, so we study a nice subspace $G^\times \subseteq G^*$ first. As usual, we start with some new notation.

Notation. For a vertex v of G , write $O(v)$ for its open neighbourhood $\overset{\circ}{E}(v) \sqcup \{v\}$ in G consisting of all half-open incident edges at v , and write

$$O_{\beta G}(v) := \overline{O(v)}^{\beta G} \setminus \overline{N(v)}^{\beta G}.$$

Due to $\beta G = \overline{O(v)}^{\beta G} \cup \overline{G \setminus O(v)}^{\beta G}$ and $\overline{O(v)}^{\beta G} \cap \overline{G \setminus O(v)}^{\beta G} = \overline{N(v)}^{\beta G}$ the set $O_{\beta G}(v)$ is open in βG , and it meets G precisely in $O(v)$. The set $O_{\beta G}(v)$ is also known as $\text{Ex } O(v) = \beta G \setminus \overline{G \setminus O(v)}$, the largest open subset of βG whose intersection with G is $O(v)$, cf. [36, p. 388].

3. Tangles and the Stone-Ćech compactification of infinite graphs

Observation 3.5.3. Put $F = E(v)$ and write H for the subspace $\bigcup F \subseteq G$. Since H is the 1-complex of a star, the set $O_{\beta G}(v)$ is homeomorphic to the space from Example 3.4.13 without the “endpoints” (also see Fig. 3.4.1):

$$\begin{aligned} O_{\beta G}(v) &= \overline{H}^{\beta G} \setminus \overline{N(v)}^{\beta G} \cong \beta H \setminus \overline{N(v)}^{\beta H} \\ &\cong (\beta \mathbb{M}_F / \{0_U \mid U \in \beta F\}) \setminus \{1_U \mid U \in \beta F\} \end{aligned}$$

Definition 3.5.4. The auxiliary remainder of G is the space

$$G^\times := \beta G \setminus O_{\beta G}[V] \subseteq G^*$$

where we write $O_{\beta G}[W] = \bigcup_{v \in W} O_{\beta G}(v)$ for all $W \subseteq V$.

Fact 3.5.5. Since βG is compact Hausdorff, so is G^\times .

Lemma 3.5.6. The vertex set V of any graph satisfies $V^* \subseteq G^\times$.

Proof. We show that, for every vertex $v \in V$, the set $O_{\beta G}(v)$ avoids V^* :

$$\begin{aligned} \overline{UE(v)}^{\beta G} \cap V^* &= \left(\overline{UE(v)}^{\beta G} \cap \overline{V}^{\beta G} \right) \setminus G = \overline{\{v\} \sqcup N(v)}^{\beta G} \setminus G \\ &= \left(\{v\} \sqcup \overline{N(v)}^{\beta G} \right) \setminus G = N(v)^* \subseteq \overline{N(v)}^{\beta G} \quad \square \end{aligned}$$

3.5.3. The components of the remainder can be distinguished by finite separators

For the tangle compactification it is true that every open set $\mathcal{O}_{|G|_\Theta}(X, \mathcal{C})$ gives rise to a clopen bipartition of the tangle space, namely

$$\begin{aligned} &(\mathcal{O}_{|G|_\Theta}(X, \mathcal{C}) \cap \Theta) \oplus (\mathcal{O}_{|G|_\Theta}(X, \mathcal{C}_X \setminus \mathcal{C}) \cap \Theta), \\ \text{i.e. } &\{\tau \in \Theta \mid \mathcal{C} \in U(\tau, X)\} \oplus \{\tau \in \Theta \mid \mathcal{C} \notin U(\tau, X)\}. \end{aligned}$$

In fact, for every two distinct \aleph_0 -tangles there exists such a clopen bipartition of the tangle space separating the two. Our next target is to prove that any two components of the remainder of a graph are—just as the \aleph_0 -tangles—distinguished by a finite order separation. That is why we start by studying a possible analogue $\mathcal{O}_{\beta G}(X, \mathcal{C})$ of $\mathcal{O}_{|G|_\Theta}(X, \mathcal{C})$ for βG .

Notation. Given $X \in \mathcal{X}$ and $\mathcal{C} \subseteq \mathcal{C}_X$ we write $G[X, \mathcal{C}]$ for $G[X \cup V[\mathcal{C}]]$. If τ is an \aleph_0 -tangle of G and γ is an element of Γ , then we write $G[\tau, \gamma]$ for $G[X(\gamma), \mathcal{C}(\tau, \gamma)]$.

For every $X \in \mathcal{X}$ and $\mathcal{C} \subseteq \mathcal{C}_X$ we let

$$\mathcal{O}_{\beta G}(X, \mathcal{C}) := \overline{G[X, \mathcal{C}]}^{\beta G} \setminus G[X]$$

which is open in βG as a consequence of $\beta G = \overline{G[X, \mathcal{C}]} \cup \overline{G[X, \mathcal{C}_X \setminus \mathcal{C}]}$ and $\overline{G[X, \mathcal{C}]} \cap \overline{G[X, \mathcal{C}_X \setminus \mathcal{C}]} = \overline{G[X]} = G[X]$ (see Theorem 3.4.1 (v)). Before we check that $\mathcal{O}_{\beta G}(X, \mathcal{C})$ gives rise to clopen bipartitions of G^* and G^\times , we prove a lemma:

3. Tangles and the Stone-Ćech compactification of infinite graphs

Lemma 3.5.7. *For all $X \in \mathcal{X}$ and $\mathcal{C} \subseteq \mathcal{C}_X$ we have*

$$\overline{G[X, \mathcal{C}]^{\beta G}} \subseteq O_{\beta G}[X] \sqcup \overline{\bigcup \mathcal{C}}^{\beta G}.$$

In particular, for all $\gamma \in \Gamma$ we have

$$\beta G = O_{\beta G}[X(\gamma)] \sqcup \bigsqcup_{\mathcal{C} \in P(\gamma)} \overline{\bigcup \mathcal{C}}^{\beta G}.$$

Proof. Due to $\beta G = \bigcup_{\mathcal{C} \in P(\gamma)} \overline{G[X(\gamma), \mathcal{C}]^{\beta G}}$ it suffices to show the first statement:

$$\begin{aligned} \overline{G[X, \mathcal{C}]} &= G[X] \cup \bigcup_{x \in X} \overline{\bigcup E(x, \bigcup \mathcal{C})} \cup \overline{\bigcup \mathcal{C}} \\ &\subseteq \bigcup_{x \in X} (O_{\beta G}(x) \sqcup \overline{N(x) \cap \bigcup \mathcal{C}}) \cup \overline{\bigcup \mathcal{C}} = O_{\beta G}[X] \sqcup \overline{\bigcup \mathcal{C}} \end{aligned}$$

where at the “ \subseteq ” we used Theorem 3.4.1 (v) for

$$\begin{aligned} \overline{\bigcup E(x, \bigcup \mathcal{C})} &= \left(\overline{\bigcup E(x, \bigcup \mathcal{C})} \setminus \overline{N(x)} \right) \sqcup \left(\overline{\bigcup E(x, \bigcup \mathcal{C})} \cap \overline{N(x)} \right) \\ &\subseteq O_{\beta G}(x) \sqcup \overline{\left(\bigcup E(x, \bigcup \mathcal{C}) \right) \cap N(x)} = O_{\beta G}(x) \sqcup \overline{N(x) \cap \bigcup \mathcal{C}}. \quad \square \end{aligned}$$

Lemma and Definition 3.5.8. *Let any $(X, P) \in \Gamma$ be given. Then*

- (i) $P_* := \{ \overline{G[X, \mathcal{C}]^{\beta G}} \cap G^* \mid \mathcal{C} \in P \}$ and
- (ii) $P_\times := \{ \overline{\bigcup \mathcal{C}}^{\beta G} \cap G^\times \mid \mathcal{C} \in P \}$

are finite separations of G^ and G^\times into clopen subsets.*

Proof. (i). First observe that

$$\beta G = \overline{G} = \overline{\bigcup_{\mathcal{C} \in P} G[X, \mathcal{C}]} = \bigcup_{\mathcal{C} \in P} \overline{G[X, \mathcal{C}]}.$$

At the same time, however, since every $G[X, \mathcal{C}]$ is a subgraph, and hence a closed subset of G , for all $\mathcal{C} \neq \mathcal{C}' \in P$ it follows from Theorem 3.4.1 (v) that

$$\overline{G[X, \mathcal{C}]} \cap \overline{G[X, \mathcal{C}']} = \overline{G[X, \mathcal{C}] \cap G[X, \mathcal{C}']} = \overline{G[X]} = G[X] \subseteq G$$

where the last equality follows from the fact that compact subsets of Hausdorff spaces are closed. Hence, we see that G^* is a disjoint union of finitely many closed sets $G^* = \bigsqcup_{\mathcal{C} \in P} (\overline{G[X, \mathcal{C}]} \cap G^*)$.

(ii) follows from (i) with Lemma 3.5.7. □

Notation. We write \approx_* and \approx_\times for the equivalence relations on G^* and G^\times whose classes are precisely the connected components of G^* and G^\times respectively. If C is a component of G^\times we write \hat{C} for the unique component of G^* including it.

3. Tangles and the Stone-Ćech compactification of infinite graphs

Our next lemma, the so-called *Separating Lemma*, can be considered as our main technical result of this chapter, yielding that distinct components of G^* can be distinguished by a finite order separation of the graph G , see Corollaries 3.5.13 and 3.5.15 below. However, we state the lemma in a slightly more general form, so that we can also apply it in Section 3.6 when proving Theorem 3.1. For this, we shall need the following notion of “tame”:

Definition 3.5.9. We call a subset $A \subseteq \beta G$ *tame* if it is \approx_x -closed and for every component C of G^* meeting A in a point of $O_{\beta G}[V]$ (cf. Def. 3.5.4) we have $C \subseteq A$.

Here, for a set M we say that M is R -closed for R an equivalence relation on any set N , if for all $m \in M$ and $n \in N$ with mRn there holds $n \in M$ (phrased differently, whenever M meets an R -class then it contains that class entirely, but M might also contain points that do not lie in any R -class).

Example 3.5.10. All \approx_* -closed subsets of βG and all \approx_x -closed subsets of G^\times are tame, but both $G^* \setminus G^\times$ and $O_{\beta G}[V]$ are not tame as soon as G is not locally finite.

Lemma 3.5.11 (Šura-Bura Lemma [36, Theorem 6.1.23]). *If C_1 and C_2 are distinct components of a compact Hausdorff space X , there is a clopen bipartition $A \oplus B$ of X with $C_1 \subseteq A$ and $C_2 \subseteq B$.*

Lemma 3.5.12 (Separating Lemma). *Let $A, B \subseteq \beta G$ be two disjoint closed and tame subsets. Then there is a finite $X \subseteq V(G)$ and a bipartition $\{\mathcal{C}_1, \mathcal{C}_2\}$ of \mathcal{C}_X with $A \subseteq \overline{G[X, \mathcal{C}_1]}^{\beta G}$ and $B \subseteq \overline{G[X, \mathcal{C}_2]}^{\beta G}$.*

Proof. Given A and B we use normality of G^\times and a compactness argument to deduce from Lemma 3.5.11 that there is a clopen bipartition $K_A \oplus K_B$ of G^\times with $A \cap G^\times \subseteq K_A$ and $B \cap G^\times \subseteq K_B$. Put $A' = A \cup K_A$ and $B' = B \cup K_B$ so A' and B' are closed and disjoint subsets of βG . Using that βG is normal we find disjoint open sets $O_A, O_B \subseteq \beta G$ with $A' \subseteq O_A$ and $B' \subseteq O_B$. Next, since

$$\bigcap_{v \in V} (\beta G \setminus O_{\beta G}(v)) = G^\times = K_A \oplus K_B \subseteq O_A \sqcup O_B$$

is an intersection of closed sets which is contained in the open set $O_A \sqcup O_B$, it follows from compactness that there are finitely many vertices v_1, \dots, v_n such that

$$\bigcap_{i=1}^n (\beta G \setminus O_{\beta G}(v_i)) \subseteq O_A \sqcup O_B.$$

Put $\Xi = \{v_1, \dots, v_n\}$. Then $O_A \sqcup O_B$ induces a clopen bipartition $K'_A \oplus K'_B$ of the closed subspace $\beta G \setminus O_{\beta G}[\Xi]$ of βG which in turn induces a bipartition $Q = \{\mathcal{A}, \mathcal{B}\}$ of \mathcal{C}_Ξ via

$$\mathcal{A} = \{C \in \mathcal{C}_\Xi \mid C \subseteq K'_A\} \quad \text{and} \quad \mathcal{B} = \{C \in \mathcal{C}_\Xi \mid C \subseteq K'_B\}.$$

3. Tangles and the Stone-Ćech compactification of infinite graphs

In particular, we have

$$\overline{\bigcup \mathcal{A}}^{\beta G} \subseteq K'_A \quad \text{and} \quad \overline{\bigcup \mathcal{B}}^{\beta G} \subseteq K'_B. \quad (3.5.1)$$

Moreover, by Lemma 3.5.8, Q_\times must be the clopen bipartition $K_A \oplus K_B$ of G^\times .

Now we want that

$$A \subseteq \overline{G[\Xi, \mathcal{A}]}^{\beta G} \quad \text{and} \quad B \subseteq \overline{G[\Xi, \mathcal{B}]}^{\beta G}, \quad (3.5.2)$$

but with the help of Lemma 3.5.7 and (3.5.1) we only get

$$\begin{aligned} A &\subseteq \beta G \setminus \overline{\bigcup \mathcal{B}}^{\beta G} = \overline{G[\Xi, \mathcal{A}]}^{\beta G} \cup O_{\beta G}[\Xi] \\ \text{and } B &\subseteq \beta G \setminus \overline{\bigcup \mathcal{A}}^{\beta G} = \overline{G[\Xi, \mathcal{B}]}^{\beta G} \cup O_{\beta G}[\Xi] \end{aligned}$$

with A and B possibly meeting $O_{\beta G}[\Xi]$. To resolve this issue, we will find a way to widen Ξ by adding only finitely many vertices, and adjusting \mathcal{A} and \mathcal{B} accordingly so as to make (3.5.2) true.

For this, we note first that

$$A \cap G^* \subseteq \overline{G[\Xi, \mathcal{A}]}^{\beta G}. \quad (3.5.3)$$

Indeed, we know that A is tame, that each component of G^* meets $V^* \subseteq G^\times$ (see Lemmas 3.5.2 and 3.5.6), and that

$$A \cap G^\times \subseteq K_A = \overline{\bigcup \mathcal{A}}^{\beta G} \cap G^\times \subseteq \overline{G[\Xi, \mathcal{A}]}^{\beta G}$$

where $\overline{G[\Xi, \mathcal{A}]}^{\beta G} \cap G^*$ is clopen (Lemma 3.5.8); combining these facts yields (3.5.3).

Second, we show that there exists a finite set F_A of edges of G with

$$A \cap G \subseteq G[\Xi, \mathcal{A}] \cup \bigcup F_A. \quad (3.5.4)$$

Indeed, by $A \subseteq \overline{G[\Xi, \mathcal{A}]}^{\beta G} \cup O_{\beta G}[\Xi]$, it suffices to show that

$$F_A := \{e \in E(X, \bigcup \mathcal{B}) \mid \dot{e} \text{ meets } A\}$$

is finite. Suppose for a contradiction that F_A is infinite, and for every edge $e \in F_A$ pick some $i_e \in \dot{e} \cap A$. Then $\overline{\{i_e \mid e \in F_A\}}^{\beta G} \subseteq A$ meets $G^* \cap \overline{G[X, \mathcal{B}]}^{\beta G} \cap O_{\beta G}[\Xi]$ in some component C of G^* . But we noted earlier that each component of G^* meets $V^* \subseteq G^\times$, so the tame set A meeting C means that $\emptyset \neq C \cap G^\times \subseteq A \cap K_B$, a contradiction. Of course, corresponding versions of (3.5.3) and (3.5.4) hold for B .

Finally, we use (3.5.3) and (3.5.4) to yield a true version of (3.5.2). For this, we let X be the finite vertex set obtained from Ξ by adding the endvertices of the edges in $F_A \cup F_B$, and we put $\mathcal{C}_1 = c_{X, \Xi}^{-1}(\mathcal{A})$ and $\mathcal{C}_2 = c_{X, \Xi}^{-1}(\mathcal{B})$. Due to $G[X, \mathcal{C}_1] \supseteq G[\Xi, \mathcal{A}] \cup \bigcup F_A$ and $G[X, \mathcal{C}_2] \supseteq G[\Xi, \mathcal{B}] \cup \bigcup F_B$, we may use (3.5.3) and (3.5.4) to deduce that

$$A \subseteq \overline{G[X, \mathcal{C}_1]}^{\beta G} \quad \text{and} \quad B \subseteq \overline{G[X, \mathcal{C}_2]}^{\beta G}. \quad \square$$

3. Tangles and the Stone-Ćech compactification of infinite graphs

Using Lemma 3.5.7 we obtain the following corollary:

Corollary 3.5.13. *For every pair of distinct components C_1, C_2 of G^\times there is a finite $X \subseteq V(G)$ and a bipartition $P = \{\mathcal{C}_1, \mathcal{C}_2\}$ of \mathcal{C}_X such that the components $\hat{C}_1 \supseteq C_1$ and $\hat{C}_2 \supseteq C_2$ of G^* are separated by the clopen bipartition P_* of G^* . \square*

Lemma 3.5.14. *The map $C \mapsto \hat{C}$ defines a bijection between $G^\times / \approx_\times$ and G^* / \approx_* .*

Proof. Each component of G^* meets $V^* \subseteq G^\times$ (see Lemmas 3.5.2 and 3.5.6), so the map $C \mapsto \hat{C}$ is onto. It is injective by Corollary 3.5.13. \square

Corollary 3.5.13 and Lemma 3.5.14 yield another important result:

Corollary 3.5.15. *For every pair of distinct components C_1, C_2 of G^* there is a finite $X \subseteq V(G)$ and a bipartition P of \mathcal{C}_X such that the clopen bipartition P_* of G^* separates C_1 and C_2 . \square*

Corollary 3.5.16. *The quotients $G^\times / \approx_\times$ and G^* / \approx_* are Hausdorff. \square*

Theorem 3.5.17. *For any graph G , we have $G^\times / \approx_\times \cong G^* / \approx_*$.*

Proof. Let $\hat{\iota}: G^\times / \approx_\times \rightarrow G^* / \approx_*$ map C to \hat{C} . By Lemma 3.5.14 this is a bijection. Denote the quotient map $G^* \rightarrow G^* / \approx_*$ by q_* . Clearly, the diagram

$$\begin{array}{ccc} G^\times & \xrightarrow{\iota} & G^* \\ \downarrow & & \downarrow q_* \\ G^\times / \approx_\times & \xrightarrow{\hat{\iota}} & G^* / \approx_* \end{array}$$

commutes. Since $G^\times / \approx_\times$ is compact and G^* / \approx_* is Hausdorff (Corollary 3.5.16), to show that $\hat{\iota}$ is a homeomorphism it suffices to verify continuity. But note that by the quotient topology, $\hat{\iota}$ is continuous if and only if $q_* \circ \iota$ is continuous. \square

3.5.4. Comparing \mathcal{P} with G^\times

Now that we are able to distinguish distinct components of the remainder by some $\gamma \in \Gamma$, the next step is to use this to show $\Theta \cong G^* / \approx_*$. Technically, we will achieve this by showing $\mathcal{P} \cong G^\times / \approx_\times$ instead.

For every $\gamma \in \Gamma$ let $\sigma_\gamma: G^\times \rightarrow \mathcal{P}_\gamma$ map every point $x \in G^\times$ to the $\mathcal{C} \in \mathcal{P}_\gamma$ whose induced clopen partition class $\overline{\bigcup \mathcal{C}^{\beta G}} \cap G^\times \in P(\gamma)_\times$ containing x , i.e. including the connected component of G^\times containing x .

Lemma 3.5.18. *The maps σ_γ are continuous surjections.*

Proof. To see that σ_γ is continuous, observe that

$$\sigma_\gamma^{-1}(\mathcal{C}) = \overline{\bigcup \mathcal{C}^{\beta G}} \cap G^\times \in P(\gamma)_\times$$

and recall that partition classes of $P(\gamma)_\times$ are clopen in G^\times .

The map σ_γ is surjective: since every $\mathcal{C} \in \mathcal{P}_\gamma$ is such that $V[\mathcal{C}]$ is infinite, Lemma 3.5.6 ensures that $\overline{\bigcup \mathcal{C}^{\beta G}} \cap G^\times$ is non-empty. \square

3. Tangles and the Stone-Ćech compactification of infinite graphs

Lemma 3.5.19. *The maps σ_γ are compatible.*

Proof. For this assertion it suffices to show that whenever $(X, P) \leq (X', P')$, then we have $P_\times \preceq P'_\times$, i.e. the finite clopen partition P'_\times refines that partition of G^\times induced by P_\times . To see this, consider any $\mathcal{C}' \in P'_\times$. Since P' refines $P \upharpoonright X'$, there is a unique $\mathcal{C} \in P$ with $\mathcal{C}' \upharpoonright X \subseteq \mathcal{C}$. Thus $\bigcup \mathcal{C}'^{\beta G} \cap G^\times \subseteq \bigcup \mathcal{C}^{\beta G} \cap G^\times$ follows. \square

We put $\sigma = \varprojlim \sigma_\gamma: G^\times \rightarrow \mathcal{P}$, and we aim to show that σ gives rise to a homeomorphism between $G^\times / \approx_\times$ and \mathcal{P} .

Lemma 3.5.20. *The map $\sigma: G^\times \rightarrow \mathcal{P}$ is a continuous surjection.*

Proof. We combine Lemmas 3.5.18 and 3.5.19 with the fact that compatible continuous surjections from a compact space onto Hausdorff spaces combine into one continuous surjection onto the inverse limit of their image spaces (cf. [36, Corollary 3.2.16]). \square

Lemma 3.5.21. *The fibres of σ are precisely the connected components of G^\times .*

Proof. First, it is clear by the definition of the σ_γ that every σ_γ is constant on connected components of G^\times . Conversely, we need to argue that for any pair of distinct components C_1 and C_2 of G^\times there is some σ_γ with $\sigma_\gamma \upharpoonright C_1 \neq \sigma_\gamma \upharpoonright C_2$. Such a σ_γ is provided by Corollary 3.5.13. \square

Proposition 3.5.22. $G^\times / \approx_\times \cong \mathcal{P}$.

Proof. It is well-known that every continuous surjection $f: X \rightarrow Y$ from a compact space X onto a Hausdorff space Y gives rise to a homeomorphism between the quotient $X / \{f^{-1}(y) \mid y \in Y\}$ over the fibres of f , and the space Y . Thus, it follows from Lemmas 3.5.20 and 3.5.21, that

$$G^\times / \approx_\times = G^\times / \{\sigma^{-1}(\xi) \mid \xi \in \mathcal{P}\} \cong \mathcal{P}. \quad \square$$

We now have all ingredients to prove the following key result that is essential for the proof our main theorem:

Theorem 3.5.23. *The tangle space Θ of any graph G is homeomorphic to the quotient G^* / \approx_* of the Stone-Ćech remainder G^* of G , where each connected component of G^* is collapsed to a single point.*

Proof of Theorem 3.5.23. Theorem 3.5.17, Proposition 3.5.22 and Theorem 3.3.3 yield

$$G^* / \approx_* \cong G^\times / \approx_\times \cong \mathcal{P} \cong \Theta. \quad \square$$

3. Tangles and the Stone-Ćech compactification of infinite graphs

We write τ_* for the component of G^* corresponding to τ and τ_\times for the component $\tau_* \cap G^\times$ of G^\times corresponding to τ (cf. Theorem 3.5.23 and Lemma 3.5.14).

Theorem 3.5.24. *If τ is an \aleph_0 -tangle of G , then*

- (i) $\tau_* = G^* \cap \bigcap_{\gamma \in \Gamma} \overline{G[\tau, \gamma]}^{\beta G}$ and
- (ii) $\tau_\times = \bigcap_{\gamma \in \Gamma} \overline{\mathcal{C}(\tau, \gamma)}^{\beta G} = G^\times \cap \bigcap_{\gamma \in \Gamma} \overline{G[\tau, \gamma]}^{\beta G} = \tau_* \cap G^\times$

are the components of G^* and G^\times corresponding to τ respectively.

In statement (i) of the theorem, the intersection with G^* is really necessary—we will see the reason for this in Proposition 3.7.3.

Proof of Theorem 3.5.24. We show (ii) first. The first equality is evident from the definition of σ , and the centre equality follows from Lemma 3.5.7 with

$$G^\times = \bigcap_{\gamma \in \Gamma} (\beta G \setminus \mathcal{O}_{\beta G}[X(\gamma)]).$$

(i). By Corollary 3.5.15, the right-hand side contains at most one connected component of G^* . We have $\tau_\times \subseteq \overline{G[\tau, \gamma]}$ for all $\gamma \in \Gamma$ by (ii), so $\tau_* = \hat{\tau}_\times \subseteq \overline{G[\tau, \gamma]}$ holds for all γ as well (see Lemma 3.5.8), finishing the proof. \square

3.6. Obtaining the tangle compactification from the Stone-Ćech compactification

Now that we know $\Theta \cong G^*/\approx_*$, our next target is the proof of our main result, Theorem 3.1. For this, recall that $\mathcal{O}_{\beta G}(X, \mathcal{C}) = \overline{G[X, \mathcal{C}]}^{\beta G} \setminus G[X]$, and that Lemma 3.5.8 and Theorem 3.5.24 ensure that $\mathcal{O}_{\beta G}(X, \mathcal{C})$ is \approx_* -closed and includes precisely the components τ_* of G^* with $\mathcal{C} \in U(\tau, X)$.

Lemma 3.6.1. *Let $A \subseteq \beta G$ be closed and \approx_* -closed, and let τ be an \aleph_0 -tangle of G . If A avoids τ_* , then there are $X \in \mathcal{X}$ and $\mathcal{C} \subseteq \mathcal{C}_X$ with $\tau_* \subseteq \mathcal{O}_{\beta G}(X, \mathcal{C}) \subseteq \beta G \setminus A$.*

Proof. By the Separating Lemma 3.5.12 there is $X \in \mathcal{X}$ and a bipartition $\{\mathcal{C}_1, \mathcal{C}_2\}$ of \mathcal{C}_X with $A \subseteq \overline{G[X, \mathcal{C}_1]}$ and $\overline{\tau_*} \subseteq \overline{G[X, \mathcal{C}_2]}$. Then $\tau_* \subseteq \mathcal{O}_{\beta G}(X, \mathcal{C}_2) \subseteq \beta G \setminus A$. \square

We write $\widehat{\beta G}$ for the topological space obtained from βG by declaring G to be open, and we write \hat{G} for the quotient $\widehat{\beta G}/\approx_*$. Since βG contains G as a subspace, all the open sets of G are open in $\widehat{\beta G}$ as well; and since \approx_* does not affect G , all the open sets of G are also open in \hat{G} . As a consequence, the open sets of βG plus the open sets of G form a basis for the topology of $\widehat{\beta G}$, yielding that

Lemma 3.6.2. *The open sets of $(\beta G)/\approx_*$ plus the open sets of G form a basis for the topology of \hat{G} .* \square

We define a bijection $\Psi: \hat{G} \rightarrow |G|_\Theta$ by letting it be the identity on G and letting it send each \approx_* -class τ_* to its corresponding \aleph_0 -tangle τ .

3. Tangles and the Stone-Čech compactification of infinite graphs

Lemma 3.6.3. *The map Ψ is continuous.*

Proof. Since the open sets of G are open in both $|G|_{\Theta}$ and \hat{G} , it suffices to show that the preimage of any $\mathcal{O}_{|G|_{\Theta}}(X, \mathcal{C})$ is open in \hat{G} , and it is:

$$\Psi^{-1}(\mathcal{O}_{|G|_{\Theta}}(X, \mathcal{C})) = \mathcal{O}_{\beta G}(X, \mathcal{C})/\approx_*. \quad \square$$

Lemma 3.6.4. *The map Ψ is closed.*

Proof. Let A be any closed subset of \hat{G} ; we show that $\Psi[A]$ is closed in $|G|_{\Theta}$. For this, let ξ be any point of $|G|_{\Theta} \setminus \Psi[A]$, and let \mathcal{B} be the basis for the topology of \hat{G} provided by Lemma 3.6.2.

If ξ is a point of G , then we find an open neighbourhood O of ξ in G avoiding A since A is closed in \hat{G} . Then O witnesses $\xi \notin \overline{\Psi[A]}$ as well.

Otherwise ξ is an \aleph_0 -tangle $\tau \in \Theta \setminus \Psi[A]$. The set A is closed in \hat{G} , but it need not be closed in $(\beta G)/\approx_*$. Let us consider the closure B of A in $(\beta G)/\approx_*$ and show $B \setminus A \subseteq G$ (actually, one can even show that B adds only some vertices of infinite degree to A , but $B \setminus A \subseteq G$ suffices for our cause). Each point of $\hat{G} \setminus G$ that is not contained in A has an open neighbourhood from the basis \mathcal{B} avoiding A . Since all these neighbourhoods are not included in G , they must be open sets of $(\beta G)/\approx_*$, yielding $B \setminus A \subseteq G$. Therefore, the closed set $B' = \bigcup B$ of βG avoids the component τ_* of G^* corresponding to τ , and since B' is also \approx_* -closed our Lemma 3.6.1 yields $X \in \mathcal{X}$ and $\mathcal{C} \subseteq \mathcal{C}_X$ such that $\tau_* \subseteq \mathcal{O}_{\beta G}(X, \mathcal{C}) \subseteq \beta G \setminus B'$. Therefore, the open neighbourhood $\mathcal{O}_{|G|_{\Theta}}(X, \mathcal{C})$ of τ avoids $\Psi[A]$. \square

Theorem 3.1. *The tangle compactification $|G|_{\Theta}$ of any graph G is homeomorphic to the quotient $(\beta G, \tau')/\approx_*$ where τ' is the finer topology on βG obtained from βG by declaring G to be open in βG and then collapsing each connected component of G^* to a single point.*

Proof. Lemma 3.6.3 and Lemma 3.6.4 yield a homeomorphism. \square

3.7. Three observations about the Stone-Čech compactification

Given \mathbb{M}_E and an ultrafilter $U \in \beta E$ we write P_U for the collection of all points of \mathbb{I}_U that are of the form x_U for some family $(x_e \mid e \in E)$ of points $x_e \in \mathbb{I}_e$. By [48, Proposition 2.6], the set $P_U \setminus \{0_U, 1_U\}$ is dense in \mathbb{I}_U .

Theorem 3.7.1. *If G is an infinite graph that is not locally finite, then no compactification of G can both be Hausdorff and have a totally disconnected remainder.*

Proof. Suppose for a contradiction that αG is a Hausdorff compactification of G with totally disconnected remainder, and let v be a vertex of G of infinite degree. Consider a representation \mathbb{M}_E/\sim of G , so Theorem 3.4.9 yields $\beta G = (\beta \mathbb{M}_E)/\sim_{\beta}$ and we find a free ultrafilter $U \in E^*$ with $[0_U]_{\sim_{\beta}} = v$, say. The set $P_U \setminus \{0_U, 1_U\}$

3. Tangles and the Stone-Ćech compactification of infinite graphs

is dense in \mathbb{I}_U , so every open neighbourhood of v in βG meets $P_U \setminus \{0_U, 1_U\}$. In order to use this to derive a contradiction, we need to know more about αG first.

The Hausdorff compactification αG can be obtained from βG as a quotient $\beta G / \approx$ where \approx is an equivalence relation on G^* . Since αG has a totally disconnected remainder and since the (continuous) restriction of the quotient map to components of G^* preserves connectedness, we deduce that the equivalence relation \approx must refine \approx_* . Consequently, the connected subspace $\check{\mathbb{I}}_U$ of G^* (cf. Corollary 3.4.7) is included in a single \approx -class x , say. To yield a contradiction, it suffices to show that every open neighbourhood O of v in αG contains x . And indeed: if we view αG as the quotient $(\beta G) / \approx$ of βG , then $\bigcup O$ is open in βG and \approx -closed. Using that $\bigcup O$ meets $P_U \setminus \{0_U, 1_U\}$ and \approx refines \approx_* we deduce that $x \subseteq \bigcup O$, i.e. $x \in O$. \square

For our the second observation we need a short lemma and some notation: Since G is dense in βG , so is the locally compact subspace formed by the inner edge points and the vertices of finite degree, and hence [36, Theorem 3.3.9] yields:

Lemma 3.7.2. *If G is a graph, then $\mathring{E} \subseteq G$ is open in βG .* \square

Given an end ω of G we write $\Delta\omega$ for the set of those vertices dominating it. Our second observation describes explicitly how the connected components of the Stone-Ćech remainder of G interact with G .

Proposition 3.7.3. *Let G be any graph, and let \mathbb{M}_E / \sim be a representation of G .*

- (i) *If τ is an ultrafilter tangle of G , then $\overline{\tau}^{\beta G} = \tau_* \sqcup X_\tau$, and for each $x \in X_\tau$ there is an ultrafilter $U \in E^*$ with $[0_U]_{\sim_\beta} = x$, say, and with $\check{\mathbb{I}}_U \subseteq \tau_*$.*
- (ii) *If ω is an end of G , then $\overline{\omega}^{\beta G} = \omega_* \sqcup \Delta\omega$, and for each $x \in \Delta\omega$ there is an ultrafilter $U \in E^*$ with $[0_U]_{\sim_\beta} = x$, say, and with $\check{\mathbb{I}}_U \subseteq \omega_*$.*

Proof. (i). First, we show that $\overline{\tau}^{\beta G}$ avoids $G \setminus X_\tau$ (where G is the 1-complex). Since \mathring{E} is open in βG (Lemma 3.7.2) we may assume that $\overline{\tau}^{\beta G} \cap G \subseteq V$. Let v be any vertex of G that is not in X_τ , and let C be the (graph) component of $G - X_\tau$ with $v \in C$. Then $v \notin \overline{G[X_\tau, \mathcal{C}_{X_\tau} \setminus \{C\}]}$ implies $v \notin \overline{\tau}^{\beta G}$ by Theorem 3.5.24 as desired. Therefore, $\overline{\tau}^{\beta G} \cap G \subseteq X_\tau$.

Now suppose that any vertex $x \in X_\tau$ is given. Write Γ_x for the set of all $\gamma \in \Gamma$ with $x \in X(\gamma)$, and given $\gamma \in \Gamma_x$ put $F_\gamma = E(x, \bigcup \mathcal{C}(\tau, \gamma))$. The sets F_γ are infinite due to [62, Lemma 4.4]. We consider the filter on $E(x)$ that is given by the up-closure of the collection $\{F_\gamma \mid \gamma \in \Gamma_x\} \subseteq 2^{E(x)}$ (from the directedness of Γ_x it follows that this collection is directed by reverse inclusion, which is enough to ensure that we get a filter). Next, we extend this filter to an ultrafilter U on $E(G)$, and note that U must be free. Due to $E(x) \in U$ we may assume without loss of generality that there is some $F \in U$ with $F \subseteq E(x)$ and $\{0_e \mid e \in F\} \subseteq x$ where we view x as a \sim -class of \mathbb{M}_E . Then $0_U \in \overline{\{0_e \mid e \in F\}}^{\beta \mathbb{M}_E}$ implies $[0_U]_{\sim_\beta} = x$ as a consequence of $\beta G = (\beta \mathbb{M}_E) / \sim_\beta$, Theorem 3.4.9. If we can show that $\check{\mathbb{I}}_U$ is included in $\overline{G[\tau, \gamma]}^{\beta G}$ for all $\gamma \in \Gamma_x$, then we are done since Γ_x is cofinal in Γ and τ_* can be written as the directed intersection $G^* \cap \bigcap_{\gamma \in \Gamma} \overline{G[\tau, \gamma]}$ (cf. Theorem 3.5.24).

3. Tangles and the Stone-Ćech compactification of infinite graphs

For this, let any $\gamma \in \Gamma_x$ be given. Since $\check{\mathbb{I}}_U \subseteq G^*$ is connected (cf. Corollary 3.4.7) and $G^* \cap \overline{G[\tau, \gamma]}$ is clopen in G^* (cf. Lemma 3.5.8), it suffices to show that $\check{\mathbb{I}}_U$ meets $\overline{G[\tau, \gamma]}$ in $(\frac{1}{2})_U$. And indeed we have

$$(\frac{1}{2})_U \in \overline{\{(\frac{1}{2})_e \mid e \in F_\gamma\}}^{\beta M_E}$$

which implies $(\frac{1}{2})_U \in \overline{G[\tau, \gamma]}$ as desired.

(ii). This is proved similar to (i), where to show $\overline{\omega_*^{\beta G}} \cap G \subseteq \Delta\omega$ we use that for every vertex v of G not dominating ω there is $X \in \mathcal{X}$ separating v from $C(X, \omega)$ in that $v \notin X \cup C(X, \omega)$ so in particular $v \notin \overline{G[X, \{C(X, \omega)\}]}^{\beta G} \supseteq \overline{\omega_*^{\beta G}}$. \square

For the study of locally finite connected graphs, the so-called *Jumping Arc Lemma* (cf. [26, Lemma 8.5.3]) plays an important role. By considering subcontinua of the Stone-Ćech compactification instead of arcs in the Freudenthal compactification, we obtain the following quite strong generalisation of this lemma:

Lemma 3.7.4 (Jumping ‘Arc’ Lemma for the Stone-Ćech compactification).

Let $F \subseteq E$ be a cut of G with sides V_1, V_2 .

- (i) If F is finite, then $\overline{G[V_1]} \oplus \overline{G[V_2]}$ is a clopen bipartition of $(\beta G) \setminus \overset{\circ}{F}$, and there is no subcontinuum of $(\beta G) \setminus \overset{\circ}{F}$ meeting both V_1 and V_2 .
- (ii) If F is infinite, then $(\beta G) \setminus \overset{\circ}{F}$ might contain a subcontinuum meeting both V_1 and V_2 . This is the case, e.g., if both $G[V_1]$ and $G[V_2]$ are connected.

Moreover, two vertices of G lie in the same component (subcontinuum) of $(\beta G) \setminus \overset{\circ}{E}$ if and only if they lie on the same side of every finite cut of the graph G .

Proof. (i) is immediate from Theorem 3.4.1 (v). For (ii), note that if both $G[V_1]$ and $G[V_2]$ are connected then F is a bond so $(\beta G) \setminus \overset{\circ}{F} = \text{cl}_{\beta G}(G \setminus \overset{\circ}{F})$ is a continuum for every finite $F' \subseteq F$ by Lemmas 3.4.3, 3.7.2 and Theorem 3.4.1 (v). Hence $(\beta G) \setminus \overset{\circ}{F}$ is also a continuum as directed intersection of the continua $(\beta G) \setminus \overset{\circ}{F}'$, see Lemma 3.4.4.

Finally, note that, by (i), for the ‘moreover’ part it suffices to show the backward direction. For this, find infinitely many edge-disjoint paths P_0, P_1, \dots between the two vertices inductively, and note that by Lemmas 3.4.3, 3.4.4 and 3.7.2 the intersection

$$\bigcap_{n \in \mathbb{N}} \overline{\bigcup_{m > n} P_m}^{\beta G} \subseteq (\beta G) \setminus \overset{\circ}{E}$$

is a continuum containing the two vertices as desired. \square

4. A tree-of-tangles theorem for infinite-order tangles

4.1. Introduction

The *tree-of-tangles theorem*, one of the cornerstones of Robertson and Seymour’s proof of their graph-minor theorem, says (in the terminology of [26, §12.5]):

Theorem. *Every finite graph G has a nested set of separations which efficiently distinguishes all the finite-order tangles in G that can be distinguished.*

This is Theorem 12.5.4 in [26], the original article is [73].

Recently, Carmesin [19] has extended the tree-of-tangles theorem to the infinite-order tangles of infinite graphs that are locally finite. The precise statement of Carmesin’s result reads:

Theorem. *Every infinite connected graph G has a nested set of separations which efficiently distinguishes all the ends of G .*

Note that, in the wording of his theorem, Carmesin does not require the graph to be locally finite, and he speaks of ends where one expects infinite-order tangles. This is because his result is more general than an extension of the tree-of-tangles theorem to the infinite-order tangles of locally finite infinite graphs. To understand the difference, let us look at how the ends of a graph are related to its infinite-order tangles.

An end ω of a graph G (see [26]) orients every finite-order separation $\{A, B\}$ of G towards the side that contains a tail from every ray in ω . Since these orientations are, for distinct separations, consistent in a number of ways, they form an infinite-order tangle of G . Conversely, every infinite-order tangle of a locally finite and connected graph G is defined by an end in this way [25, 30]. Thus, if G is locally finite and connected, there is a canonical bijection between its infinite-order tangles and its ends. In this way, Carmesin’s result extends the tree-of-tangles theorem to the infinite-order tangles of locally finite graphs.

When G is not locally finite, however, there can be infinite-order tangles that are not defined by an end. Then Carmesin’s result no longer extends the tree-of-tangles theorem to the infinite-order tangles of G .

The infinite-order tangles that do not come from ends of the graph are fundamentally different from ends. They are closely related to free ultrafilters, and are called *ultrafilter tangles* [25]. More explicitly, by a recent result from [62], there is a canonical bijection between the ultrafilter tangles and the *ultrafilter tangle blueprints*: pairs (X, U) of a critical vertex set X and a free ultrafilter U on \mathcal{C}_X , where a finite set $X \subseteq V(G)$ is *critical* if the collection \mathcal{C}_X of the components of $G - X$ whose neighbourhood is equal to X is infinite. Therefore, every ultrafilter tangle $\tau = (X, U)$ has two aspects: Its combinatorial aspect is captured by its blueprint’s critical vertex set X , and its ultrafilter aspect is encoded by the free ultrafilter U (see Section 4.2.2 for details). Since every vertex in a critical vertex set

4. A tree-of-tangles theorem for infinite-order tangles

has infinite degree, it follows that locally finite connected graphs have no ultrafilter tangles, so all their infinite-order tangles are ends.

Ultrafilter tangles are interesting also for topological reasons. Every locally finite connected graph can be naturally compactified by its ends to form its well known end compactification [26] introduced by Freudenthal [37]. But for a non-locally finite graph, adding its ends no longer suffices to compactify it. Adding its ends *plus its ultrafilter tangles*, however, (i.e. adding all its infinite-order tangles) does again compactify the graph. This is Diestel’s *tangle compactification* [25]. The tangle compactification generalises the end compactification twofold. On the one hand, it defaults to the end compactification when the graph is locally finite and connected. And on the other hand, the relation between the end compactification of locally finite connected graphs and their Stone-Ćech compactification extends to all graphs when ends are generalised to tangles; see Chapter 3.

As our main result, we extend Robertson and Seymour’s tree-of-tangles theorem to the infinite-order tangles of infinite graphs (and thus, we extend Carmesin’s result from ends to all infinite-order tangles):

Theorem 4.1. *Every infinite connected graph G has a nested set of finite-order separations that efficiently distinguishes all the inequivalent infinite-order tangles of G and is oriented in the same way by equivalent infinite-order tangles.*

Here, two ultrafilter tangles are *equivalent* if their blueprints’ critical vertex sets coincide. Therefore, our nested set of separations distinguishes precisely those ultrafilter tangles that differ in their combinatorial aspects. As we will show, our result is best possible in the following sense. If a graph G has an ultrafilter tangle τ , then no nested set of finite-order separations of G efficiently distinguishes all the ultrafilter tangles of G that are equivalent to τ .

Applications

Our work has four applications.

Elbracht, Kneip and Teegen need it in their paper [34]. So do Burger and the second author in Chapter 8.

Our third application is the following structural connectivity result for infinite graphs, which generalises the way in which the cutvertices of a graph decompose it into its blocks in a tree-like fashion. Call a graph *tough* if deleting finitely many vertices from it never leaves more than finitely many components. By the pigeonhole principle a graph is tough if and only if it has no critical vertex set.

Theorem 4.2. *Every connected graph G has a nested set of separations whose separators are precisely the critical vertex sets of G and all whose torsos are tough.*

(See Section 4.2.3 for definitions.)

Theorem 4.2 is interesting also from the perspective of topological infinite graph theory, in view of the following two results. Diestel and Kuhn [24] showed that a graph is compactified by its ends if and only if it is tough (i.e., if and only if it has no critical vertex sets), and in [62] it was shown that every graph is compactified

4. A tree-of-tangles theorem for infinite-order tangles

by its ends plus critical vertex sets. So a graph is compactified by points that come in two types, ends and critical vertex sets, and the second type decomposes the graph into a nested set of separations all whose torsos are compactified by the points of the first type.

Our fourth application answers a question that arises from the work of Polat and of Sprüssel. End spaces of graphs, in general, are not compact. However, Polat [68] and Sprüssel [80] independently showed that end spaces of graphs are normal. Polat even showed that end spaces of graphs are collectionwise normal, which is stronger than normal but weaker than compact Hausdorff. (In a *collectionwise* normal space one can at once pairwise separate any *collection* of closed disjoint sets with disjoint open neighbourhoods, cf. Definition 4.6.2.)

The infinite-order tangle space, endowed with the subspace topology of the tangle compactification, contains the end space as a subspace. As Diestel [25] showed, the infinite-order tangle space is compact Hausdorff, which implies collectionwise normality by general topology.

The ultrafilter tangle space, endowed with the subspace topology of the infinite-order tangle space, is not usually compact. Since the infinite-order tangle space is the disjoint union of the end space and the ultrafilter tangle space, the question arises whether the ultrafilter tangle space is collectionwise normal as well. We answer this question in the affirmative:

Theorem 4.3. *The ultrafilter tangle space of a graph is collectionwise normal.*

Our chapter is organised as follows. Background knowledge is supplied in Section 4.2. In Section 4.3 we study examples and show that our main result is best possible. In Section 4.4 we give an overview on our overall proof strategy. Our main technical results are stated and proved in Section 4.5. In Section 4.6 we provide the applications of our main technical results. In Section 4.7 we introduce an equivalence relation on a tree set given a consistent orientation of that tree set. This is the foundation for the definition of the modified torsos and proxies as well as for a ‘lifting’ process that we need in Section 4.8. In Section 4.8, finally, we introduce the modified torsos and prove our main result.

Throughout this chapter, $G = (V, E)$ is any connected graph.

Acknowledgement. We are grateful to Nathan Bowler, Christian Elbracht, Konstantinos Stavropoulos and Maximilian Teegen for stimulating discussions that contributed to the genesis of this chapter.

4.2. Ends, tangles and tree sets

4.2.1. Definition of \aleph_0 -tangles

In this subsection we formally introduce tangles for a particular type of separation system. More precisely, we introduce a definition of \aleph_0 -tangles provided by Diestel [25] which, as he proved, is equivalent to the original one due to Robertson and Seymour [73]. A more detailed summary of [25] that does not rely on [23] can be found in [62].

4. A tree-of-tangles theorem for infinite-order tangles

The *interior* of a star $\{(A_i, B_i) \mid i \in I\} \subseteq \vec{S}_{\aleph_0}$ is the intersection $\bigcap_{i \in I} B_i$.

Definition 4.2.1. An \aleph_0 -*tangle* (of G) is a consistent orientation of S_{\aleph_0} that contains no finite star of finite interior as a subset. We write Θ for the set of all \aleph_0 -tangles.

4.2.2. Properties of \aleph_0 -tangles

If ω is an end of G , then letting

$$\tau_\omega := \{ (A, B) \in \vec{S}_{\aleph_0} \mid C(A \cap B, \omega) \subseteq G[B \setminus A] \}$$

defines an injection $\Omega \hookrightarrow \Theta$, $\omega \mapsto \tau_\omega$. The \aleph_0 -tangles of the form τ_ω are called *end tangles*. All other \aleph_0 -tangles are *ultrafilter tangles*. For a better explanation of ultrafilter tangles we need some notation first.

Given a subset $X \subseteq V(G)$ we write \mathcal{C}_X for the collection of components of $G - X$, and moreover if Y is a subset of X we write $\mathcal{C}_X(Y)$ for the set $\{C \in \mathcal{C}_X \mid N(C) = Y\}$ of components of $G - X$ that have their neighbourhood precisely equal to Y . In the special case of $X = Y$ we abbreviate $\mathcal{C}_X(X)$ to $\check{\mathcal{C}}_X$.

In this chapter, partition classes are required to be non-empty as usual, with the exception that whenever we speak of a *bipartition* we do not formally mean a partition and allow for at most one empty class. Now if τ is an \aleph_0 -tangle of the graph, then for every $X \in \mathcal{X}$ it chooses one *big side* from each bipartition $\{\mathcal{C}, \mathcal{C}'\}$ of \mathcal{C}_X , namely the $\mathcal{D} \in \{\mathcal{C}, \mathcal{C}'\}$ with $(X, \mathcal{D}) \in \tau$. Since it chooses these sides consistently, it induces an ultrafilter

$$U(\tau, X) = \{ \mathcal{C} \subseteq \mathcal{C}_X \mid (X, \mathcal{C}) \in \tau \}$$

on \mathcal{C}_X , one for every $X \in \mathcal{X}$. Diestel showed that the map

$$\tau \mapsto (U(\tau, X) \mid X \in \mathcal{X})$$

is a bijection between Θ and the inverse limit $\varprojlim \beta(\mathcal{C}_X)$; see the previous chapter. He also showed that the end tangles are precisely the \aleph_0 -tangles with all induced ultrafilters principal. Consequently, an ultrafilter tangle induces for some $X \in \mathcal{X}$ a free ultrafilter on \mathcal{C}_X , and Diestel showed that each of these free ultrafilters alone determines that tangle. Moreover, he showed that for an ultrafilter tangle τ the collection \mathcal{X}_τ of all $X \in \mathcal{X}$ with $U(\tau, X)$ free does have a least element X_τ of which it is the up-closure. In [62], these insights have been employed to yield a new view on ultrafilter tangles, as follows.

An *ultrafilter tangle blueprint* is a pair (X, U) of a critical vertex set X and a free ultrafilter U on $\check{\mathcal{C}}_X$, where a set $X \in \mathcal{X}$ is *critical* if $\check{\mathcal{C}}_X$ is infinite. Then by a recent result, [62, Theorem 4.10], the map

$$\tau \mapsto (X_\tau, U(\tau, X_\tau) \cap 2^{\check{\mathcal{C}}_{X_\tau}})$$

is a bijection between the ultrafilter tangles and the ultrafilter tangle blueprints. In particular, every ultrafilter tangle τ contains $(X_\tau, \check{\mathcal{C}}_{X_\tau})$, the separation with which τ naturally comes as stated in the introduction. Here is the inverse of the bijection:

4. A tree-of-tangles theorem for infinite-order tangles

Theorem 4.2.2 ([25, Theorem 3.5]). *For every ultrafilter tangle τ and each $X \in \mathcal{X}_\tau$ the free ultrafilter $U(\tau, X)$ determines τ in that*

$$\tau = \{ (A, B) \in \vec{S}_{\aleph_0} \mid \exists \mathcal{C} \in U(\tau, X) : V[\mathcal{C}] \subseteq B \setminus A \}.$$

We will need the following notation from [62] for critical vertex sets. For every $X \in \mathcal{X}$ and all critical Y that are not entirely contained in X we write $C_X(Y)$ for the unique component of $G - X$ meeting Y (equivalently: including $\bigcup \mathcal{C}_{X \cup Y}(Y)$). The collection of all critical vertex sets of G is denoted by $\text{crit}(G)$.

Lemma 4.2.3 ([62, Lemma 4.8]). *For every ultrafilter tangle τ and each $X \in \mathcal{X} \setminus \mathcal{X}_\tau$ we do have $X_\tau \subseteq X \cup C_X(X_\tau)$ and the ultrafilter $U(\tau, X)$ is generated by $\{C_X(X_\tau)\}$.*

The following lemma will be useful:

Lemma 4.2.4 ([25, Lemma 1.10]). *Let τ be an \aleph_0 -tangle of G and $(A, B) \in \tau$. Let (A', B') be a separation of G with $A \Delta A'$ and $B \Delta B'$ finite. Then $(A', B') \in \tau$.*

We say that two ultrafilter tangles τ_1, τ_2 of G are *equivalent* and write $\tau_1 \sim \tau_2$ if $X_{\tau_1} = X_{\tau_2}$. Identifying all equivalent \aleph_0 -tangles yields the quotient $\Theta/\sim = \Omega \sqcup \text{crit}(G)$ which is yet again a tangle space. For this, we need the concept of tame finite-order separations:

Definition 4.2.5. A finite-order separation $\{X, \mathcal{C}\}$ of G and its orientations are *tame* if for no $Y \subseteq X$ both $\mathcal{C}_X(Y) \cap \mathcal{C}$ and $\mathcal{C}_X(Y) \cap (\mathcal{C}_X \setminus \mathcal{C})$ are infinite.

We write S_t for the set of all tame finite-order separations of G . By [62, Theorem 5.10], the \aleph_0 -tangles of S_t , infinite-order tangles that only orient the separations in S_t , correspond precisely to the ends plus critical vertex sets which, in turn, correspond precisely to the elements of the quotient Θ/\sim . For details, see Section 5 of [62]. When we construct the tree set for our main result, the following fact will be useful:

Observation 4.2.6. *If T is a set of tame finite-order separations of G , then equivalent \aleph_0 -tangles induce the same orientation on T (i.e. ‘live’ in the same part).*

4.2.3. Tree sets

A tree set of G is a tree set of separations of G with the usual partial ordering and involution. If T is a tree set of separations of G and O is a consistent orientation of T , then the intersection $\Pi = \bigcap \{ B \mid (A, B) \in O \}$ is called the *part* of O . And the graph that is obtained from $G[\Pi]$ by adding an edge xy whenever $x \neq y \in \Pi$ lie together in the separator of some separation of O is called the *torso* of O (or of Π if O is clear from context). We denote the torso of O by $\text{torso}(G, O)$.

We will need the following lemma and its corollaries (the lemma is folklore and has been proved, e.g., in [19]; we present an alternative proof for convenience):

4. A tree-of-tangles theorem for infinite-order tangles

Lemma 4.2.7. *If Π is a part of a tree set of G , then for every $G[\Pi]$ -path P there is some separation of the tree set whose separator contains both endvertices of P .*

Proof. Let O be any consistent orientation of a tree set of G , write Π for its part and suppose that $P = xv_1 \dots v_n y$ is a $G[\Pi]$ -path (so $n \geq 1$). For every $k \in [n]$ pick an oriented separation $(A_k, B_k) \in O$ with $v_k \in A_k \setminus B_k$ (so that (A_k, B_k) witnesses $v_k \notin \Pi$). Let N consist of the \leq -maximal separations from the collection $\{(A_k, B_k) \mid k \in [n]\}$. Then for every v_k there is a separation $(A, B) \in N$ with $v_k \in A \setminus B$. Our aim is to show that N is a singleton, since then the separator of the sole separation in N must contain both x and y , so we would be done. By the choice of N , every two oriented separations in N are \leq -incomparable. As O is a consistent orientation of a tree set, this means that N must be a star. Then $|N| = 1$ is evident, since otherwise the sides $G[A \setminus B]$ for $(A, B) \in N$ altogether induce a disconnection of the subpath $v_1 \dots v_n$ of P contradicting its connectedness. \square

Corollary 4.2.8. *If Π is a part of a tree set of G and ω is an end of G in the closure of Π while $G[\Pi]$ coincides with the torso of Π , then ω has a ray in $G[\Pi]$.*

Proof. If ω lies in the closure of Π , we find a comb in G with its spine R in ω and all of its teeth in Π . Without loss of generality the comb meets Π precisely in its teeth. Then, as $G[\Pi]$ coincides with the torso of Π , it has an edge between every two consecutive teeth by Lemma 4.2.7, and so contains a ray equivalent to R . \square

Corollary 4.2.9. *If Π is a part of a tree set of G and two rays of $G[\Pi]$ are equivalent in G , then they are equivalent in the torso of Π as well.*

Proof. Given two rays of $G[\Pi]$ that are equivalent in G , we inductively construct infinitely many pairwise vertex-disjoint paths in G between them, and then employ Lemma 4.2.7 to turn these into paths of the torso. \square

The next corollary has already been known to Carmesin [19]:

Corollary 4.2.10. *The intersection of a connected set of vertices of G with a part of a tree set of separations of G induces a connected subgraph of the part's torso.* \square

Finally, we state Carmesin's result that implies the tree-of-tangles theorem for the infinite-order tangles of locally finite infinite graphs:

Theorem 4.2.11 ([19, Corollary 5.17]). *Every connected graph G has a tree set of finite order separations of G that efficiently distinguishes all the ends of G .*

Very recently, Carmesin, Hamann and Miraftab [20] showed a canonical version of this theorem, see their paper for definitions.

For more on tree sets, see e.g. [29] or [53].

4.3. Example section

The aim of this section is twofold. First, we verify that our main result, Theorem 4.1, is indeed best possible as claimed in the introduction. More precisely, in Subsection 4.3.1 we show that tree sets of finite-order separations cannot distinguish all the ultrafilter tangles from the same equivalence class at once—for any G .

Second, we study the candidate for a starting tree set that is formed by the separations $\{X, \check{\mathcal{C}}_X\}$ with X critical in G (recall that these are precisely the separations which naturally accompany the ultrafilter tangles). More precisely, in Subsection 4.3.2 we will see two example graphs showing that it is necessary to modify the tree set candidate: For the first example graph, the separations $\{X, \check{\mathcal{C}}_X\}$ form a tree set but do not distinguish any two ultrafilter tangles at all. For the second example graph, the separations $\{X, \check{\mathcal{C}}_X\}$ are not even nested.

4.3.1. Ultrafilters and tree sets

In this subsection we show that, as soon as a graph G has some ultrafilter tangle τ , it already cannot admit a tree set of finite-order separations that distinguishes all the ultrafilter tangles that are equivalent to τ . As our first step, we translate the problem from graphs to bipartitions of sets.

For this, we need to make some things formal first. Suppose that K is a non-empty set. We let $\vec{\mathcal{B}}(K) := 2^K$. Thus, every subset of K is an oriented ‘separation’. The partial ordering \leq of $\vec{\mathcal{B}}(K)$ will be \supseteq , the involution $*$ on $\vec{\mathcal{B}}(K)$ will be complementation in the set K . If desired, we can think of a separation $Z \subseteq K$ as the oriented bipartition (Z^*, Z) of K , and then $\mathcal{B}(K)$ is the set of bipartitions of K . Note that two separations $Z_1, Z_2 \in \vec{\mathcal{B}}(K)$ are nested if $Z_1 \subseteq Z_2$ or $Z_1 \supseteq Z_2$ or $Z_1 \cup Z_2 = K$ or $Z_1 \cap Z_2 = \emptyset$. A *tree set of bipartitions* of K is a tree set contained in $\vec{\mathcal{B}}(K)$ with the induced partial ordering and involution. Note that \emptyset is the sole small separation in $\vec{\mathcal{B}}(K)$ for $Z \subseteq K \setminus Z$ implies $Z = \emptyset$. Since an ultrafilter on K happens to be an orientation of $\mathcal{B}(K)$, a tree set \vec{T} of bipartitions of K distinguishes two distinct ultrafilters $U \neq U'$ on K if there is some $Z \in \vec{T}$ with $Z \in U$ and $Z^* \in U'$. We are almost ready for the translation, we only need one more lemma:

Lemma 4.3.1. *Let τ be any ultrafilter tangle of G with blueprint (X, U) and let any separation $(Y, \mathcal{D}) \in \tau$ be given. Write \mathcal{C} for the set of those components in $\check{\mathcal{C}}_X$ that avoid Y .*

- (i) *If Y includes X , then $(Y, \mathcal{D}) \leq (X, \mathcal{D} \cap \check{\mathcal{C}}_X) \in \tau$.*
- (ii) *Otherwise $(Y, \mathcal{D}) \leq (X, \mathcal{C}) \in \tau$.*

In particular, the set $\{(X, \mathcal{C}) \mid \mathcal{C} \in U\}$ is cofinal in τ .

Proof. Since Y is finite, \mathcal{C} is a cofinite subset of $\check{\mathcal{C}}_X$, giving $\mathcal{C} \in U$.

(i) The intersection $\mathcal{D} \cap \check{\mathcal{C}}_X$ can be written more complicated as $(\mathcal{D} \upharpoonright X) \cap \mathcal{C}$ where $\mathcal{C} \in U$ as noted above and

$$\mathcal{D} \upharpoonright X := \{C \in \check{\mathcal{C}}_X \mid \exists D \in \mathcal{D} : D \supseteq C\}$$

4. A tree-of-tangles theorem for infinite-order tangles

is known to be contained in U by [25, Lemma 2.2]. Hence $(X, \mathcal{D} \cap \check{\mathcal{C}}_X) \in \tau$. It is straightforward to check $(Y, \mathcal{D}) \leq (X, \mathcal{D} \cap \check{\mathcal{C}}_X)$.

(ii) From $\mathcal{C} \in U$ we get $(X, \mathcal{C}) \in \tau$. Lemma 4.2.3 deduces from $(Y, \mathcal{D}) \in \tau$ that $C_Y(X) \in \mathcal{D}$. Finally, we calculate $(Y, \mathcal{D}) \leq (Y, C_Y(X)) \leq (X, \mathcal{C})$ where for the second inequality we use that every component in \mathcal{C} sends an edge to the non-empty $X \setminus Y \subseteq C_Y(X)$ to deduce $\bigcup \mathcal{C} \subseteq C_Y(X)$. \square

Now we are ready for the translation:

Lemma 4.3.2. *Let X be a critical vertex set of G . Then every tree set of finite-order separations of G that distinguishes all the ultrafilter tangles τ of G with $X_\tau = X$ does induce a tree set of bipartitions of $\check{\mathcal{C}}_X$ that distinguishes all the free ultrafilters on $\check{\mathcal{C}}_X$.*

Proof. Let \vec{T} be a tree set of finite order separations of G that distinguishes all the ultrafilter tangles of G with $X_\tau = X$. Without loss of generality every separation $(Y, \mathcal{D}) \in \vec{T}$ distinguishes some two such ultrafilter tangles, and so $X \subseteq Y$ follows for all $(Y, \mathcal{D}) \in \vec{T}$.

The candidate for a tree set of bipartitions of $\check{\mathcal{C}}_X$ is $\{\bar{\mathcal{D}} \mid (Y, \mathcal{D}) \in \vec{T}\}$ where $\bar{\mathcal{D}} = \mathcal{D} \cap \check{\mathcal{C}}_X$. But when (Y, \mathcal{D}') is the inverse of (Y, \mathcal{D}) it can happen that $\bar{\mathcal{D}}'$ is not the inverse of $\bar{\mathcal{D}}$ in $\check{\mathcal{B}}(\check{\mathcal{C}}_X)$. For example, this happens when a finite component $C \in \check{\mathcal{C}}_X$ is contained in Y , for then both $\bar{\mathcal{D}}'$ and $\bar{\mathcal{D}}$ are missing C .

We overcome this obstacle as follows. First, we choose any consistent orientation O of \vec{T} (such an orientation exists, e.g. by [23, Lemma 4.1] which essentially applies Zorn's lemma to find an inclusionwise maximal partial orientation). Then, we define $N_O := \{\mathcal{D} \mid (\mathcal{D}, Y) \in O\}$. Finally, we claim that $\vec{N} := N_O \cup N_O^*$ is a tree set of bipartitions of $\check{\mathcal{C}}_X$ that distinguishes all the free ultrafilters on $\check{\mathcal{C}}_X$.

To verify that \vec{N} is a tree set we show that \vec{N} is nested. For this, consider any two separations $(\mathcal{D}_1, Y_1), (\mathcal{D}_2, Y_2) \in O$. Then, say, either $(\mathcal{D}_1, Y_1) \leq (\mathcal{D}_2, Y_2)$ implies $\bar{\mathcal{D}}_1 \subseteq \bar{\mathcal{D}}_2$ or $(\mathcal{D}_1, Y_1) \leq (Y_2, \mathcal{D}_2)$ implies $\bar{\mathcal{D}}_1 \subseteq (\bar{\mathcal{D}}_2)^*$. So \vec{N} is a tree set.

Now let $U \neq U'$ be any distinct two free ultrafilters on $\check{\mathcal{C}}_X$. Then there is a separation $(\mathcal{D}, Y) \in O$ that distinguishes the ultrafilter tangles τ_U and $\tau_{U'}$ corresponding to (X, U) and (X, U') , say with $(\mathcal{D}, Y) \in \tau_U$ and $(Y, \mathcal{D}) \in \tau_{U'}$. By Lemma 4.3.1 we have $\bar{\mathcal{D}} \in U'$. Similarly $\bar{\mathcal{D}}_Y \setminus \bar{\mathcal{D}} \in U$, which then via the inclusion $\bar{\mathcal{D}}_Y \setminus \bar{\mathcal{D}} \subseteq (\bar{\mathcal{D}})^*$ implies $(\bar{\mathcal{D}})^* \in U$. \square

As a consequence of this lemma, it suffices to show

Theorem 4.3.3. *If K is an infinite set, then no tree set of bipartitions of K distinguishes all the free ultrafilters on K .*

in order to obtain our desired result:

Corollary 4.3.4. *If τ is an ultrafilter tangle of G , then no tree set of finite-order separations of G distinguishes all the ultrafilter tangles that are equivalent to τ . \square*

Theorem 4.3.3 above has been proved independently from us by Bowler [6] in 2014 who did not publish his findings. The proof presented below is ours. For

4. A tree-of-tangles theorem for infinite-order tangles

the proof we need the following lemma which is a tree set version of the fact that every connected infinite graph contains either a ray or a vertex of infinite degree, [26, Proposition 8.2.1]:

Lemma 4.3.5. *Every regular infinite tree set contains either an ω -chain or an infinite splitting star.*

Proof. If a tree set contains no ω -chain, then it is isomorphic to the edge tree set of a rayless tree by Kneip's Theorem 2.3.1. This tree, then, must have an infinite degree vertex if the tree set is infinite. \square

If U is an ultrafilter on a set K and \mathcal{K} is a partition of K , then we write $U \cdot \mathcal{K}$ for the *induced* ultrafilter on \mathcal{K} given by $\{\mathcal{A} \subseteq \mathcal{K} \mid \bigcup \mathcal{A} \in U\}$. Notably, if U is principal, then so is $U \cdot \mathcal{K}$. Conversely, every ultrafilter \mathcal{U} on \mathcal{K} gives a filter

$$[\{\bigcup \mathcal{A} \mid \mathcal{A} \in \mathcal{U}\}]_K := \{A \subseteq K \mid \exists \mathcal{A} \in \mathcal{U} : A \supseteq \bigcup \mathcal{A}\}$$

on K , and every ultrafilter U on K that extends this filter induces \mathcal{U} in that $\mathcal{U} = U \cdot \mathcal{K}$. Phrased differently, the map $U \mapsto U \cdot \mathcal{K}$ is a surjection from the set of ultrafilters on K onto the set of ultrafilters on \mathcal{K} . Notably, free ultrafilters on \mathcal{K} are induced only by free ultrafilters on K .

Proof of Theorem 4.3.3. Let any infinite set K be given and assume for a contradiction that \vec{T} is a tree set of bipartitions of K that distinguishes all the free ultrafilters on K . If \vec{T} is finite, then there are only finitely many orientations of \vec{T} . But there are infinitely many free ultrafilters on K , so a finite tree set cannot possibly distinguish all of them. Therefore, \vec{T} must be infinite. Since the empty set does not distinguish any two ultrafilters on K we may assume without loss of generality that \vec{T} is regular. Then by Lemma 4.3.5 we know that \vec{T} contains either an ω -chain or an infinite splitting star.

Suppose first that \vec{T} contains an ω -chain; that is to say that we find a sequence $(Z_n)_{n < \omega}$ in \vec{T} with $Z_n \supsetneq Z_{n+1}$ for all n . As \vec{T} is a tree-set, $K \setminus Z_0$ is non-empty. Put $Z_\omega := \bigcap_{n < \omega} Z_n$. Then Z_ω is nested with every separation in \vec{T} . More precisely, every separation in T has an orientation Z such that either $Z \supseteq Z_n$ for some $n < \omega$ or $Z_\omega \supseteq Z$. We turn the transfinite sequence $(Z_\alpha)_{\alpha \leq \omega}$ into a partition of K , as follows. For every $n < \omega$ set $K_n = Z_n \setminus Z_{n+1}$; and put $K_\omega := (K \setminus Z_0) \cup Z_\omega$. Then $\mathcal{K} := \{K_\alpha \mid \alpha \leq \omega\}$ is an infinite partition of K . Let \mathcal{U} be any free ultrafilter on \mathcal{K} , and pick some free ultrafilter U on K with $\mathcal{U} = U \cdot \mathcal{K}$. The free ultrafilter \mathcal{U} contains all cofinite subsets $\{K_m \mid n \leq m < \omega\} \subseteq \mathcal{K}$ with $n < \omega$, and so U contains all $Z_n \setminus Z_\omega$ with $n < \omega$. Recall that every separation in T has an orientation Z such that either $Z \supseteq Z_n$ for some $n < \omega$ or $Z_\omega \supseteq Z$. Hence for every separation $\{Z^*, Z\} \in T$ we have that either $Z \supseteq Z_n$ with $Z_n \setminus Z_\omega \in U$ implies $Z \in U$, or $Z_\omega \supseteq Z$ with $Z_0 \setminus Z_\omega \in U$ implies $Z^* \in U$. Therefore, if \mathcal{U}' is any free ultrafilter on \mathcal{K} other than \mathcal{U} , and U' is a free ultrafilter on K inducing \mathcal{U}' , then U' orients every separation in T the same way as U . But then \vec{T} does not distinguish U and U' from each other, a contradiction.

Finally suppose that T contains an infinite splitting star $\sigma = \{K_i \mid i \in I\}$. If $\mathcal{K} := \{K_i^* \mid i \in I\}$ is not yet a partition of K , then we add the non-empty interior

4. A tree-of-tangles theorem for infinite-order tangles

$\bigcap_{i \in I} K_i$ of σ to \mathcal{K} to turn \mathcal{K} into one. Let \mathcal{U} be any free ultrafilter on \mathcal{K} , and pick some free ultrafilter U on K inducing \mathcal{U} . The free ultrafilter \mathcal{U} contains all collections $\mathcal{K} - K_i^*$, and hence U contains all K_i . Now every separation in T has an orientation Z with $Z \supseteq K_i$ for some $i \in I$ as σ is splitting, and then $K_i \in U$ implies $Z \in U$. Therefore, if \mathcal{U}' is any free ultrafilter on \mathcal{K} other than \mathcal{U} , and U' is a free ultrafilter on K inducing \mathcal{U}' , then U' orients every separation in T the same way as U . But then \vec{T} does not distinguish U and U' , a contradiction. \square

We remark that the proof above even shows the following stronger version of Theorem 4.3.3: *If K is an infinite set, then for every tree set of bipartitions of K there is a collection of at least $2^{2^{\aleph_0}} = 2^c$ many free ultrafilters on K all of which induce the same orientation of the tree set.*¹ So if G has precisely one critical vertex set X with \mathcal{C}_X countable, then for every tree set of finite order separations of G there is a collection \mathcal{O} of ultrafilter tangles of G such that all ultrafilter tangles in \mathcal{O} induce the same orientation of the tree set and the cardinal $|\mathcal{O}|$ is equal to the total number 2^c of ultrafilter tangles of G .

4.3.2. The problem case

This subsection is dedicated to examples that show why we do need the function $X \mapsto \mathcal{K}(X)$ in our main result. More precisely, we will see two graphs whose critical vertex sets give a very bad starting set

$$\{ \{X, \mathcal{C}_X\} \mid X \in \text{crit}(G) \}.$$

In both cases, all the critical vertex sets interact with each other in a particular way, made precise as follows. Let us say that two critical vertex sets X and Y of G form a *problem case* if X and Y are incomparable as sets and additionally both $C_X(Y) \in \mathcal{C}_X$ and $C_Y(X) \in \mathcal{C}_Y$ hold.




Figure 4.3.1.: This graph's critical vertex sets give an infinite star of small separations

Example 4.3.6. If G is the graph shown in Figure 4.3.1, then the collection

$$\{ (X, \mathcal{C}_X) \mid X \in \text{crit}(G) \}$$

¹By improving Lemma 4.3.5 it might be possible to replace $2^{2^{\aleph_0}}$ with $2^{2^{|\mathcal{K}|}}$.

4. A tree-of-tangles theorem for infinite-order tangles

is an infinite star of small separations, as we shall show in a moment. As every \aleph_0 -tangle contains all the small separations (A, V) with A finite (because these can be written as (A, \mathcal{C}_A) and $\mathcal{C}_A \in U(\tau, A)$ for every \aleph_0 -tangle τ), it follows that every ultrafilter tangle contains this star as a subset, and so no two ultrafilter tangles are distinguished by this star's underlying tree set.

Before we take a closer look at the critical vertex sets of G , however, we describe G more precisely. For this, we define graphs G_n , one for each $n \in \mathbb{N}$, by letting G_n be a copy of K_{2, \aleph_0} with 2-class $\{x_n, a\}$, say, such that G_n meets all G_m with $n \neq m$ precisely in a . Then G is obtained from the union of all G_n by adding a new vertex b and joining it precisely to every x_n . Now $Y := \{a, b\}$ and the sets $X_n := \{x_n, a\}$ are critical, and these are all critical vertex sets. Moreover, we have $(Y, \mathcal{C}_Y) = (Y, V)$ and $(X_n, \mathcal{C}_{X_n}) = (X_n, V)$ with $(Y, V) \leq (V, X_n)$ and $(X_n, V) \leq (V, X_m)$. Notably, every two distinct critical vertex sets of G form a problem case. \square




Figure 4.3.2.: This graph's critical vertex sets do not give nested separations

Example 4.3.7. If G is the graph shown in Figure 4.3.2, then the collection

$$\{ \{X, \mathcal{C}_X\} \mid X \in \text{crit}(G) \}$$

is not even nested. Indeed, X and Y are the only two critical vertex sets of G . Write V' for $V - u$. Then $\{X, \mathcal{C}_X\} = \{X + u, V'\}$ and $\{Y, \mathcal{C}_Y\} = \{Y + u, V'\}$. Now these two separations cannot be nested: as X and Y are incomparable as sets, we have neither $(X + u, V') \leq (Y + u, V')$ nor $(V', X + u) \leq (V', Y + u)$. But $(X + u, V') \leq (V', Y + u)$ and $(V', X + u) \leq (Y + u, V')$ are impossible as well since $X + u$ and $Y + u$ are both incomparable with V' as sets. As in the previous example we note that X and Y form a problem case. \square

4.4. The overall proof strategy

Our overall strategy to achieve our main result, Theorem 4.1, roughly goes as follows. Let G be any infinite connected graph. Recall that every ultrafilter tangle $\tau = (X, U)$ of G naturally comes with a finite-order separation $(X, \mathcal{C}_X) \in \tau$. As our first step, we carefully extend and refine the set of these separations into a starting tree set T that already distinguishes all the inequivalent ultrafilter tangles of G , but does not necessarily do so efficiently.

4. A tree-of-tangles theorem for infinite-order tangles




Figure 4.4.1.: The separator Z of a separation efficiently distinguishing two inequivalent tangles τ_1 and τ_2 ; and a modified torso H with ends η_1 and η_2 representing the two tangles.

Next, we modify the torsos of T so that every \aleph_0 -tangle of G is represented in every modified torso by some end of that modified torso. We then show the following assertion (also see Figure 4.4.1): *Let τ_1 and τ_2 be any two inequivalent \aleph_0 -tangles of G which are not efficiently distinguished by the starting tree set T . For every separator Z efficiently separating the τ_i in G there is a modified torso H of T in which the ends η_i representing the tangles τ_i are efficiently separated by Z .* Now we apply Carmesin's theorem as a black box in all the modified torsos H of T . That is, for every modified torso H of T we obtain a tree set T_H of finite-order separations of H that efficiently distinguishes all the ends of H . Finally, we lift all of Carmesin's tree sets compatibly with each other and with T to obtain a tree set T' of finite-order separations of G that extends T . In the end, every separation in T_H which efficiently distinguishes two ends η_i in H , with the η_i as in the assertion above, gets lifted to a separation in T' that efficiently distinguishes the τ_i in G .

Phrased differently, we reflect the problem of efficiently distinguishing two inequivalent \aleph_0 -tangles down to the modified torsos of T . There, the problem reduces to efficiently distinguishing two proxy ends, a problem that has already been solved by Carmesin. Finally, we lift the solutions for the modified torsos of T up to the original graph G to solve the original problem.

4.5. From principal collections of separators to tree sets

In this section, we show how the separations $\{X, \mathcal{C}_X\}$ can be slightly modified to give rise to a tree set that comes with quite a list of useful properties. Even though our initial intention is to consider these separations for critical vertex sets X of G , we can prove a much stronger result by more generally considering what we call principal collections of vertex sets:

Definition 4.5.1. Given a collection \mathcal{Y} of 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.

Notation. If $X \subseteq V(G)$ meets precisely one component of $G - Y$ for some $Y \subseteq V(G)$, then we denote this component by $C_Y(X)$.

Definition 4.5.2. A set $X \in \mathcal{X}$ is *principal* if it is \mathcal{X} -principal.

4. A tree-of-tangles theorem for infinite-order tangles

Example 4.5.3. An $X \in \mathcal{X}$ is principal, e.g., if it induces a clique $G[X]$ or is included in a critical vertex set of G .

Since principal vertex sets behave like cliques it is possible to alter the graph G so that all principal vertex sets actually induce cliques while the finite-order separations stay the same:

Lemma 4.5.4. *Suppose that \mathcal{Y} is a collection of principal vertex sets of G and let $G_{\mathcal{Y}}$ be obtained from G by turning each $G[X]$ with $X \in \mathcal{Y}$ into a clique. Then the finite-order separations of G are precisely the finite-order separations of $G_{\mathcal{Y}}$. In particular, $\Theta(G) = \Theta(G_{\mathcal{Y}})$.*

Proof. If $\{A, B\}$ is a finite-order separation of G , then each principal $X \in \mathcal{Y}$ meets at most one component of $G - (A \cap B)$. Therefore, no X adds an $(A \setminus B) - (B \setminus A)$ edge in $G_{\mathcal{Y}}$, so $\{A, B\}$ is also a finite-order separation of $G_{\mathcal{Y}}$. The converse holds due to $E(G_{\mathcal{Y}}) \supseteq E(G)$. \square

We will use this lemma in Section 4.8 to assume without loss of generality that, for a certain tree set, the torsos coincide with the parts. Our next definition extends ‘forming a problem case’ from critical vertex sets to arbitrary vertex sets:

Definition 4.5.5. Two vertex sets X and Y of G with $\{X, Y\}$ principal are said to *form a problem case* if X and Y are incomparable as sets and additionally $C_X(Y) \in \mathcal{C}_X$ and $C_Y(X) \in \mathcal{C}_Y$ hold.




Figure 4.5.1.: Two incomparable sets X and Y such that $\{X, Y\}$ is principal. Note that every component of $G - X$ which is neither $C_X(Y)$ nor contained in $C_Y(X)$ has its neighbourhood in $X \cap Y$ and is thus also a component of $G - Y$ (black circles). Also, not every component of $G - Y$ which is contained in $C_X(Y)$ has to be contained in $\check{\mathcal{C}}_Y$, as is depicted by the blue circle on the right. If $C_Y(X) \notin \mathcal{H}(Y)$, then $\mathcal{H}(Y)$ is a subset of $\check{\mathcal{C}}_Y \setminus \{C_Y(X)\}$ and thus $(X, C_X(Y)) \leq (Y, \mathcal{H}(Y))$.

The following lemma will keep proofs short:

Lemma 4.5.6. *If $\{X, \mathcal{C}\}$ and $\{Y, \mathcal{D}\}$ are separations of G satisfying $X \cup V[\mathcal{C}] \supseteq Y \cup V[\mathcal{D}]$ and that each component in \mathcal{D} avoids X , then $(X, \mathcal{C}) \leq (Y, \mathcal{D})$.*

4. A tree-of-tangles theorem for infinite-order tangles

Proof. It remains to show $V \setminus V[\mathcal{C}] \subseteq V \setminus V[\mathcal{D}]$ which is tantamount to $V[\mathcal{D}] \subseteq V[\mathcal{C}]$, which in turn is evident from the assumptions. \square

In the previous section, we have seen that for two distinct critical (in particular principal) vertex sets $X \neq Y$ their separations $\{X, \mathcal{C}_X\}$ and $\{Y, \mathcal{C}_Y\}$ need not be nested. This may happen, for example, if X and Y form a problem case. The following two lemmas show that actually this may happen only if X and Y form a problem case.

Lemma 4.5.7. *Let $X \subsetneq Y$ be two vertex sets of G such that $\{X, Y\}$ is principal. Then all of the components in \mathcal{C}_Y are properly contained in the component $C_X(Y)$. Notably, $C_X(Y) \in \mathcal{C}_X$ if \mathcal{C}_Y is non-empty. Moreover, if we are given subsets $\mathcal{C} \subseteq \mathcal{C}_X$ and $\mathcal{D} \subseteq \mathcal{C}_Y$, then*

$$\vec{s} \leq (X, C_X(Y)) \leq (Y, \mathcal{C}_Y) \leq (Y, \mathcal{D}) \text{ where } \begin{cases} \vec{s} = (X, \mathcal{C}) & \text{if } C_X(Y) \in \mathcal{C} \\ \vec{s} = (\mathcal{C}, X) & \text{otherwise} \end{cases}$$

so in particular $\{X, \mathcal{C}\}$ and $\{Y, \mathcal{D}\}$ are nested with each other. If additionally \mathcal{D} is non-empty, then $(X, \mathcal{C}) \not\leq (\mathcal{D}, Y)$.

Proof. Since every component $C \in \mathcal{C}_Y$ has neighbourhood precisely equal to Y , it follows from $X \subsetneq Y$ that $\bigcup(\mathcal{C}_Y \upharpoonright X) \subsetneq C_X(Y)$. Hence Lemma 4.5.6 yields $(X, C_X(Y)) \leq (Y, \mathcal{C}_Y)$. From this, the rest is evident. \square

Our next lemma is also illustrated in Figure 4.5.1.

Lemma 4.5.8. *Let X and Y be two incomparable vertex sets of G such that $\{X, Y\}$ is principal. If we are given subsets $\mathcal{C} \subseteq \mathcal{C}_X$ and $\mathcal{D} \subseteq \mathcal{C}_Y$ with $C_Y(X) \notin \mathcal{D}$, then*

$$\vec{s} \leq (X, C_X(Y)) \leq (Y, \mathcal{D}) \text{ where } \begin{cases} \vec{s} = (X, \mathcal{C}) & \text{if } C_X(Y) \in \mathcal{C} \\ \vec{s} = (\mathcal{C}, X) & \text{otherwise} \end{cases}$$

so in particular $\{X, \mathcal{C}\}$ and $\{Y, \mathcal{D}\}$ are nested with each other and we have $(X, \mathcal{C}) \not\leq (\mathcal{D}, Y)$.

Proof. The assumption $C_Y(X) \notin \mathcal{D}$ ensures that every component in \mathcal{D} avoids X . Let y be any vertex in $Y \setminus X$. As every component in \mathcal{D} avoids X and sends an edge to $y \in Y \setminus X$, we deduce that $(Y \setminus X) \cup \bigcup \mathcal{D} \subseteq C_X(Y)$. Hence Lemma 4.5.6 yields $(X, C_X(Y)) \leq (Y, \mathcal{D})$. From this, the rest is evident. \square

We are now ready to prove our main technical result, Theorem 4.5.11. To allow for more flexibility in its applications, we have extracted the following definition and second main technical result from Theorem 4.5.11:

Definition 4.5.9. Suppose that \mathcal{Y} is a principal collection of vertex sets of G . A function that assigns to every $X \in \mathcal{Y}$ a subset $\mathcal{H}(X) \subseteq \mathcal{C}_X$ is called *admissible* 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 admissible* for \mathcal{Y} .

4. A tree-of-tangles theorem for infinite-order tangles

Theorem 4.5.10. *For every principal collection of vertex sets of a connected graph there is a strongly admissable function.*

Proof. Let \mathcal{Y} be a principal collection of vertex sets of a connected graph G . We write \mathcal{P} for the collection of those principal vertex sets in \mathcal{Y} that form a problem case with some other principal vertex set in \mathcal{Y} . Let us fix any well-ordering of \mathcal{P} and view \mathcal{P} as well-ordered set from now on.

For each $X \in \mathcal{P}$ we put $K(X) := C_X(Y)$ for the first $Y \in \mathcal{P}$ which forms a problem case with X . Let us put $\mathcal{K}(X) := \mathcal{C}_X \setminus \{K(X)\}$ for every $X \in \mathcal{P}$, and $\mathcal{K}(X) := \mathcal{C}_X$ for all other vertex sets $X \in \mathcal{Y}$. We claim that \mathcal{K} is strongly admissable for \mathcal{Y} .

For this, let $X \neq Y$ be any two distinct vertex sets in \mathcal{Y} that form a problem case. We show that at least one of $K(X) = C_X(Y)$ and $K(Y) = C_Y(X)$ holds. Let $Z \in \mathcal{P}$ be the first vertex set that forms a problem case with one of X and Y . Without loss of generality we may assume that Z forms a problem case with X , so we have $K(X) = C_X(Z)$ by the minimal choice of Z . Since we are done if Y and Z meet the same component of $G - X$, we may assume that $C_X(Y) \neq C_X(Z)$. This means that the three sets X, Y, Z are pairwise incomparable. Our plan is to show that Y forms a problem case with Z , and that this gives $K(Y) = C_Y(Z) = C_Y(X)$ as desired.

We already know that Y and Z are incomparable. Next, let us verify that $C_Y(Z) \in \mathcal{C}_Y$. For this, pick any vertex $x \in X \setminus Y$. As X and Z form a problem case we have $C_X(Z) \in \mathcal{C}_X$, so the vertex x sends some edge e to the component $C_X(Z)$. Now x is not in Y and the component $C_X(Z)$ avoids Y as Y and Z live in distinct components of $G - X$ by assumption, so $C_X(Z) + e$ is a connected subgraph of $G - Y$ that meets both X and Z , yielding $C_Y(Z) = C_Y(X)$. Since Y and X form a problem case, giving $C_Y(X) \in \mathcal{C}_Y$, we get $C_Y(Z) \in \mathcal{C}_Y$ as required. By symmetry we have $C_Z(Y) \in \mathcal{C}_Z$, so Y and Z form a problem case as desired and $K(Y) = C_Y(Z)$ follows from the minimal choice of Z . To see that $K(Y) = C_Y(X)$ holds, recall that we proved $C_Y(Z) = C_Y(X)$ three sentences earlier. \square

Finally, we go for Theorem 4.5.11, which considers tree sets of the following form:

Notation. Given a principal collection \mathcal{Y} of vertex sets of G and an admissable function \mathcal{K} for \mathcal{Y} we write

$$T(\mathcal{Y}, \mathcal{K}) := \{ \{X, \mathcal{K}(X)\}, \{X, K\} \mid X \in \mathcal{Y} \text{ and } K \in \mathcal{K}(X) \}.$$

For every vertex set $X \in \mathcal{Y}$ we write $\sigma_X^{\mathcal{K}}$ for the star that consists of the separation $(X, \mathcal{K}(X))$ and all the separations (K, X) with $K \in \mathcal{K}(X)$. Notably, each star $\sigma_X^{\mathcal{K}}$ has interior X .

4. A tree-of-tangles theorem for infinite-order tangles




Figure 4.5.2.: A principal set $\mathcal{Y} = \{W, X, Y, Z\}$ of pairwise disjoint sets and the separations of the form $(\mathcal{K}(X'), X')$ for $X' \in \mathcal{Y}$ where \mathcal{K} is some admissible function for \mathcal{Y} . Note that in accordance with part (i) of Theorem 4.5.11 the depicted separations form a partial consistent orientation.

Theorem 4.5.11. *Let G be any connected graph, let \mathcal{Y} be a principal collection of vertex sets of G and let \mathcal{K} be an admissible function for \mathcal{Y} . Abbreviate $T(\mathcal{Y}, \mathcal{K}) = T$ and $\sigma_X^{\mathcal{K}} = \sigma_X$. Then the following assertions hold:*

- (i) *For every distinct two $X, Y \in \mathcal{Y}$, after possibly swapping X and Y , either*

$$\begin{aligned} & (\mathcal{K}(X), X) \leq (Y, \mathcal{K}(Y)) \\ \text{or } & (X, \mathcal{K}(X)) \leq (X, C_X(Y)) \leq (Y, \mathcal{K}(Y)). \end{aligned}$$

The collection of all separations $(\mathcal{K}(X), X)$ with $\mathcal{K}(X) \neq \emptyset$ forms a consistent partial orientation of T .

- (ii) *The collection T of separations is nested.
It is a regular tree set if $\emptyset \subsetneq \mathcal{K}(X) \subsetneq \mathcal{C}_X$ holds for all $X \in \mathcal{Y}$.*
- (iii) *Every star σ_X with $X \in \mathcal{Y}$ is a splitting star of \vec{T} .*

Moreover, if all the vertex sets in \mathcal{Y} are finite, then we may add:

- (iv) *If τ is an ultrafilter tangle of G with $X_\tau \in \mathcal{Y}$ and $\mathcal{K}(X_\tau) \in U(\tau, X_\tau)$, then τ induces via $\tau \mapsto \tau \cap \vec{T}$ on T the consistent orientation which is given by the infinite splitting star σ_{X_τ} in that $\tau \cap \vec{T} = [\sigma_{X_\tau}]$.*
- (v) *If $\text{crit}(G) \subseteq \mathcal{Y}$ and $\mathcal{C}_X \setminus \mathcal{K}(X)$ is finite for all $X \in \text{crit}(G)$, then T distinguishes every two inequivalent ultrafilter tangles τ_1 and τ_2 of G via separations in $\sigma_{X_{\tau_1}}$ and $\sigma_{X_{\tau_2}}$, and it distinguishes every end from every ultrafilter tangle τ via a separation in σ_{X_τ} .*

Proof. (i) If X and Y are comparable with $X \subsetneq Y$, say, then we are done by Lemma 4.5.7. Otherwise X and Y are incomparable, and then we are done by Lemma 4.5.8 since \mathcal{K} is admissible.

4. A tree-of-tangles theorem for infinite-order tangles

(ii) That T is nested follows from (i). For the ‘moreover’ part note that requiring $\emptyset \subsetneq \mathcal{H}(X) \subsetneq \mathcal{C}_X$ ensures that $\{X, \mathcal{H}(X)\}$ has no small orientation.

(iii) It suffices to show that every separation in T with separator $Y \neq X$ has an orientation that lies below some element of σ_X . So consider any $Y \in \mathcal{Y}$ other than X . Since σ_Y is a star, it suffices to show that some separation in $(\sigma_Y)^*$ lies below some element of σ_X . By (i) it suffices to consider the following cases. If $(\mathcal{H}(Y), Y) \leq (X, \mathcal{H}(X))$ we are done. Otherwise either

$$\begin{aligned} (X, \mathcal{H}(X)) &\leq (X, C_X(Y)) \leq (Y, \mathcal{H}(Y)) \\ \text{or } (Y, \mathcal{H}(Y)) &\leq (Y, C_Y(X)) \leq (X, \mathcal{H}(X)). \end{aligned}$$

In the first case we are fine since $(\mathcal{H}(Y), Y) \leq (C_X(Y), X) \in \sigma_X$. And in the second case we are done by the second inequality.

(iv) Let τ be any ultrafilter tangle of G with $X_\tau \in \mathcal{Y}$ and write $X = X_\tau$. First, we show that σ_X is included in $O := \tau \cap \vec{T}$. The assumption $\mathcal{H}(X) \in U(\tau, X)$ means $(X, \mathcal{H}(X)) \in O$. Moreover, we have $(K, X) \in \tau$ for every $K \in \mathcal{H}(X_\tau)$ as $U(\tau, X)$ is a free ultrafilter. Thus $\sigma_X \subseteq O$, and so $[\sigma_X] \subseteq O$ by consistency. Conversely, $O \subseteq [\sigma_X]$ since $\sigma_X \subseteq O$ is a splitting star of \vec{T} by (iii).

(v) If τ_1 and τ_2 are two ultrafilter tangles of G with $X_{\tau_1} \neq X_{\tau_2}$, then the induced orientations $\tau_i \cap \vec{T}$ come from distinct splitting stars $\sigma_{X_{\tau_i}}$ of \vec{T} by (iii). Now if ω is an end of G and τ is an ultrafilter tangle, then ω avoids the star σ_{X_τ} since it has finite interior (cf. [25, Corollary 1.7]) while τ contains it by (iii). \square

We close this section by showing that in general it is not possible to find an admissible function \mathcal{H} for which $\vec{T}(\text{crit}(G), \mathcal{H})$ is a tree set that is even isomorphic to the edge tree set of a tree.




Figure 4.5.3.: A T_{\aleph_0} (black) with 2^{\aleph_0} many copies of K_{2, \aleph_0} as ‘tops’ (visualised in red for the right-most ray)

Example 4.5.12. If G is the graph shown in Figure 4.5.3, then there is no function assigning to each critical vertex set X of G a cofinite subset $\mathcal{H}(X) \subseteq \mathcal{C}_X$ such

4. A tree-of-tangles theorem for infinite-order tangles

that

$$N := \{ \{X, \mathcal{K}(X)\} \mid X \in \text{crit}(G) \}$$

gives rise to a tree set \vec{N} that is isomorphic to the edge tree set of a tree (so in particular it cannot be induced by an S_{\aleph_0} -tree or tree-decomposition of G). First, however, we describe G more precisely. The graph G is obtained from the \aleph_0 -regular tree $T = T_{\aleph_0}$ by fixing any root r and then proceeding as follows. For every ray $R \subseteq T$ starting at the root r we add a new copy of K_{2, \aleph_0} with 2-class $\{x_R, y_R\}$, say, and join x_R to every vertex of the ray R . Readers familiar with the ‘binary tree with tops’ will note that G extends a ‘ T_{\aleph_0} with tops’.

Let us check that there really is no suitable function $X \mapsto \mathcal{K}(X)$ as claimed. Assume for a contradiction that there is. Then \vec{N} is a tree set that, by Theorem 2.3.1, has no $(\omega + 1)$ -chains. Hence to yield a contradiction, it suffices to find an $(\omega + 1)$ -chain. If t is a node of $T \subseteq G$, then its down-closure $\lceil t \rceil$ in T is a critical vertex set of G , and the components in $\mathcal{C}_{\lceil t \rceil}^{\check{}}$ are of the following form. If t' is an upward neighbour of t in T , then the vertex set of the component of $G - \lceil t \rceil$ containing t' is given by the union of $\lceil t' \rceil \subseteq T$ with all the copies of K_{2, \aleph_0} whose corresponding ray has a tail in $\lceil t' \rceil$. This gives a bijection between the upward neighbours of t in T and the components in $\mathcal{C}_{\lceil t \rceil}^{\check{}}$. Next, we claim that there is a ray $R^* = t_0 t_1 t_2 \dots \subseteq T$ starting at the root r such that for all $n > 0$ the node t_n corresponds to a component in $\mathcal{K}(\lceil t_{n-1} \rceil)$ for its predecessor t_{n-1} . Indeed, since $\mathcal{K}(\lceil t \rceil) \subseteq \mathcal{C}_{\lceil t \rceil}^{\check{}}$ is infinite for all $t \in T$, such a ray can be constructed inductively. But then we get a strictly ascending sequence

$$(\lceil t_0 \rceil, \mathcal{K}(\lceil t_0 \rceil)) < (\lceil t_1 \rceil, \mathcal{K}(\lceil t_1 \rceil)) < (\lceil t_2 \rceil, \mathcal{K}(\lceil t_2 \rceil)) < \dots$$

i.e. we get an ω -chain in \vec{N} . And this ω -chain extends to an $(\omega + 1)$ -chain as the separation $(Z, \mathcal{K}(Z))$ with $Z = \{x_{R^*}, y_{R^*}\}$ that comes from the K_{2, \aleph_0} for R^* is greater than all separations $(\lceil t_n \rceil, \mathcal{K}(\lceil t_n \rceil))$. \square

4.6. Applications

This section is dedicated to the applications of our work mentioned in the introduction. All of the four applications are, in fact, applications of Theorems 4.5.10 and 4.5.11. Elbracht, Kneip and Teegen [34] use the following corollary of our two theorems:

Corollary 4.6.1. *Suppose that \mathcal{Y} is a principal collection of vertex sets of G . Then there is a function \mathcal{K} assigning to each $X \in \mathcal{Y}$ a subset $\mathcal{K}(X) \subseteq \mathcal{C}_X^{\check{}}$ that misses at most one component from \mathcal{C}_X , such that the collection*

$$\{ \{V \setminus K, X \cup K\} \mid X \in \mathcal{Y} \text{ and } K \in \mathcal{K}(X) \}$$

is nested. \square

Chapter 8 will use Theorems 4.5.10 and 4.5.11 directly. In the remainder of this section, we present applications three and four: a structural connectivity result for infinite graphs, and the collectionwise normality of ultrafilter tangle spaces.

4. A tree-of-tangles theorem for infinite-order tangles

4.6.1. A structural connectivity result for infinite graphs

We have already explained this application in detail in our introduction, now we prove it:

Theorem 4.2. *Every connected graph G has a tree set whose separators are precisely the critical vertex sets of G and all whose torsos are tough.*

Proof. By Theorems 4.5.10 and 4.5.11 it suffices to show that for $\mathcal{Y} := \text{crit}(G)$ and a strongly admissible function \mathcal{K} the torsos of the tree set $T(\mathcal{Y}, \mathcal{K})$ are tough. For this, let O be any consistent orientation of $T(\mathcal{Y}, \mathcal{K})$, let Π be its part and H its torso. In order to show that H is tough, let Ξ be a finite subset of $V(H)$. Let $\mathcal{C} \subseteq \mathcal{C}_\Xi$ consist of those components of $G - \Xi$ that meet Π . Then \mathcal{C} must be finite: otherwise Ξ contains a critical vertex set Ξ' of G with $\mathcal{C}' := \mathcal{C}_{\Xi'} \cap \mathcal{C}$ infinite; and then $(\Xi', C) \in O$ for all $C \in \mathcal{C}' \cap \mathcal{K}(\Xi')$ as these C meet Π , contradicting the consistency of O . Thus $G - \Xi$ has only finitely many components meeting Π . By Corollary 4.2.10 each of these components induces a component of $H - \Xi$, and so $H - \Xi$ has only finitely many components. \square

4.6.2. Collectionwise normality of the ultrafilter tangle space

For this subsection, we recall the following definitions from general topology (cf. [36]):

Definition 4.6.2 (Normality and collectionwise normality). Let X be a topological space in which all singletons are closed.

The space X is said to be *normal* if for every two disjoint closed subsets A_1 and A_2 of X there are disjoint open subsets O_1 and O_2 of X with $A_1 \subseteq O_1$ and $A_2 \subseteq O_2$.

A collection $\{A_i \mid i \in I\}$ of subsets $A_i \subseteq X$ is *discrete* if every point $x \in X$ has an open neighbourhood that meets at most one of the A_i .

The space X is said to be *collectionwise normal* if for every discrete collection $\{A_i \mid i \in I\}$ of pairwise disjoint closed subsets $A_i \subseteq X$ there exists a collection $\{O_i \mid i \in I\}$ of pairwise disjoint open subsets $O_i \subseteq X$ with $A_i \subseteq O_i$ for all $i \in I$.

The following implications are true for every topological space (the first implication is [36, Theorems 5.1.1 and 5.1.18] whereas the second is clear):

$$\text{compact Hausdorff} \Rightarrow \text{collectionwise normal} \Rightarrow \text{normal}.$$

The end space $\Omega(G)$ of a graph G is endowed with the topology whose basic open sets are of the form $\{\omega \in \Omega(G) \mid C(A \cap B, \omega) \subseteq G[B \setminus A]\}$, one for every oriented finite-order separation (A, B) of G . In general, the end space $\Omega(G)$ is not compact, e.g., if G is a union of infinitely many rays sharing precisely their initial vertices.

Polat [68] and Sprüssel [80] independently showed that the end space of every graph G is normal, and Polat even showed that the end space is collectionwise normal (this is Lemma 4.14 in [68], see Chapter 10 for a modern proof):

4. A tree-of-tangles theorem for infinite-order tangles

Theorem. *Every graph G has a collectionwise normal end space $\Omega(G)$.*

The \aleph_0 -tangle space $\Theta(G)$ has been endowed with a natural topology by Diestel [25] that makes it compact Hausdorff while containing the end space as a subspace. The basic open sets of Diestel's topology are of the form $\{\tau \in \Theta(G) \mid (A, B) \in \tau\}$, one for every oriented finite-order separation (A, B) of G . Since every compact Hausdorff space is collectionwise normal, the \aleph_0 -tangle space is collectionwise normal as well:

Theorem. *Every graph G has a collectionwise normal \aleph_0 -tangle space $\Theta(G)$.*

This result, however, does not imply that the end space is collectionwise normal, for usually the end space is not closed in the \aleph_0 -tangle space.

As the \aleph_0 -tangle space is the disjoint union $\Theta(G) = \Omega(G) \sqcup \Upsilon(G)$ of the end space $\Omega(G)$ and the ultrafilter tangle space $\Upsilon(G)$, the question arises whether the ultrafilter tangle space is collectionwise normal as well. Like the end space, the ultrafilter tangle space usually is not closed in the \aleph_0 -tangle space, so the ultrafilter tangle space does not obviously inherit the collectionwise normality from the \aleph_0 -tangle space.

In this subsection we show that the ultrafilter tangle space is collectionwise normal, Theorem 4.3 (i). For readers who are familiar with the compactification $|G|_\Gamma = G \sqcup \text{crit}(G) \sqcup \Omega(G)$ from [62] we remark that our proof also shows that the critical vertex set space (with the subspace topology from $|G|_\Gamma$) is collectionwise normal as well, Theorem 4.3 (ii).

Theorem 4.3. *For every connected graph G the following two assertions hold:*

- (i) *For every discrete collection $\{A_i \mid i \in I\}$ of pairwise disjoint closed subsets $A_i \subseteq \Upsilon(G)$ there exists a collection $\{O_i \mid i \in I\}$ of pairwise disjoint open subsets $O_i \subseteq |G|_\Theta$ such that $A_i \subseteq O_i$ for all $i \in I$.
In particular, the ultrafilter tangle space of G is collectionwise normal.*
- (ii) *For every discrete collection $\{A_i \mid i \in I\}$ of pairwise disjoint closed subsets $A_i \subseteq \text{crit}(G)$ there exists a collection $\{O_i \mid i \in I\}$ of pairwise disjoint open subsets $O_i \subseteq |G|_\Gamma$ such that $A_i \subseteq O_i$ for all $i \in I$.
In particular, the critical vertex set space of G is collectionwise normal.*

For the remainder of this subsection we assume familiarity with Section 2 of [62] and use notation introduced therein.

Our proof of Theorem 4.3 will employ the following short lemma:

Lemma 4.6.3. *For every two finite-order separations $(X, \mathcal{C}) \leq (Y, \mathcal{D})$ of G we have $\mathcal{O}_{|G|_\Theta}(X, \mathcal{C}) \supseteq \mathcal{O}_{|G|_\Theta}(Y, \mathcal{D})$.*

Proof. Clearly, $G \cap \mathcal{O}_{|G|_\Theta}(X, \mathcal{C}) \supseteq G \cap \mathcal{O}_{|G|_\Theta}(Y, \mathcal{D})$. And from the consistency of \aleph_0 -tangles we deduce $\Theta \cap \mathcal{O}_{|G|_\Theta}(X, \mathcal{C}) \supseteq \Theta \cap \mathcal{O}_{|G|_\Theta}(Y, \mathcal{D})$. \square

Proof of Theorem 4.3. (i) For this, let $\{A_i \mid i \in I\}$ be any discrete collection of closed subsets $A_i \subseteq \Upsilon(G)$. We are going to find a suitable collection $\{O_i \mid i \in I\}$. To get started, we view the \aleph_0 -tangle space as inverse limit

$$\Theta = \varprojlim (\beta(\mathcal{C}_X) \mid X \in \mathcal{X})$$

4. A tree-of-tangles theorem for infinite-order tangles

where each \mathcal{C}_X is endowed with the discrete topology. Since Θ is compact and all $\beta(\mathcal{C}_X)$ are Hausdorff, it follows from general topology that all of the (continuous) projections $\text{pr}_Y: \Theta = \varprojlim \beta(\mathcal{C}_X) \rightarrow \beta(\mathcal{C}_Y)$ are open. Now consider any critical vertex set X of G . The Stone-Ćech remainder $(\check{\mathcal{C}}_X)^* = \beta(\check{\mathcal{C}}_X) \setminus \check{\mathcal{C}}_X$ is a closed subspace of $\beta(\check{\mathcal{C}}_X) = \text{cl}_{\beta(\check{\mathcal{C}}_X)}(\check{\mathcal{C}}_X) \subseteq \beta(\mathcal{C}_X)$. (This follows from general topology, but it can also be seen more directly by considering the standard basis for the Stone-Ćech compactification of discrete spaces.) And for every $U \in (\check{\mathcal{C}}_X)^*$ the preimage $\text{pr}_X^{-1}(U)$ is a singleton that consists precisely of the ultrafilter tangle of which (X, U) is the blueprint. Therefore, for every $i \in I$ the set

$$A_{i,X} := \text{pr}_X(A_i) \cap (\check{\mathcal{C}}_X)^* = \text{pr}_X(\overline{A_i}^\Theta) \cap (\check{\mathcal{C}}_X)^*$$

is closed in $\beta(\check{\mathcal{C}}_X)$. Moreover, $\{A_{i,X} \mid i \in I\}$ is a discrete collection of pairwise disjoint closed subsets of $\beta(\check{\mathcal{C}}_X)$. Now the Stone-Ćech compactification $\beta(\check{\mathcal{C}}_X)$ is collectionwise normal since it is compact Hausdorff, and so we find a collection $\{O_{i,X} \mid i \in I\}$ of pairwise disjoint open subsets $O_{i,X} \subseteq \beta(\check{\mathcal{C}}_X)$ satisfying the inclusion $A_{i,X} \subseteq O_{i,X}$ for all $i \in I$.

Next, we use Theorem 4.5.10 to find a strongly admissible function \mathcal{K} for the principal collection $\text{crit}(G)$. For every index $i \in I$ and every ultrafilter tangle $\tau \in A_i$ we choose a component collection $\mathcal{C}(\tau) \in U(\tau, X_\tau)$ such that

- $\mathcal{C}(\tau) \subseteq \mathcal{K}(X_\tau)$;
- $\mathcal{C}(\tau) \subseteq O_{i,X_\tau}$;
- $O_{i,\tau} := \mathcal{O}_{|G|_\Theta}(X_\tau, \mathcal{C}(\tau))$ avoids all A_j with $j \neq i$.

We find $\mathcal{C}(\tau)$ as follows. First, we recall that $\mathcal{K}(X_\tau)$ is contained in the free ultrafilter $U(\tau, X_\tau)$. Second, we note that $O_{i,X_\tau} \cap \check{\mathcal{C}}_{X_\tau}$ is contained in $U(\tau, X_\tau)$ as well, for O_{i,X_τ} is an open neighbourhood of $U = \text{pr}_{X_\tau}(\tau) \in A_{i,X_\tau}$ in $\beta(\check{\mathcal{C}}_{X_\tau})$ and U is contained in $U(\tau, X_\tau)$ as a subset. Therefore, if we find a component collection $\mathcal{C} \subseteq \check{\mathcal{C}}_{X_\tau}$ such that $\mathcal{O}_{|G|_\Theta}(X_\tau, \mathcal{C})$ avoids all A_j with $j \neq i$, then $\mathcal{C}(\tau) := \mathcal{K}(X_\tau) \cap O_{i,X_\tau} \cap \mathcal{C}$ will satisfy all three requirements (for the third requirement we apply Lemma 4.6.3 to $(X_\tau, \mathcal{C}) \leq (X_\tau, \mathcal{C}(\tau))$). To find a suitable component collection \mathcal{C} , we proceed as follows. The union of all sets A_j with $j \in I$ and $j \neq i$ is closed in $\Upsilon(G)$ since $\{A_i \mid i \in I\}$ is a discrete collection of closed sets. Hence there exists an open neighbourhood $\mathcal{O}_{|G|_\Theta}(Y, \mathcal{D})$ of τ in $|G|_\Theta$ which avoids this union. Applying Lemma 4.3.1 to $(Y, \mathcal{D}) \in \tau$ then yields a component collection $\mathcal{C} \subseteq \check{\mathcal{C}}_{X_\tau}$ satisfying $(Y, \mathcal{D}) \leq (X_\tau, \mathcal{C}) \in \tau$. In particular, $\mathcal{O}_{|G|_\Theta}(X_\tau, \mathcal{C}) \subseteq \mathcal{O}_{|G|_\Theta}(Y, \mathcal{D})$ (Lemma 4.6.3 again) avoids all A_j with $j \neq i$.

Letting $O_i := \bigcup \{O_{i,\tau} \mid \tau \in A_i\}$ for every $i \in I$, we claim that the collection $\{O_i \mid i \in I\}$ is as desired. For this, it suffices to show that for all indices $i \neq j$ and ultrafilter tangles $\tau \in A_i$ and $\tau' \in A_j$ the open neighbourhoods $O_{i,\tau}$ and $O_{j,\tau'}$ are disjoint. By Theorem 4.5.11 (i) and by symmetry, only the following three cases can possibly occur.

In the first case we have $X_\tau = X_{\tau'}$ and write $X = X_\tau$. Then $O_{i,X}$ and $O_{j,X}$ are disjoint, ensuring that $\mathcal{C}(\tau)$ and $\mathcal{C}(\tau')$ are disjoint. (If we had not involved the open sets $O_{i,X}$ and $O_{j,X}$, then the component collections $\mathcal{C}(\tau)$ and $\mathcal{C}(\tau')$ might

4. A tree-of-tangles theorem for infinite-order tangles

possibly have a non-empty finite intersection.) In particular, $O_{i,\tau}$ and $O_{j,\tau'}$ are disjoint as well.

In the second case we have $X_\tau \neq X_{\tau'}$ and $(\mathcal{K}(X_\tau), X_\tau) \leq (X_{\tau'}, \mathcal{K}(X_{\tau'}))$, which implies that $O_{i,\tau}$ and $O_{j,\tau'}$ are disjoint.

In the third case we have $X_\tau \neq X_{\tau'}$ and

$$(X_\tau, \mathcal{K}(X_\tau)) \leq (X_\tau, C) \leq (X_{\tau'}, \mathcal{K}(X_{\tau'}))$$

where C is the component $C_{X_\tau}(X_{\tau'})$. Since $O_{i,\tau}$ avoids $A_j \ni \tau'$ we deduce that the component C is not contained in $\mathcal{C}(\tau)$. Hence $(\mathcal{C}(\tau), X_\tau) \leq (X_{\tau'}, \mathcal{C}(\tau'))$ which implies that $O_{i,\tau}$ and $O_{j,\tau'}$ are disjoint.

(ii) It is possible to deduce (ii) from (i) by a careful analysis of the properties of the sets $O_{i,\tau}$ constructed in the proof of (i) above. But it is also possible to follow the strategy of the proof of (i) and show (ii) directly, as follows. For this, let $\{A_i \mid i \in I\}$ be any discrete collection of closed subsets $A_i \subseteq \text{crit}(G)$. Using Theorem 4.5.10 we find a strongly admissible function \mathcal{K} for the principal collection $\text{crit}(G)$. For every $i \in I$ and $X \in A_i$ we let $\mathcal{C}(X)$ be a cofinite subset of $\mathcal{K}(X)$ such that $\mathcal{O}_{|G|_\Gamma}(X, \mathcal{C}(X)) =: O_{i,X}$ avoids all A_j with $j \neq i$. Then letting $O_i := \bigcup \{O_{i,X} \mid X \in A_i\}$ for all $i \in I$ yields the desired collection as we verify using Theorem 4.5.11. \square

4.7. Consistent orientation and lifting from torsos

For this section, fix a graph G , a regular tree set N of finite-order separations of G , and a consistent orientation O of N . Also define $\Pi = \bigcap_{(C,D) \in O} D$.

This section deals with the problem of translating separations of $\text{torso}(G, O)$ to separations of G , as described in Section 4.4. More precisely, given a separation (A, B) of $\text{torso}(G, O)$, we want to find an *extension* of it in G , a separation (U, W) of G towards which all elements of O point such that $U \cap W \subseteq \Pi$ and $(U \cap \Pi, W \cap \Pi) = (A, B)$. Note that every extension (U, W) of (A, B) satisfies $U \cap W = A \cap B$. In general, extensions are not unique. However, the information contained in O already puts strong restrictions on the structure of extensions.

On the one hand, if x and y are vertices of G and (C, D) is a separation in O with $\{x, y\} \subseteq C$ then every extension (U, W) of a separation of $\text{torso}(G, O)$ has to satisfy $(C, D) \leq (U, W)$ or $(C, D) \leq (W, U)$ and thus $\{x, y\} \subseteq U$ or $\{x, y\} \subseteq W$. So here we have a relation on $\bigcup_{(C,D) \in O} C$ and related vertices cannot be separated by extensions of separations of $\text{torso}(G, O)$.

On the other hand, if (C, D) and (C', D') are separations in O such that O also contains some (C'', D'') with $(C, D) \leq (C'', D'')$ and $(C', D') \leq (C'', D'')$, then (C, D) and (C', D') cannot lie on different sides of (U, W) because (C'', D'') points towards every extension (U, W) of (A, B) . So here we have a relation on O and no extension of a separation of $\text{torso}(G, O)$ can separate two related separations in O .

It turns out that the two relations describe two points of view on the same idea: In this chapter we define \sim as a relation on the set of separations of O , as that

4. A tree-of-tangles theorem for infinite-order tangles

fits better in our framework of tree sets. But it is possible just as well to work with the relation on vertices, as is done e.g. in [19], and several lemmas in this section are inspired by similar lemmas in that paper. Indeed, we will associate with every equivalence class γ of \sim a set of vertices A_γ , thereby associating an equivalence class of \sim of separations with an equivalence class of vertices, and we will work with both γ and A_γ .

Lemma 4.7.1. *Define a relation \sim on O where $(C, D) \sim (C', D')$ if and only if there is a separation in O above both (C, D) and (C', D') . Then \sim is an equivalence relation.*

Proof. By definition the relation is reflexive and symmetric. In order to show transitivity, assume that (C, D) and (C', D') are related, as witnessed by $(U, W) \in O$, and that (C', D') and (C'', D'') are related, as witnessed by $(U', W') \in O$. As O is a consistent orientation, we have $(U, W) \leq (U', W')$ or $(W', U') \leq (W, U)$ or $(U, W) \leq (W', U')$. But $(U, W) \leq (W', U')$ implies $(C', D') \leq (U, W) \leq (W', U') \leq (D', C')$ and thus that $(C', D') \leq (D', C')$ which contradicts the fact that N is regular. So either $(U, W) \leq (U', W')$ or $(U', W') \leq (U, W)$ and in both cases the bigger one of these separations shows that (C, D) and (C'', D'') are related. \square

Definition 4.7.2. An equivalence class of the relation from Lemma 4.7.1 is a *corridor* of O . For a corridor γ let A_γ be the union of all sets C where $(C, D) \in \gamma$.

Remark. Let γ be a corridor and (A, B) the supremum of all elements of γ . Then $A = A_\gamma$ and $A \cap B = A_\gamma \cap \Pi$.

Lemma 4.7.3. *If (C, D) and (C', D') are elements of O and $C \setminus D'$ is non-empty then (C, D) and (C', D') are comparable.*

Proof. Because O is consistent and nested, any two separations in O either point towards each other or are comparable. Let w be a vertex contained in $C \setminus D'$. Then w witnesses that $(C, D) \not\leq (C', D')$, hence (C, D) and (C', D') do not point towards each other. \square

Remark. Lemma 4.7.1 also holds in abstract separation systems with the same proof. In particular, corridors are well-defined for abstract separation systems.

Lemma 4.7.4. *Let γ be a corridor of O and U a finite subset of A_γ . Then there is a separation (C, D) in γ such that C contains U and $C \setminus D$ contains $U \setminus \Pi$.*

Proof. First we consider the special case that U contains only one vertex $v \notin \Pi$. As v is a vertex of A_γ there is a separation (C, D) in γ such that $v \in C$. Furthermore because v is not contained in Π there is a separation (C', D') in O such that v is contained in $C' \setminus D'$. By Lemma 4.7.3 the separations (C, D) and (C', D') are comparable and thus contained in the same corridor, so (C', D') is contained in γ .

Now consider an arbitrary finite subset U of A_γ . For every vertex v of U there is a separation (C_v, D_v) in γ such that $v \in C_v$. We just showed that if v is not contained in Π then (C_v, D_v) can be chosen such that D_v does not contain v . As γ is a corridor and U is finite, there is a separation (C, D) in γ which is bigger than or equal to all separations (C_v, D_v) . In particular $v \in C_v \subseteq C$ for all $v \in U$ and $v \in C_v \setminus D_v \subseteq C \setminus D$ for all $v \in U \setminus \Pi$. \square

4. A tree-of-tangles theorem for infinite-order tangles

Lemma 4.7.5. *The sets $A_\gamma \setminus \Pi$ partition $V(G) \setminus \Pi$.*

Proof. By definition of Π every vertex $v \in V(G) \setminus \Pi$ is contained in A_γ for some corridor γ , so the sets $A_\gamma \setminus \Pi$ cover $V(G) \setminus \Pi$. To prove their disjointedness, assume that some vertex is contained in A_γ and $A_{\gamma'}$ for two corridors γ and γ' . By Lemma 4.7.4 there are separations (C, D) in γ and (C', D') in γ' respectively such that both $C \setminus D$ and $C' \setminus D'$ contain v . Thus by Lemma 4.7.3 the separations (C, D) and (C', D') are contained in the same corridor and hence $\gamma = \gamma'$. \square

Corollary 4.7.6. *For a separation (C, D) of N and a corridor γ we have $(C, D) \in \gamma$ if and only if $C \setminus D \subseteq A_\gamma$.* \square

Lemma 4.7.7. *Let U be a connected set of vertices avoiding Π . Then there is a corridor γ with $U \subseteq A_\gamma$.*

Proof. By Lemma 4.7.5 it is sufficient to show the statement for U with exactly two elements. Let v and w be two neighbours not in Π , and let (C, D) be a separation in O such that $v \in C \setminus D$. Because w is a neighbour of v and (C, D) is a separation, w is contained in U and thus for the corridor γ containing (C, D) we have that A_γ contains both v and w . \square

Corollary 4.7.8. *Let F be a finite connected set of vertices not meeting Π . Then there is a separation $(C, D) \in O$ such that $F \subseteq C \setminus D$.*

Proof. By Lemma 4.7.7 we may apply Lemma 4.7.4. \square

Lemma 4.7.9. *Let γ be a corridor and assume that all separators of separations in N are cliques. Then $A_\gamma \cap \Pi$ is a clique, too.*

Proof. Let v and w be two distinct vertices of $A_\gamma \cap \Pi$. Then by Lemma 4.7.4 there is a separation $(C, D) \in \gamma$ such that C contains both v and w . Because v and w are contained in Π which in turn is a subset of D , both v and w are contained in $C \cap D$. Because $C \cap D$ is a clique by assumption, v is a neighbour of w . \square

4.8. Extending the tree set of the principal vertex sets

In this section we prove our main result, Theorem 4.1. To obtain a starting tree set T as described in our overall proof strategy in Section 4.4, we apply our technical main result Theorem 4.5.11 (combined with Theorem 4.5.10) to a carefully chosen collection \mathcal{Y} of principal vertex sets of G . For choosing \mathcal{Y} we need the following definition:

Definition 4.8.1. A separation $\{X, \mathcal{C}\}$ is *generous* if both \mathcal{C} and the complement $\mathcal{C}_X \setminus \mathcal{C}$ contain components whose neighbourhoods are precisely equal to X , i.e. if \mathcal{C}_X meets both \mathcal{C} and $\mathcal{C}_X \setminus \mathcal{C}$. A set X of vertices of G is *generous* if it is the separator of some generous separation, i.e. if $|\mathcal{C}_X| \geq 2$.

4. A tree-of-tangles theorem for infinite-order tangles

Now we are ready to set up our starting tree set T and more, as follows. **Throughout this section we fix the following notation.** We let \mathcal{Y} be the collection of all generous subsets of the critical vertex sets of G , in formula:

$$\mathcal{Y} = \{ X \in \mathcal{X} \mid X \text{ is generous and } \exists Y \in \text{crit}(G) : X \subseteq Y \}.$$

Notably, $\text{crit}(G) \subseteq \mathcal{Y}$. We assume, without loss of generality by Lemma 4.5.4, that each $X \in \mathcal{Y}$ induces a clique $G[X]$. Using Theorems 4.5.10 and 4.5.11 we obtain a strongly admissible function \mathcal{K} for \mathcal{Y} that deviates from all \mathcal{C}_X with $X \in \mathcal{Y}$ by precisely one component, in formula $|\mathcal{C}_X \setminus \mathcal{K}(X)| = 1$ for all $X \in \mathcal{Y}$. This way we ensure that $T := T(\mathcal{Y}, \mathcal{K})$ is a regular tree set of generous finite-order separations of G . For $X \in \mathcal{Y}$ we abbreviate $\sigma_X = \sigma_X^{\mathcal{K}}$. Moreover, O always denotes a consistent orientation of T , and then $\Pi \subseteq V(G)$ denotes the part of O . At some point in this section the concept of a ‘modified torso’ of O will be defined. From that point onward, H will always denote the modified torso of O . Whenever we speak of Π or H we tacitly assume that they stem from some O . This completes the list of fixed notation for this section.

Next, we consider two inequivalent \aleph_0 -tangles τ_1 and τ_2 of G , we pick a finite-order separation $\{A_1, A_2\}$ of G that efficiently distinguishes τ_1 and τ_2 , and we write $Z = A_1 \cap A_2$ for its separator. If Z is included entirely in a critical vertex set of G , then T efficiently distinguishes τ_1 and τ_2 :

Lemma 4.8.2. *Let $\{Z, \mathcal{D}\}$ efficiently distinguish two \aleph_0 -tangles τ_1 and τ_2 of G . Then $\{Z, \mathcal{D}\}$ is generous. If additionally τ_1 and τ_2 are inequivalent and Z is included in some critical vertex set of G , then T efficiently distinguishes τ_1 and τ_2 .*

Proof. Let $\{\mathcal{D}_1, \mathcal{D}_2\} := \{\mathcal{D}, \mathcal{C}_Y \setminus \mathcal{D}\}$ such that $(Z, \mathcal{D}_i) \in \tau_i$ for both $i = 1, 2$. Our proof starts with a more general analysis of the situation, as follows. Consider any $i \in \{1, 2\}$ and put $j = 3 - i$.

If τ_i lives in a component C of $G - Z$ in that $(Z, C) \in \tau_i$, then by the consistency of τ_j we deduce from $(C, Z) \leq (Z, \mathcal{D}_j) \in \tau_j$ that $(C, Z) \in \tau_j$, so $\{Z, C\}$ distinguishes τ_1 and τ_2 . But then so does $\{N(C), C\}$ by Lemma 4.2.4, and hence $N(C) = Z$ follows by the efficiency of Z .

Otherwise τ_i is an ultrafilter tangle and $X_i := X_{\tau_i}$ is contained in Z . Then, as $U(\tau_i, Z)$ is a free ultrafilter, we have $(Z, \mathcal{D}'_i) \in \tau_i$ for $\mathcal{D}'_i := \mathcal{D}_i \cap \mathcal{C}_Z(X_i)$. Hence $(X_i, \mathcal{D}'_i) \in \tau_i$ by Lemma 4.2.4. And $(\mathcal{D}'_i, X_i) \leq (Z, \mathcal{D}_j) \in \tau_j$ implies $(\mathcal{D}'_i, X_i) \in \tau_j$ by the consistency of τ_j . Therefore, $\{X_i, \mathcal{D}'_i\}$ distinguishes τ_1 and τ_2 , so $X_i = Z$ follows by the efficiency of Z .

From the two cases above we deduce that $\{Z, \mathcal{D}\}$ is generous. It remains to show that if additionally τ_1 and τ_2 are inequivalent and Z is contained in a critical vertex set of G , then T efficiently distinguishes τ_1 and τ_2 . First, we have $Z \in \mathcal{Y}$ as Z is generous. Next, we note that not both τ_1 and τ_2 can be ultrafilter tangles with $X_1, X_2 \subseteq Z$ for otherwise $X_1 = Z = X_2$ follows from our considerations above, contradicting that τ_1 and τ_2 are inequivalent. So at least one of τ_1 and τ_2 lives in a component C of $G - Z$, say $(Z, C) \in \tau_1$, and then $C \in \mathcal{C}_Z$ follows from our considerations above. If $C \in \mathcal{K}(Z)$ then $\{Z, C\} \in T$ efficiently distinguishes τ_1 and τ_2 . Otherwise $\{C\} = \mathcal{C}_Z \setminus \mathcal{K}(Z)$, and we claim that $\{Z, \mathcal{K}(Z)\} \in T$

4. A tree-of-tangles theorem for infinite-order tangles

efficiently distinguishes τ_1 and τ_2 . On the one hand, $(\mathcal{K}(Z), Z) \leq (Z, C) \in \tau_1$ implies $(\mathcal{K}(Z), Z) \in \tau_1$ by the consistency of τ_1 . On the other hand, τ_2 either lives in a component in $\mathcal{K}(Z)$ or τ_2 is an ultrafilter tangle with $X_2 = Z$, and in both cases we deduce $(Z, \mathcal{K}(Z)) \in \tau_2$. \square

Therefore, we may assume that Z is not contained entirely in any critical vertex set of G . Then Z is contained in a part of T , as follows.

Lemma 4.8.3. *Let $Z \in \mathcal{X}$ be generous. If X is a principal vertex set of G that does not contain Z entirely, then there is a unique component of $G - X$ that Z meets.*

Proof. As Z is not contained in X as a subset, there is a component C of $G - X$ which Z meets. Assume for a contradiction that there is another component D of $G - X$ meeting Z . Pick vertices $c \in Z \cap C$ and $d \in Z \cap D$. Now note that every component $K \in \mathcal{C}_Z$ must meet X , for K plus its $K-c$ and $K-d$ edges admits a $c-d$ path connecting the distinct components C and D of $G - X$. But since X is principal it meets at most one component of $G - Z$, namely $C_Z(X)$, contradicting that $|\mathcal{C}_Z| \geq 2$. \square

By Lemma 4.8.3 above the separator Z meets precisely one side from every separation in T , and then orienting each separation in T towards that side results in a consistent orientation O of T whose part Π contains Z . The remainder of this section is dedicated to modifying the torso H of O so that

- τ_1 and τ_2 are ‘represented’ by ‘proxy’ ends η_1 and η_2 in H ; and
- applying Carmesin’s theorem in H yields a tree set that lifts compatibly with T to a tree set of tame finite-order separations of G that efficiently distinguishes τ_1 and τ_2 .

4.8.1. Modified torsos, proxies of corridors and lifting from modified torsos

In this subsection we introduce modified torsos and show that there is an elegant way to lift tree sets from modified torsos to the graph G itself. Proxies of corridors are introduced as a technical tool whose purpose is twofold: first, they are key to the elegant lifting of tree sets. And second, they will be employed in the next subsection to define proxies for \aleph_0 -tangles.

Definition 4.8.4 (Modified torso). Whenever Π is non-empty we define the *modified torso* H of O , as follows. Consider the set \mathcal{Z} of all finite subsets of Π that are separators of suprema of corridors of O . Then we obtain H from $G[\Pi]$ by disjointly adding for each $X \in \mathcal{Z}$ a copy of K^{\aleph_0} that we join completely to X .

We remark that Π being non-empty ensures that the empty set is not an element of \mathcal{Z} , so modified torsos are connected. Since the copies of K^{\aleph_0} are joined to finite cliques of $G[\Pi]$ by Lemma 4.7.9, no two ends of $G[\Pi]$ are merged when we move on to the modified torso H :

4. A tree-of-tangles theorem for infinite-order tangles

Lemma 4.8.5. *Every finite-order separation of $G[\Pi]$ extends to some finite-order separation of H . Thus sending each end η of $G[\Pi]$ to the end $\iota(\eta)$ of $H \supseteq G[\Pi]$ with $\eta \subseteq \iota(\eta)$ defines an injection $\iota: \Omega(G[\Pi]) \hookrightarrow \Omega(H)$. Moreover, the ends of H that do not lie in the image of ι correspond bijectively to the copies of K^{\aleph_0} that were added to $G[\Pi]$ in order to obtain H . \square*

Now we tend to the lifting of separations from H to G . It is desirable to have the separator of a separation remain unchanged when lifting it. But H usually will contain many vertices that are not vertices of the original graph G . We solve this as follows. When we consider finite-order separations of H , we are only interested in ones that efficiently distinguish some two ends of H . And these H -relevant separations have their separators consist of vertices of the original graph G :

Definition 4.8.6 (H -relevant). If a separation of H has finite order and efficiently distinguishes some two ends of H , then we call it and its orientations H -relevant.

Lemma 4.8.7. *If $\{A, B\}$ is H -relevant, then $A \cap B \subseteq \Pi$.*

Proof. Assume for a contradiction that $A \cap B$ meets an added copy K of a K^{\aleph_0} in a vertex v and write $X = N_H(K)$. Notably $H[X]$ is a clique, and hence so is $H[X \cup K]$. Without loss of generality $H[X \cup K] \subseteq H[A]$, so K meets $A \setminus B$ while $H[X \cup K]$ avoids $B \setminus A$. Now $v \in A \cap B \cap K$ sends its edges only to K and X , and in particular v sends no edges to $B \setminus A$. So $\{A, B - v\}$ is again a separation of H , but of order $|A \cap B| - 1$, and this separation still distinguishes all the ends of H that were distinguished by $\{A, B\}$, contradicting that $\{A, B\}$ is H -relevant. \square

Now we are almost ready to define lifts of separations of H , all we miss is

Definition 4.8.8 (Proxy of a corridor). Suppose that Π is non-empty and γ is a corridor of O . The *proxy* of γ in the modified torso H is the end η of H that is defined as follows. Consider the separator X of the supremum of γ . If X is finite, then η is the end of H containing the rays of the K^{\aleph_0} that was added for X . Otherwise $G[X] \subseteq H$ is an infinite clique by Lemma 4.7.9, and then η is the end of H that contains the rays of $G[X]$.

Finally, we can lift separations from H to G :

Definition 4.8.9 (Lift from a modified torso). Let (A, B) be an H -relevant separation of a modified torso H . By Lemma 4.8.7 the separator $A \cap B$ is included in Π entirely. The *lift* $(\ell(A), \ell(B))$ of (A, B) is defined as follows. The set $\ell(A) \subseteq V(G)$ agrees with A on Π , and a vertex of $G - \Pi$ is contained in $\ell(A)$ whenever its corridor's proxy in H lives on the A -side. The set $\ell(B)$ is defined analogously.

We remark that $\{\ell(A), \ell(B)\}$ does not depend on the orientation of $\{A, B\}$. In order to verify that the lifts work as intended we need the following lemma:

Lemma 4.8.10. *If $\{A, B\}$ is H -relevant and γ is a corridor of O whose proxy lives on the A -side, then $A_\gamma \subseteq \ell(A)$.*

4. A tree-of-tangles theorem for infinite-order tangles

Proof. We have $A_\gamma \setminus \Pi \subseteq \ell(A)$ by the definition of $\ell(A)$. It remains to show $\Xi \subseteq A$ for the separator $\Xi = A_\gamma \cap \Pi$ of the supremum of γ . If Ξ is infinite, then the proxy η of γ living on the A -side means $G[\Xi] \subseteq H[A]$. Otherwise Ξ is finite. Then η stems from a copy $K \subseteq H$ of K^{\aleph_0} that is joined completely to the clique $G[\Xi]$, and so η living on the A -side means $K \subseteq H[A]$. Consequently, the infinite clique $H[K \cup \Xi]$ is contained in $H[A]$ as well, giving $\Xi \subseteq A$ as desired. \square

Now we can check for ourselves that lifts work:

Lemma 4.8.11. *The lift of an H -relevant separation is a separation of G with the same separator.*

Proof. Let $\{A, B\}$ be any H -relevant separation, and recall that $A \cap B \subseteq \Pi$ by Lemma 4.8.7. Every vertex of $G - \Pi$ lies in A_γ for a unique corridor γ of O , and hence is contained in precisely one of $\ell(A)$ and $\ell(B)$. Thus $\ell(A) \cap \ell(B) = A \cap B$. It remains to verify that G has no edge between $\ell(A) \setminus \ell(B)$ and $\ell(B) \setminus \ell(A)$. For this, let $e = xy$ be any edge of G . If both x and y are contained in Π , then $e \subseteq A$ say, and hence $e \subseteq \ell(A)$. Otherwise one of x and y lies outside of Π , say $x \in \ell(A) \setminus \Pi$. Let γ be the corridor of O with $x \in A_\gamma \setminus \Pi$, so the proxy η of γ lives on the A -side. From $x \in A_\gamma \setminus \Pi$ we infer $y \in A_\gamma$. Then $e \subseteq A_\gamma \subseteq \ell(A)$ by Lemma 4.8.10. \square

Starting with an intuitive lemma we verify that our lifts are compatible with T and lifts of other modified torsos:

Lemma 4.8.12. *Let γ be a corridor of O and let η be the proxy of γ in H . If $\{A, B\}$ is H -relevant with η living on the A -side, then $\vec{s} \leq (\ell(A), \ell(B))$ for all $\vec{s} \in \gamma$. In particular, the lift of an H -relevant separation is nested with T .*

Proof. Consider any $(C, D) \in \gamma$. We have to show $(C, D) \leq (\ell(A), \ell(B))$. For the inclusion $C \subseteq \ell(A)$ we start with $C \subseteq A_\gamma$ and employ Lemma 4.8.10 for $A_\gamma \subseteq \ell(A)$. Now the inclusion $\ell(B) \subseteq D$ is tantamount to $C \setminus D \subseteq \ell(A) \setminus \ell(B)$ which is immediate from $C \subseteq \ell(A)$ as $C \setminus D$ avoids $\Pi \supseteq \ell(A) \cap \ell(B)$ (cf. Lemma 4.8.11). \square

Corollary 4.8.13. *If H' is the modified torso of a consistent orientation O' of T other than O , then all lifts of H -relevant separations are nested with all lifts of H' -relevant separations.* \square

Lemma 4.8.14. *If (A, B) and (C, D) are H -relevant with $(A, B) \leq (C, D)$, then their lifts satisfy $(\ell(A), \ell(B)) \leq (\ell(C), \ell(D))$. In particular, the lifts of two nested H -relevant separations are again nested.* \square

We close this subsection with the lemma that ensures that when we construct the tree set for our main result, we are able to ensure the ‘moreover’ part stating equivalent \aleph_0 -tangles orient the tree set the same way.

Lemma 4.8.15. *Every H -relevant separation lifts to a tame separation of G .*

4. A tree-of-tangles theorem for infinite-order tangles

Proof. If H stems from a consistent orientation O of T that contains some star σ_X with $X \in \mathcal{Y}$, then $O = \lceil \sigma_X \rceil$ (with the down-closure taken in \vec{T}) by Theorem 4.5.11. Consequently, H was obtained from the finite clique $G[X]$ by disjointly adding precisely one copy of a K^{\aleph_0} and joining it completely to X , so the one-ended H has no H -relevant separations. Therefore, we may assume that O avoids all of the stars σ_X with $X \in \mathcal{Y}$.

Let $\{A, B\}$ be an H -relevant separation and recall that we have $A \cap B \subseteq \Pi$ by Lemma 4.8.7. And let $X \subseteq A \cap B$ be a critical vertex set of G .

If there is a component $K \in \mathcal{K}(X)$ with $(X, K) \in O$, then the proxy of the corridor of O that contains (X, K) ensures that all the components in $\mathcal{K}(X) \setminus \{K\}$ are contained in the same side of $\{\ell(A), \ell(B)\}$.

Otherwise, since O avoids the star σ_X , we have $(\mathcal{K}(X), X) \in O$. Then the proxy of the corridor of O that contains $(\mathcal{K}(X), X)$ ensures that all the components in $\mathcal{K}(X)$ are contained in the same side of $\{\ell(A), \ell(B)\}$.

In either case, all but finitely many of the components in $\check{\mathcal{C}}_X$ lie on the same side of $\{\ell(A), \ell(B)\}$. Since $A \cap B \supseteq X$ meets at most finitely many components in $\check{\mathcal{C}}_X$, the collection $\mathcal{C}_{A \cap B}(X)$ forms a cofinite subset of $\check{\mathcal{C}}_X$, and therefore all but finitely many components in $\mathcal{C}_{A \cap B}(X)$ lie on the same side of $\{\ell(A), \ell(B)\}$ as desired. \square

4.8.2. Proxies of \aleph_0 -tangles

We start this subsection by introducing the technical notion of ‘walking a corridor’ and prove two technical lemmas about ends. This framework, together with proxies of corridors, then enables us to give a comprehensible definition of proxies of \aleph_0 -tangles. We emphasise that this technical layering is highly important to save the key segments of our overall proof from being swamped with terrible amounts of case distinctions.

Definition 4.8.16 (Walking). We say that an end ω of G *walks* a corridor γ of O if for the supremum (A, B) of γ the end ω has a ray contained in $G[A \setminus B]$. And we say that an ultrafilter tangle τ of G *walks* a corridor γ of O if τ contains the inverse of some separation in γ .

Lemma 4.8.17. *Suppose that N is a tree set of generous finite-order separations of G all whose separators induce cliques. Let ω be an end of G , let Π be the part of the orientation $O = \omega \cap \vec{N}$ that ω induces on N , and suppose that ω walks a corridor γ of O . If the separator $A_\gamma \cap \Pi$ of the supremum of γ is infinite, then $G[\Pi]$ contains a ray from ω .*

Proof. By Lemma 4.7.9 it is sufficient to show that there are infinitely many pairwise disjoint paths from R to $A_\gamma \cap \Pi$. We will recursively construct such paths P_n ($n \in \mathbb{N}$) of which only the last vertex v_n is contained in Π . Assume that P_0, \dots, P_{n-1} have already been defined. Then there is a finite non-empty initial segment R' of R such that $R' \cup P_0 \dot{v}_0 \cup \dots \cup P_{n-1} \dot{v}_{n-1}$ is connected. Let $(A, B) \in O$ be a separation such that all vertices of $R' \cup P_0 \dot{v}_0 \cup \dots \cup P_{n-1} \dot{v}_{n-1}$ are contained in $A \setminus B$ (such a separation exists by Corollary 4.7.8). Then (A, B) is contained

4. A tree-of-tangles theorem for infinite-order tangles

in γ . Every vertex v_k with $k < n$ is a neighbour of a vertex in $A \setminus B$ and thus contained in A .

As $A_\gamma \cap \Pi$ is infinite, it contains a vertex v which is not contained in $A \cap B$ and thus not contained in A . In particular, v is not contained in any path P_k with $k < n$. Because $v \in A_\gamma$, there is a separation (A', B') in γ such that $v \in A'$ and thus $v \in A' \cap B'$. Let (C, D) be a separation in γ which is bigger than both (A, B) and (A', B') . Then all vertices contained in some P_k with $k < n$ are contained in $(C \setminus D) \cup \Pi$. Furthermore $v \in C \cap D$ and R contains a vertex of $C \setminus D$. Because $(C, D) \in O$, some tail of R is contained in $D \setminus C$ and thus R also contains vertices of $D \setminus C$. As (C, D) is a separation and R connected this implies that some vertex w of R is contained in $C \cap D$. Because w is not a vertex of Π it is also not a vertex of some P_k with $k < n$.

As (D, C) is generous, there is a component of $G - (C \cap D)$ which is contained in $D \setminus C$ and whose neighbourhood is precisely equal to $C \cap D$. Thus there is a path P from w to v whose inner vertices are contained in $D \setminus C$. We already established that v and w are not vertices of any P_k with $k < n$. Hence P is disjoint from all P_k with $k < n$. Let v_n be the first vertex of P in Π and let $P_n := wPv_n$. By Corollary 4.7.8 there is a separation $(I, J) \in O$ such that the vertices of $P_n \setminus v_n$ are contained in $I \setminus J$. Then $(I, J) \in \gamma$ and $v_n \in I$, so $v_n \in A_\gamma$. As also $v_n \in \Pi$ we have $v_n \in A_\gamma \cap \Pi$ as required. \square

Lemma 4.8.18. *If an end ω of G does not lie in the closure of Π , then ω walks a unique corridor of O .*

Proof. Since ω does not lie in the closure of Π , we in particular find a ray $R \in \omega$ that avoids Π . As R is connected, it defines a corridor γ of O with $R \subseteq A_\gamma \setminus \Pi$. Then ω walks the corridor γ , and so it remains to show that γ is unique.

If $O \not\subseteq \omega$, then ω contains the inverse \overleftarrow{s} of some separation $\vec{s} \in O$, and then γ is determined as the corridor of O containing \vec{s} . Otherwise $O \subseteq \omega$. Then we assume for a contradiction that there is another ray $R' \in \omega$ that walks a corridor γ' of O other than γ . Since the suprema of γ and γ' both separate R and R' , their separators cannot be finite, and so they are infinite. But then applying Lemma 4.8.17 to either γ or γ' yields a ray of ω in $G[\Pi]$, contradicting the assumption that ω does not lie in the closure of Π . \square

Finally, we are ready for the definition of proxies of \aleph_0 -tangles. We split the definition and consider ends and ultrafilter tangles separately.

Definition 4.8.19 (Proxy of an end). If ω is an end of G , then the *proxy* of ω in H is the end η of H that is defined as follows.

- If ω lies in the closure of Π , then ω has a ray in $G[\Pi]$ by Corollary 4.2.8, and η is the end of such a ray in H (this is well-defined by Corollary 4.2.9).
- Otherwise ω does not lie in the closure of Π and by Lemma 4.8.18 walks a unique corridor γ of O ; then η is the proxy of γ in H .

Definition 4.8.20 (Proxy of an ultrafilter tangle). If τ is an ultrafilter tangle of G and O avoids the star σ_{X_τ} , then τ walks a unique corridor γ of O and the *proxy* of τ in H is the end η of H that is the proxy of γ in H .

4. A tree-of-tangles theorem for infinite-order tangles




Figure 4.8.1.: A graph G , a vertex set Z efficiently separating two inequivalent \aleph_0 -tangles τ_1 and τ_2 of G , and a modified torso $H(Z)$ which contains Z . The \aleph_0 -tangles τ_i walk corridors γ_i and $H(Z)$ has proxies η_1 and η_2 for τ_1 and τ_2 .

We close this subsection with a lemma on the interaction of lifts with proxies:

Lemma 4.8.21. *Let τ be an \aleph_0 -tangle of G and suppose that the proxy η of τ in H is defined. If $\{A, B\}$ is H -relevant and $(A, B) \in \eta$, then $(\ell(A), \ell(B)) \in \tau$.*

Proof. If τ is an end of G that lies in the closure of Π , then this follows from the fact that some ray of $G[\Pi]$ is contained in both τ and η . Otherwise τ is an \aleph_0 -tangle of G that walks a unique corridor γ of O . If additionally τ is an end, then every ray in τ that avoids Π is contained in $\ell(B)$, ensuring $(\ell(A), \ell(B)) \in \tau$. So we may assume that τ is an ultrafilter tangle. As the proxy η of τ is defined, we know that O avoids the star σ_{X_τ} so that τ walks a unique corridor γ of O . By definition, this means that τ contains the inverse of some oriented separation from γ . Then $(\ell(A), \ell(B)) \in \tau$ follows from Lemma 4.8.12 and the consistency of τ . \square

4.8.3. Efficiently distinguishing the proxies

In this subsection we provide the final key segments of our overall proof. We start with an overview of the situation that is of interest.

Throughout this subsection we fix the following notation in addition to the notation fixed throughout the ambient section. (See also Figure 4.8.1.) We are given two inequivalent \aleph_0 -tangles τ_1 and τ_2 of G that are efficiently distinguished by a finite-order separation $\{A_1, A_2\}$ of G with separator $Z = A_1 \cap A_2$. The separator Z is not contained in a critical vertex set of G . Hence, by Lemma 4.8.3 the separator Z meets precisely one side from every separation in T , and then orienting each separation in T towards that side results in a consistent orientation O of T whose part Π contains Z . For this special orientation we write $O(Z)$, and we write $\Pi(Z)$ and $H(Z)$ for its part and modified torso. Moreover, η_1 and η_2 are the proxies of τ_1 and τ_2 in $H(Z)$ (note that these are defined as $O(Z)$ avoids all stars σ_X with $X \in \mathcal{Y}$). Whenever we write i we mean an arbitrary $i \in \{1, 2\}$, and we write $j = 3 - i$. If τ_i happens to be an ultrafilter tangle, then we write X_i instead of X_{τ_i} . This completes the list of fixed notation for this subsection.

The final key segments are Lemma 4.8.22 and Proposition 4.8.24 below. We start with the lemma:

4. A tree-of-tangles theorem for infinite-order tangles

Lemma 4.8.22. *Every relevant finite-order separation of H that distinguishes η_1 and η_2 does lift to a separation of G that distinguishes τ_1 and τ_2 .*

Proof. Let $\{A, B\}$ be a relevant finite-order separation of H that distinguishes η_1 and η_2 , say with $(B, A) \in \eta_1$ and $(A, B) \in \eta_2$. Then Lemma 4.8.21 gives both $(\ell(B), \ell(A)) \in \tau_1$ and $(\ell(A), \ell(B)) \in \tau_2$, so $\{\ell(A), \ell(B)\}$ distinguishes τ_1 and τ_2 . \square

For the key proposition, we need the following proposition whose proof we postpone to after the proof of the key proposition.

Proposition 4.8.23. *If τ_i walks a corridor γ_i of $O(Z)$ where Ξ_i denotes the separator of the supremum of γ_i , then $G[\Xi_i \setminus Z]$ is a non-empty clique that is entirely contained in $G[A_i \setminus A_j]$.*

The key final key segment is

Proposition 4.8.24. *The proxies η_1 and η_2 are efficiently distinguished by Z .*

Proof. If Z distinguishes the proxies η_1 and η_2 in $H(Z)$, then it does so efficiently, for otherwise the separation of order $< |Z|$ doing so lifts to one distinguishing τ_1 and τ_2 in G by Lemma 4.8.22, contradicting the efficiency of Z . Therefore, it remains to show that η_1 and η_2 are distinguished by Z . For this, we check three cases.

In the first case, both τ_1 and τ_2 lie in the closure of $\Pi(Z)$. Then τ_1 and τ_2 are distinct ends of G that lie in the closure of $\Pi(Z)$, and so their proxies stem from rays of τ_1 and τ_2 respectively. Now Z witnesses that these rays are inequivalent in G and, in particular, that they are inequivalent in $G[\Pi]$. Thus Z distinguishes η_1 and η_2 in $H(Z)$ by Lemma 4.8.5.

In the second case, neither τ_1 nor τ_2 lies in the closure of $\Pi(Z)$, and both walk corridors γ_1 and γ_2 of $O(Z)$. We let Ξ_1 and Ξ_2 be the separators of the suprema of γ_1 and γ_2 . Then, by Proposition 4.8.23, for both $i = 1, 2$ the induced subgraph $G[\Xi_i \setminus Z]$ is a non-empty clique that is entirely contained in $G[A_i \setminus A_j]$. Consequently, Z distinguishes η_1 and η_2 in $H(Z)$ by Lemma 4.8.5.

In the third case, τ_1 does not lie in the closure of $\Pi(Z)$ and walks a corridor γ_1 of $O(Z)$ while τ_2 lies in the closure of $\Pi(Z)$. Then τ_2 must be an end of G . We let Ξ_1 be the separator of the supremum of γ_1 .

By Proposition 4.8.23 the induced subgraph $G[\Xi_1 \setminus Z]$ is a non-empty clique that is entirely contained in $G[A_1 \setminus A_2]$. Since η_1 stems from the copy of K^{\aleph_0} that is attached to the clique $G[\Xi_1] \subseteq G[\Pi]$ while η_2 stems from a ray of $G[\Pi]$ in τ_2 , we deduce that Z distinguishes η_1 and η_2 in $H(Z)$ by Lemma 4.8.5. \square

In the remainder of this subsection we prove Proposition 4.8.23. For this, we introduce the concept of a pointer. Basically, the idea is to have a connected subgraph of G that can be employed as an oracle—like we employ rays as oracles for their ends.

Definition 4.8.25. (Pointer) If τ_i walks a corridor γ_i of $O(Z)$, then a *pointer* of τ_i is a connected subgraph K_i of $G[A_{\gamma_i} \setminus \Pi(Z)] \cap G[A_i \setminus A_j]$ that is of the following form. If τ_i is an end of G , then K_i is a ray in τ_i . Otherwise τ_i is an ultrafilter tangle of G , and then K_i is a component in \mathcal{C}_{X_i} .

4. A tree-of-tangles theorem for infinite-order tangles

Lemma 4.8.26. *If τ_i walks a corridor of $O(Z)$, then τ_i has a pointer.*

Proof. If τ_i is an end, then τ_i has a ray avoiding $\Pi(Z) \cup Z$ for τ_i walks a corridor of $O(Z)$ and Z is finite, and every such ray is a pointer of τ_i . Otherwise τ_i is an ultrafilter tangle. Then we let γ_i be the corridor of $O(Z)$ walked by τ_i . Let $(C, D) \in \gamma_i$ witness that τ_i walks γ_i , so $D \setminus C \subseteq A_{\gamma_i} \setminus \Pi(Z)$. Using Theorem 4.2.2 we pick $\mathcal{C} \in U(\tau_i, X_i)$ with $V[\mathcal{C}] \subseteq A_i \setminus A_j$ and $\mathcal{C}' \in U(\tau_i, X_i)$ with $V[\mathcal{C}'] \subseteq D \setminus C$. As $U(\tau_i, X_i)$ is a free ultrafilter, the intersection $\mathcal{C} \cap \mathcal{C}' \cap \mathcal{C}_{X_i} \in U(\tau_i, X_i)$ is infinite, and every component in this intersection is a pointer of τ_i . \square

Lemma 4.8.27. *If the neighbourhood $N(C_i)$ of the component C_i of $G - \Pi(Z)$ containing a pointer K_i of τ_i is finite, then $N(C_i)$ is the separator of some finite-order separation of G that distinguishes τ_1 and τ_2 .*

Proof. By the consistency of τ_j it suffices to find a separation $(A, B) \in \tau_i$ with $(B, A) \leq (A_i, A_j)$ and $A \cap B = Y$ where we write $Y = N(C_i)$. As K_i is a pointer we have $K_i \subseteq G[A_i \setminus A_j]$. Since the separator $Z = A_1 \cap A_2$ is included in $\Pi(Z)$ we have $C_i \subseteq G[A_i \setminus A_j]$ as well. If τ_i is an end then $(Y, C_i) \in \tau_i$ is as desired. Otherwise τ_i is an ultrafilter tangle. If $C_i = K_i$ then $N(C_i) = N(K_i) = X_i$, and employing Theorem 4.2.2 we may pick $\mathcal{C} \in U(\tau_i, X_i)$ with $V[\mathcal{C}] \subseteq A_i$, so $(X_i, \mathcal{C}) \in \tau_i$ is as desired. Hence we may assume that $C_i \supsetneq K_i$ must meet X_i . Then $C_i = C_Y(X_i)$, and by Lemma 4.2.3 we have $(Y, C_Y(X_i)) \in \tau_i$ as desired. \square

Proof of Proposition 4.8.23. By Lemma 4.8.26 we find a pointer K_i of τ_i , and we let C_i be the component of $G - \Pi(Z)$ containing K_i . Then $K_i \subseteq G[A_{\gamma_i} \setminus \Pi(Z)]$ implies $C_i \subseteq G[A_{\gamma_i} \setminus \Pi(Z)]$, so we have $N(C_i) \subseteq \Xi_i$. First, we show that C_i has a neighbour in $\Xi_i \setminus Z$. Otherwise $N(C_i) \subseteq \Xi_i \cap Z$, and then $N(C_i) = Z$ by Lemma 4.8.27 and the efficiency of Z . Now $Z \subseteq \Xi_i$ with Z being finite allows us to find a separation $(X, \mathcal{C}) \in \gamma_i$ with $Z \subseteq X$ contradicting this subsection's assumption on Z . Therefore, C_i has a neighbour in $\Xi_i \setminus Z$. Next, since $G[\Xi_i]$ is a clique, there is a unique component D_i of $G - Z$ containing $G[\Xi_i \setminus Z]$. Then $C_i \subseteq D_i$ as C_i has a neighbour in $\Xi_i \setminus Z$, and so $G[\Xi_i \setminus Z] \subseteq G[A_i \setminus A_j]$ follows from the pointer K_i being included in $G[A_i \setminus A_j]$. \square

4.8.4. Proof of the main result

At last, we prove our main result:

Theorem 4.1. *Every connected graph G has a tree set of tame finite-order separations that efficiently distinguishes all its inequivalent \aleph_0 -tangles. In particular, equivalent \aleph_0 -tangles induces same orientations on the tree set.*

Proof. For every modified torso H of T we employ Carmesin's Theorem 4.2.11 to obtain a tree set T_H of H -relevant separations that efficiently distinguishes all the ends of H . Then we lift all the separations in all the tree sets T_H and add these to T to obtain an extension T' of T . Then T' is again a tree set by Lemma 4.8.12, Corollary 4.8.13 and Lemma 4.8.14.

4. A tree-of-tangles theorem for infinite-order tangles

First, we show that T' efficiently distinguishes every two inequivalent \aleph_0 -tangles of G . For this, let τ_1 and τ_2 be two inequivalent \aleph_0 -tangles of G . We have to find a separation in T' that efficiently distinguishes τ_1 and τ_2 . Pick some finite-order separation of G with separator Z say that efficiently distinguishes τ_1 and τ_2 . If Z is contained in some critical vertex set of G , then by Lemma 4.8.2 we find a separation in $T \subseteq T'$ that efficiently distinguishes τ_1 and τ_2 . Otherwise Z is not contained in any critical vertex set of G . However, Z is generous by Lemma 4.8.2, and so by Lemma 4.8.3 induces a consistent orientation of T whose part contains Z . Then by Proposition 4.8.24 the \aleph_0 -tangles τ_1 and τ_2 have distinct proxies η_1 and η_2 in the modified torso H of that orientation, and Z efficiently distinguishes η_1 and η_2 in H . Thus there is a separation in T_H of order $|Z|$ that distinguishes η_1 and η_2 . By Lemma 4.8.22 this separation lifts to a separation of G that distinguishes τ_1 and τ_2 . This lift still has order $|Z|$ and lies in T' .

Second we show that all separations in T' are tame. Every separation in T is tame. And by Lemma 4.8.15 the lifts of all T_H are tame as well. \square

Part II.

Stars and combs

5. Duality theorems for stars and combs

I: Arbitrary stars and combs

5.1. Introduction

The series

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 (*)

[26, 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 [26, Lemma 8.1.2] to these classes yields a ray.

Similarly, having a rank in the sense of Schmidt [78] is complementary for infinite graphs to containing a ray, see [26, 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 [26, 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 .*

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 this chapter, 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 the remaining three chapters of this series 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 four chapters in our series, each naming the substructure for which duality theorems are proved in its title:

- I: ARBITRARY STARS AND COMBS (this chapter)
- II: DOMINATING STARS AND DOMINATED COMBS (Chapter 6)
- III: UNDOMINATED COMBS (Chapter 7)
- IV: UNDOMINATING STARS (Chapter 8)

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 parts I–III of our series.

This chapter

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 interderivable.

Our first result is obtained by techniques of Jung [52]. 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. [26].

Theorem 5.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 .

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 [26] for the definition of tree-decompositions and adhesion sets):

Theorem 5.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 [43]: G is rayless if and only if G has a rayless tree-decomposition into finite parts.

While Theorems 5.1 and 5.2 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 \mathcal{C}_X of the components of $G - X$ having their neighbourhood precisely equal to X is infinite.

Theorem 5.3. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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' .

Critical vertex sets were introduced in [62]. 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. [25, 73] and Chapter 3. Moreover, it turns out that Theorem 5.3 implies another characterisation of rayless graphs by Halin [41].

Schmidt's ranking of rayless graphs was employed by Bruhn, Diestel, Georgakopoulos and Sprüssel [13] 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:

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Theorem 5.4. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5.5. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5.3.5. Essentially, the list (\dagger) consists of the following four properties:

- its decomposition tree stems from a normal tree as in Theorem 5.1;
- it has the properties of the tree-decomposition in Theorem 5.2;
- the infinite-degree nodes of its decomposition tree correspond bijectively to the critical vertex sets of G that are relevant in Theorem 5.3;
- the rank of its decomposition tree is equal to the U -rank of G from Theorem 5.4.

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. [26]. Our first duality theorem for stars reads as follows:

Theorem 5.6. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5.1. Theorem 5.6 above compares with Theorem 5.1 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 5.6 exclude each other. To see this, let G be obtained from a ray R by

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 5.1 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 5.3.4 for details).

Our second duality theorem for stars is phrased in terms of tree-decompositions, similar to Theorem 5.2:

Theorem 5.7. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5.2 provides the tools and terminology that we use throughout this series. Section 5.3 and 5.4 are dedicated to the duality theorems for combs and stars respectively.

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

5.2. Tools and terminology

An independent set M of edges in a graph G is called a *partial 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 .

5.2.1. The star-comb lemma

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

Lemma 5.2.1 ([26, 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 5.2.2 ([26, 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, [26, Lemma 8.1.2]. Lemma 5.2.1 has been generalised to higher connectivity, [38, 51, 65], and so has Lemma 5.2.2 in [40, 45, 65]. For an overview we recommend the introduction of [40].

For locally finite trees, Lemma 5.2.2 already yields a comb:

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Lemma 5.2.3. *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 5.2.2 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 5.2.4. *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 5.2.5 (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:*

- (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 [40] introduced a way more complex version of the star-comb lemma above for the more difficult singular case, the *Frayed-Star-Comb Lemma*, [40, Corollary 8.1].

The version for regular cardinals has been proved in, e.g., [30] and [40]. We repeat the short proof here for the sake of convenience:

Proof of Lemma 5.2.5. 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.

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.

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 5.2.2. 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 [55, Lemma III.6.14].

5.2.2. Ends of graphs

We use the notation for ends from the tools of terminology chapter of the dissertation, Chapter 2.

Carmesin [19] observed that

Lemma 5.2.6. *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 [30] (Theorem 5.2.7 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 $c_{X',X}: \mathcal{C}_{X'} \rightarrow \mathcal{C}_X$ for $X' \supseteq X$ send each component of $G - X'$ to the component of $G - X$ containing it; then $\{\mathcal{C}_X, c_{X',X}, \mathcal{X}\}$ is an inverse system whose inverse limit, by construction, consists of the directions.

Theorem 5.2.7 ([30, 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 5.2.8 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:

²For details on inverse limits, see e.g. [36] or [72].

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 $r \geq q$.

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Lemma 5.2.8. *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. [36, 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 5.2.8 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 5.2.8 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.

5.2.3. Critical vertex sets

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

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

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

is infinite (cf. [25, 62] and Chapter 3). 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 $C \in \check{\mathcal{C}}_X$ 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 $\check{\mathcal{C}}_X$ 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:

⁴An overview on this young field is presented in [26, 27].

⁵If G is tough and a covering of $G \sqcup \Omega$ with basic open sets is given, first apply Lemma 5.2.8 to V to obtain a finite subcover \mathcal{O} of Ω , then apply Lemma 5.2.8 to $U = V \setminus \bigcup \mathcal{O}$ to deduce that U is finite and, therefore, $G \setminus \bigcup \mathcal{O}$ is compact.

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Lemma 5.2.9. *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 $\check{\mathcal{C}}_X$ meet U , giving (ii). Otherwise Lemma 5.2.8 gives (i). \square

5.2.4. 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 , [26, 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 [26, Lemma 1.5.5] to possibly non-spanning normal trees):

Lemma 5.2.10. *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 [26, 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 $\lfloor t \rfloor$ 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 $\lfloor t \rfloor \subseteq W$ 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

5. Duality theorems for stars and combs I: Arbitrary stars and combs

first step we conclude that C is a component of $G - W$ meeting T and spanned by $\llbracket x \rrbracket$. \square

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 , [26, Lemma 8.2.3]. More generally,

Lemma 5.2.11. *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 5.2.10 (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 $T - \lceil v_n \rceil$ for which $\llbracket v_{n+1} \rrbracket$ spans the component of $G - \lceil v_n \rceil$ that contains all but finitely many vertices of the comb. Such a vertex v_{n+1} exists by Lemma 5.2.10 (ii). And it is an upward neighbour of v_n , which can be seen by applying Lemma 5.2.10 (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 . \square

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 5.2.12. *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 5.2.10, $\lceil x \rceil$ pairwise separates all the leaves of the star above x at once; let us write L for the set of these leaves. Since $\lceil x \rceil$ is finite, all but finitely many of the infinitely many components in \mathcal{C}_X that meet L are also components of $G - \lceil x \rceil$. 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 $G - \lceil x \rceil$; in other words, X is contained in the chain $\lceil x \rceil$. \square

5.2.5. Containing vertex sets cofinally

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

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 $p \geq x$.

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Lemma 5.2.13. *Let G be any graph. If $T \subseteq G$ is a rooted tree that contains a vertex set W cofinally, then $\partial_T T = \partial_T W$.*

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. $X' := X \cup [X \cap V(T)]_T$, 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 $[v]_T$ 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 $\partial_T T$ and $\partial_T W$ coincide on $\text{crit}(G)$. From $W \subseteq T$ we infer $\partial_T W \subseteq \partial_T T$, 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 $X \in \partial_T T$ be given. We pick, for every component $C \in \mathcal{C}_X$ 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 $X \in \partial_T W$ follows. \square

5.2.6. Tree-decompositions and S -trees

We assume familiarity with [26, 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 [26, 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 5.2.14 ([26, 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 .*

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Corollary 5.2.15. *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. [26]. 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.

5.2.7. 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

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 5.2.16. *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 [19] that solved a conjecture of Diestel [28] from 1992 (in amended form) and, as a corollary, also solved a conjecture of Halin [47] from 1964 (again in amended form):

Theorem 5.2.17 (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 Lemma 6.3.7 in Chapter 6 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 .*

5.3. Combs

Jung [52] 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 ,

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 [26], 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.

5.3.1. Normal trees

In this section we prove

Theorem 5.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5.3.5, of the connected graphs that have normal spanning trees.

Proposition 5.3.1 (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 5.3.2 (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*

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Corollary 5.3.3 (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 5.3.4. *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 5.2.11. \square

Jung's abstract characterisation of the normally spanned graphs goes as follows:

Theorem 5.3.5 (Jung, [52, 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 [32].

Proof of Theorem 5.3.5. The backward implication is provided by Proposition 5.3.1. The forward implication holds as the levels of any normal tree are dispersed, Lemma 5.3.4. \square

We are now ready to prove Theorem 5.1:

Proof of Theorem 5.1. 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 5.3.4, 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 5.3.2 yields a rayless normal tree in G that contains U cofinally. \square

5.3.2. Tree-decompositions

In this section, we show how the rayless normal tree from Theorem 5.1 gives rise to a tree-decomposition that is complementary to combs.

Theorem 5.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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.

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Lemma 5.3.6. *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 5.2.4 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 5.2.14. 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 5.2. By Lemma 5.3.6 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 5.1. 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 5.2 (ii) cannot be strengthened so as to get a star as decomposition tree or to have pairwise disjoint separators:

Example 5.3.7. 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

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Corollary 5.2.15. 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 5.2.14, so V_t is infinite, and hence so is $V_t \cap U = V_t$. But this contradicts our assumptions.

5.3.3. 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 5.3. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5.3 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 5.3 implies Halin's [41, 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 5.3 in two efficient ways:

Combinatorial proof of Theorem 5.3 using Theorem 5.1 or 5.2. Clearly, at most one of (i) and (ii) can hold. And if G contains no comb attached to U , then (ii) holds by Theorem 5.1 with Lemma 5.2.4 or by Theorem 5.2 with Lemma 5.3.6. \square

Inverse limit proof of Theorem 5.3. Lemma 5.2.9 gives \neg (i) \rightarrow (ii). \square

Note that condition (ii) yields a star attached to U .

5.3.4. Rank

In 1983, Schmidt [78] introduced a notion that is now known as the *rank* of a graph, cf. Chapter 8.5 of [26]. 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. [13, 26]. Since the absence of a comb attached to U is equivalent to the existence of a *rayless* normal tree containing U , Theorem 5.1, 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 5.3.8 ([26, 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, [26, 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 5.4. *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.

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Proof of Theorem 5.4. 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 $\overset{\circ}{u}_1 P_1$. 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 5.3.9. *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 5.3.10. *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 5.3.9 (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 5.3.9 (ii)). \square

Lemma 5.3.11. *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 .*

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Proof. The equality (i) = (ii) is the subject of Lemma 5.3.10. Lemma 5.3.9 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 5.3.9 (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 5.2.10 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,

$$(U\text{-rank of } C) \leq (U\text{-rank of } T \cap C) < \alpha. \quad \square$$

5.3.5. 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 5.5. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5.3.12. *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 . \square*

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Here we remark that, in this chapter, we consider Schmidt's ranking of rayless graphs as discussed in Section 5.3.4. 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. [26, p. 242 & 243]).

The proof of the theorem above is organised as follows. First, we will state Proposition 5.3.13, which lists some useful properties of squeezed expansions. Then, we will employ this proposition in a high level proof of Theorem 5.5. 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 5.3.13.

Proposition 5.3.13. *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 5.3.14. *Every normally spanned graph has a rooted tame tree-decomposition displaying its combinatorial remainder. \square*

Now we prove Theorem 5.5 using Proposition 5.3.13 above:

Proof of Theorem 5.5. (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 5.1 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 5.3.13 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 5.2.13, Proposition 5.3.13 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 5.4 and is equal to the rank of T_{NT} by Lemma 5.3.11, which in turn is equal to the rank of T by Proposition 5.3.13. \square

Next, we provide all the definitions needed: First, we extend the definition of 'display' to include critical vertex sets (Definition 5.3.16). Second, we define the 'expansion' of a normal tree (Definition 5.3.17), which is a certain tree-decomposition. Finally we define what it means to 'squeeze' a tree-decomposition (Definition 5.3.18).

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Recall that the definition of ‘display’, as discussed in Section 5.2, 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 $\mathcal{C}_X \setminus \mathcal{C}$ 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 5.3.15. *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 $G[X \cup \bigcup \check{\mathcal{C}}_X]$ 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 5.3.15 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 5.2.7.

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.

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Definition 5.3.16. [Display $\Psi \subseteq \Gamma(G)$] Let G be any graph. A rooted tame $S_{\mathbb{N}_0}$ -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 $\vec{S}_{\mathbb{N}_0}$ 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 5.2.7. 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 $S_{\mathbb{N}_0}$ -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 5.2.7. 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 $S_{\mathbb{N}_0}$ -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 Examples 4.3.6 and 4.3.7 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 5.3.17 (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 $\beta: \vec{E}(T_{\text{NT}}) \rightarrow \vec{S}_{\mathbb{N}_0}$ 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 $S_{\mathbb{N}_0}$ -tree that displays $\partial_\Omega T_{\text{NT}} \subseteq \Omega(G)$.

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 5.3.1).




Figure 5.3.1.: 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 $[t]$.

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 $C \in \check{\mathcal{C}}_{[t]}(X) \subseteq \check{\mathcal{C}}_X$ we delete the edge tC (this is 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 $\alpha(t, X) := (X, \check{\mathcal{C}}_{[t]}(X))$, and for every component $C \in \check{\mathcal{C}}_{[t]}(X)$ 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 5.3.18 (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:

5. Duality theorems for stars and combs I: Arbitrary stars and combs

Lemma 5.3.19. *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 5.3.15 it suffices to show that for every critical vertex set X of G there is a cofinite subset $\mathcal{C} \subseteq \mathcal{C}_X$ 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 $\mathcal{C} \subseteq \mathcal{C}_X$ 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 $\mathcal{C}(i) \subseteq \mathcal{C}_X$ 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 $\mathcal{C} := \bigcap_i \mathcal{C}(i) \subseteq \mathcal{C}_X$. \square

Now that we have formally introduced all the definitions involved, we are ready to prove Proposition 5.3.13:

Proof of Proposition 5.3.13. (i) The expansion is upwards connected by definition.

(ii) Using Lemma 5.2.11 and the fact that every component of $G - T_{\text{NT}}$ has finite neighbourhood, it is straightforward to check that (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 5.2.12 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 5.3.19. 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 5.2.11.

(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

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 [19] 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 5.2.17. 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 5.3.20. *Characterise, for all connected graphs G , the subsets $\Psi \subseteq \Gamma(G)$ for which G admits a rooted tame tree-decomposition displaying Ψ .*

5.4. Stars

5.4.1. Normal trees

In this section we prove a duality theorem for stars in terms of normal trees.

Theorem 5.6. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5.6 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 5.2.3 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 5.3.3 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 6.1 in Chapter 6 which is why its proof is placed in the second chapter of our series, cf. Section 6.2. 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 5.2.3 to T .

5.4.2. 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 5.7. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 $\partial_T U$ 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 5.2.3) 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 5.2.4 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 5.4.1. *In Theorem 5.7 the graph G cannot satisfy both (i) and (ii).*

Proof. Let (T, \mathcal{V}) be a tree-decomposition as in (ii) of Theorem 5.7. 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 5.7 it remains to show $\neg(\text{i}) \rightarrow (\text{ii})$. This time, however, it is harder to see how the normal tree from Theorem 5.6 can be employed to yield a

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 [30, 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 5.4.2. *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 $\partial_\Omega U$ is non-empty and compact by Lemma 5.2.8. \square

Our next lemma generalises the fact that a vertex can be strictly separated from every end which it does not dominate.

Lemma 5.4.3. *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 5.4.4. *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 .*

5. Duality theorems for stars and combs I: Arbitrary stars and combs

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 5.4.3 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 5.2.16 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 5.7. By Lemma 5.4.1 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 5.4.2 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 5.4.4 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 5.4.2 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

6. Duality theorems for stars and combs II: Dominating stars and dominated combs

6.1. Introduction

Two properties of infinite graphs are *complementary* in a class of infinite graphs if they partition the class. In a series of four chapters 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. See Chapter 5 for a comprehensive introduction, and a brief overview of results, for the entire series of four chapters (5, 7, 8 and this chapter).

In the first chapter of this series 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 second chapter of our series 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 [26].)

6. Duality theorems for stars and combs II: Dominated combs

Theorem 6.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5, Theorems 6.2.1 and 6.2.2. 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 6.1, 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 6.1 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 [26]. ‘Essentially disjoint’ and ‘displaying’ are defined in Section 6.3. 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 6.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 6.1, 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 5, Theorems 6.3.1 and 6.3.2.

6. Duality theorems for stars and combs II: Dominated combs

In our proof of Theorem 6.2 we employ a profound theorem of Carmesin [19], which states that every graph has a tree-decomposition displaying all its undominated ends. As it will be the case in this chapter, 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 6.2 establishes our duality theorem for dominated combs in terms of normal trees. In Section 6.3 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 series, Section 5.2.

6.2. Dominated combs and normal trees

In this section we obtain the following duality theorem for dominated combs in terms of normal trees:

Theorem 6.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 6.2).

Before we provide a proof of Theorem 6.1 above, we shall discuss some consequences and applications. As a first consequence, Theorem 6.1 above builds a bridge between the duality theorems for combs (Theorem 6.2.1) and stars (Theorem 6.2.2) in terms of normal trees, which we recall here.

Theorem 6.2.1 (Theorem 5.1). *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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.

6. Duality theorems for stars and combs II: Dominated combs

Theorem 6.2.2 (Theorem 5.6). *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 6.2.1. 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 .

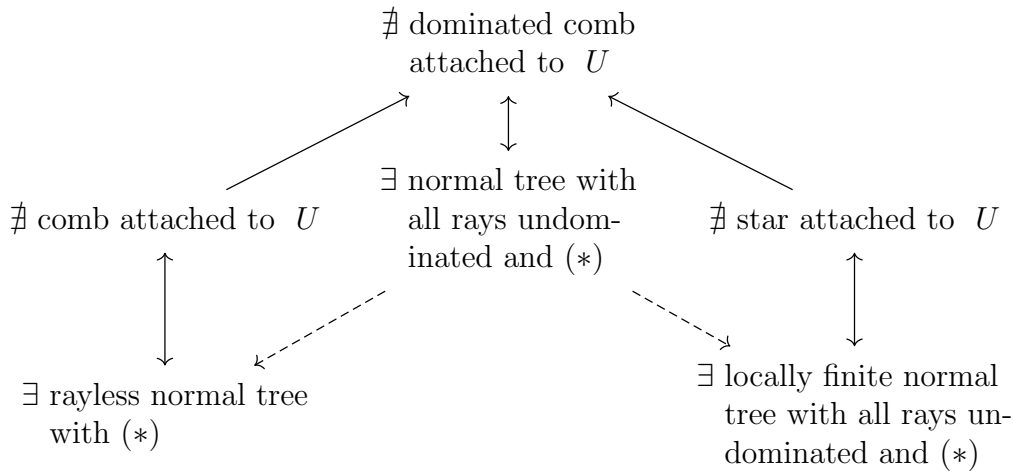


Figure 6.2.1.: The relations between the duality theorems for combs, stars and dominated combs in terms of normal trees.

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 6.2.1 and of Theorem 6.2.2 from the corresponding implication of Theorem 6.1. 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 6.1 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 6.1 yields a normal tree, which additionally must be locally finite and satisfy that all its rays are undominated. Since (i) and (ii) of Theorem 6.2.1 and of Theorem 6.2.2 exclude each other almost immediately

6. Duality theorems for stars and combs II: Dominated combs

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 6.2.2.

We proved Theorem 6.2.2 without its ‘moreover’ part in Chapter 5 of our series. There, instead of proving the ‘moreover’ part as well, we announced that we would prove it in this second chapter of the series. And here we prove it, by deriving it from the identical ‘moreover’ part of Theorem 6.1:

Proof of Theorem 6.2.2, including its ‘moreover’ part. We employ Theorem 6.1 as above. □

Another consequence of Theorem 6.1 is a fact whose previous proof, [24, Lemma 2.3], relied on the theorem of Halin [45] which states that every connected graph without a subdivided K^{\aleph_0} has a normal spanning tree:

Corollary 6.2.3. *If G is a connected graph none of whose ends is dominated, then G is normally spanned.* □

For the proof of Theorem 6.1, we shall need the following four lemmas and a result by Jung (cf. [52, Satz 6] or Theorem 5.3.5). The first lemma is from the first chapter of this series and we remark that the original statement also takes critical vertex sets in the closure of T or W into account.

Lemma 6.2.4 (see Lemma 5.2.13). *Let G be any graph. If $T \subseteq G$ is a rooted tree that contains a vertex set W cofinally, then $\partial_\Omega T = \partial_\Omega W$.*

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 $\lfloor x \rfloor$ 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 $\lfloor x \rfloor$.

Lemma 6.2.5 (Lemma 5.2.10). *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 $\lfloor x \rfloor \cap \lfloor y \rfloor$.*
- (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 $T - W$.*

Lemma 6.2.6 (Lemma 5.2.11). *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 6.2.7. *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 $w \in W$

6. Duality theorems for stars and combs II: Dominated combs

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 partial 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 partial matching. \square

Theorem 6.2.8 (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 6.1. 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 6.2.6 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 6.2.7 contrary to our assumptions. Therefore, all the sets $D_n \cap U$ with $n \in \mathbb{N}$ are dispersed. Now, Jung's Theorem 6.2.8 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 6.1 (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

6. Duality theorems for stars and combs II: Dominated combs

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 $\hat{U} \supseteq U$. For the forward inclusion, consider any end ω of G with $\omega \notin \partial_\Omega U$; we show $\omega \notin \partial_\Omega \hat{U}$. As T' contains U cofinally, it follows from Lemma 6.2.4 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 6.1 (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 6.2.4 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 partial 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 6.2.6. 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 6.2.5, 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

6. Duality theorems for stars and combs II: Dominated combs

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 6.2.9 ([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 6.1 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 6.2.2 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 6.2.9 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 [26, p. 246]).

Lemma 6.2.10. *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 6.2.6, 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 6.2.5 it is straightforward to verify that the bijection is a homeomorphism, no matter which of the three topologies we chose. \square

Theorem 6.2.11. *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 6.2.9 (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 6.2.8, 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

6. Duality theorems for stars and combs II: Dominated combs

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 \bar{T} . By Lemma 6.2.10 the spaces \bar{T} and $|T|$ are homeomorphic. Since T normally spans itself, $|T|$ is metrisable by Theorem 6.2.9 (i).

(ii) Suppose that G contains no dominated comb attached to U . By Theorem 6.1, there is a normal tree $T \subseteq G$ that contains U cofinally. Then $\bar{T} \cong |T|$ by Lemma 6.2.10, and $|T|$ is metrisable by Theorem 6.2.9 (ii).

(iii) If G contains no star attached to U , then by Theorem 6.2.2 there is a locally finite normal tree $T \subseteq G$ that contains U cofinally. By Lemma 6.2.10 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 6.2.9 (i). \square

The statements (ii) and (iii) of Theorem 6.2.11 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 6.2.10 the standard subspace that arises from R is homeomorphic to $|R|$, which in turn is metrisable by Theorem 6.2.9. But $R \subseteq G$ is a dominated comb attached to U .

6.3. Dominated combs and 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 6.2.1 and of Theorem 6.2.2 (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 6.3.1 and Theorem 6.3.2 below)—by following arrows in Figure 6.3.1 like we did in Figure 6.2.1.

Theorem 6.3.1 (Theorem 5.2). *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 5 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

6. Duality theorems for stars and combs II: Dominated combs

$\tau \upharpoonright \Psi: \Psi \rightarrow \Omega(T)$ 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 6.3.2 (Theorem 5.7). *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 6.3.1, 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 6.3.1 true. In Section 6.3.2, the situation is reversed: we will prove a duality theorem making the right but not the left dashed arrow in Figure 6.3.1 true. Here we also provide a short proof of Carmesin’s result [19], which states that every graph has a tree-decomposition displaying all its undominated ends, for normally spanned graphs. Finally, in Section 6.3.3, we will prove a duality theorem that makes both the left and the right dashed arrow in Figure 6.3.1 true. This will be achieved by combining our proof technique from Section 6.3.1 and our duality theorem from Section 6.3.2.

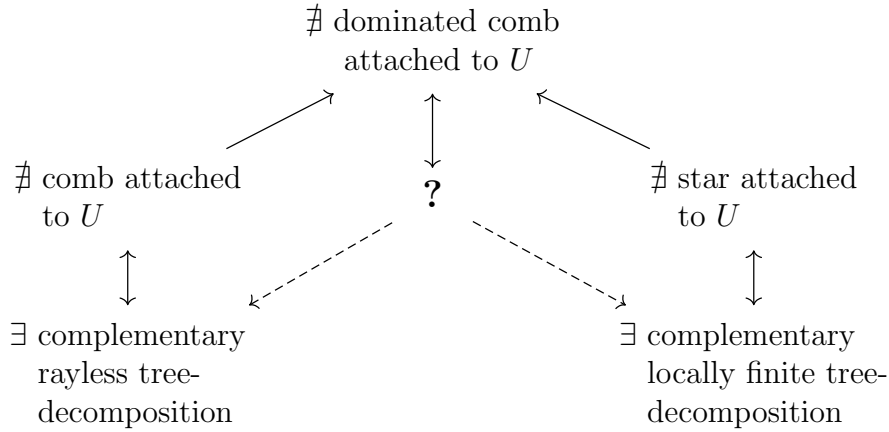


Figure 6.3.1.: The desired relation between stars, combs, dominated combs and complementary tree-decompositions.

The left and right dashed arrow describe an implication whenever there is no comb and no star attached to U , respectively.

6.3.1. 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 6.3.1 true:

6. Duality theorems for stars and combs II: Dominated combs

Theorem 6.3.3. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 6.3.1 from it (also see Figure 6.3.2 which shows the first two columns of Figure 6.3.1 in greater detail and with Theorem 6.3.3 (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 6.3.3 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 6.3.3 (ii) including the theorem’s ‘moreover’ part, we conclude that (T, \mathcal{V}) is as in Theorem 6.3.1 (ii) including the theorem’s ‘moreover’ part.

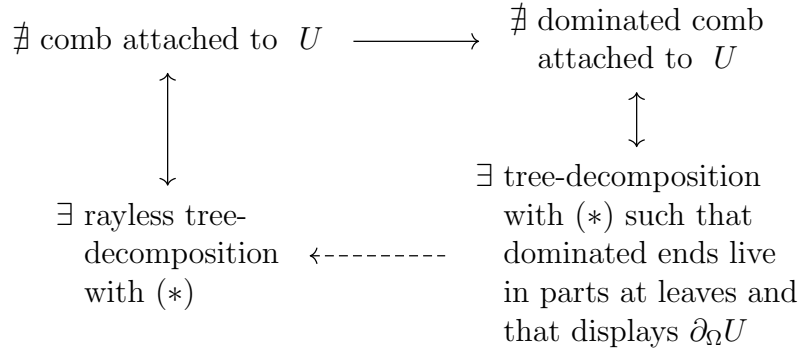


Figure 6.3.2.: The first two columns of Figure 6.3.1 with Theorem 6.3.3 (ii) including the theorem’s ‘moreover’ part inserted for ‘?’.

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.

Finally, we prove Theorem 6.3.3:

Proof of Theorem 6.3.3. 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

6. Duality theorems for stars and combs II: Dominated combs

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.

Now, to show that at least one of (i) and (ii) holds, we show $\neg(\text{i}) \rightarrow (\text{ii})$. By Theorem 6.1 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 $V_C := [t_C]_{T_{\text{NT}}} \cup V(C)$.

Since T_{NT} is normal and contains U cofinally, it follows by standard arguments employing Lemma 6.2.4 and Lemma 6.2.6 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 6.3.4. The tree-decomposition in Theorem 6.3.3 (ii) cannot be chosen to additionally have pairwise disjoint separators, which shows that the theorem does not make the right dashed arrow in Figure 6.3.1 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 6.3.3 (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 6.3.3, (T, \mathcal{V}) is also a tree-decomposition of G complementary to combs as in Theorem 6.3.1. But then (T, \mathcal{V}) cannot have pairwise disjoint separators, as pointed out in Example 5.3.7.

6.3.2. 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 6.3.1 true.

6. Duality theorems for stars and combs II: Dominated combs

Theorem 6.3.5. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 6.3.1 from it (also see Figure 6.3.3 which shows the last two columns of Figure 6.3.1 in greater detail and where Theorem 6.3.5 (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 6.3.5 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 6.3.2 from the identical properties of Theorem 6.3.5 (ii) including that theorem’s ‘moreover’ part.




Figure 6.3.3.: The last two columns of Figure 6.3.1 with Theorem 6.3.5 (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 6.3.5, we will employ the following result by Carmesin. Recall that a rooted $S_{\mathbb{N}_0}$ -tree (T, α) has *upwards disjoint* separators if for every

6. Duality theorems for stars and combs II: Dominated combs

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 6.3.6 (Carmesin 2014, Theorem 5.2.17). *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 [19] is long and complex. However, in this chapter 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 6.3.10). Furthermore, we prove that the separators of the tree-decomposition in Theorem 6.3.6 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 6.3.7. *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 6.3.6 yields the following theorem:

Theorem 6.3.8. *Every connected graph G has a tree-decomposition with pairwise disjoint finite connected separators that displays the undominated ends of G . \square*

For the proof of Lemma 6.3.7 we need the following lemma from the first chapter of our series:

Lemma 6.3.9 (Lemma 5.2.16). *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 6.3.7. 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

6. Duality theorems for stars and combs II: Dominated combs

$\tilde{\alpha}' := \alpha' \upharpoonright E(\tilde{T})$ 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 6.3.9 this correspondence is bijective, and hence $(\tilde{T}, \tilde{\alpha}')$ displays Ψ as desired. \square

Theorem 6.3.10. *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.*

6. Duality theorems for stars and combs II: Dominated combs

Proof. Given the normal tree T_{NT} , by Lemma 6.3.7 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 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 $V(T_{\text{NT}})$. It remains to prove that (T, \mathcal{V}) displays the undominated ends in the closure of T_{NT} .

By Lemma 6.3.9 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 6.3.4. At step $n+1$ suppose that (T_n, α_n) has already been defined and recall that the 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

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 $n \in \mathbb{N}$.

6. Duality theorems for stars and combs II: Dominated combs




Figure 6.3.4.: The construction of the (T_n, α_n) in the proof of Theorem 6.3.10. 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} .

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_{\aleph_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 6.3.10 at hand, we are finally able to prove Theorem 6.3.5:

Proof of Theorem 6.3.5. 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 $\neg(i) \rightarrow (ii)$. Using Theorem 6.1 we find a normal tree $T_{NT} \subseteq G$ that contains U cofinally and all whose

6. Duality theorems for stars and combs II: Dominated combs

rays are undominated in G . Furthermore, by the ‘moreover’ part of Theorem 6.1 we may assume that every component of $G - T_{\text{NT}}$ has finite neighbourhood, and by Lemma 6.2.4 we have $\partial_{\Omega}U = \partial_{\Omega}T_{\text{NT}}$. Then Theorem 6.3.10 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

6.3.3. 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 6.3.1 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 6.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 6.3.5, which shows Figure 6.3.1 in greater detail and where Theorem 6.2 (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 6.2 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 6.3.1 (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 6.2 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

6. Duality theorems for stars and combs II: Dominated combs




Figure 6.3.5.: 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 $\partial_\Omega U$.

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 6.3.5 (ii) including the theorem's 'moreover' part (it covers U cofinally as F meets every rooted ray of T while (T, \mathcal{V}) displays $\partial_\Omega U$). Then, as we have already argued below Theorem 6.3.5, 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 6.3.2 (ii) including the theorem's 'moreover' part.

As we work with contraction minors in the proof of Theorem 6.2 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 $[f](X) := [f(\cup X)]$ 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:

6. Duality theorems for stars and combs II: Dominated combs

Lemma 6.3.11. *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 5.2.7), 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 6.3.12. *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 [26, 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. [26, Lemma 12.3.3].

6. Duality theorems for stars and combs II: Dominated combs

Lemma 6.3.13. *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 6.3.12, 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 6.3.14. *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 6.2 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{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 6.3.13. \square

Proof of Theorem 6.2. 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 6.3.5 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

6. Duality theorems for stars and combs II: Dominated combs

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 6.3.12. By Lemma 6.3.14 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 6.1 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 $\mathcal{T}_{\text{DISJ}}$ whose parts meet T_{NT} induce a subtree T'_{DISJ} of $\mathcal{T}_{\text{DISJ}}$, 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 6.2.4 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 6.3.6. 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 (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 s_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




Figure 6.3.6.: The construction of (T', \mathcal{V}') in the proof of Theorem 6.2. 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 $r(t) = r_y$.

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 6.2.4, it suffices to show that (T', \mathcal{V}') displays the ends in the closure of T_{NT} , which follows from standard arguments. \square

Example 6.3.15. The tree-decomposition in Theorem 6.2 (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

6. Duality theorems for stars and combs II: Dominated combs

no comb attached to U . In particular, as we have already argued in the text below Theorem 6.2, every tree-decomposition (T, \mathcal{V}) of G complementary to dominated combs as in Theorem 6.2 is also a tree-decomposition of G complementary to combs as in Theorem 6.3.1. But then (T, \mathcal{V}) cannot be chosen with pairwise disjoint separators, as pointed out in Example 5.3.7.

7. Duality theorems for stars and combs

III: Undominated combs

7.1. Introduction

Two properties of infinite graphs are *complementary* in a class of infinite graphs if they partition the class. In a series of four chapters 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. See Chapter 5 for a comprehensive introduction, and a brief overview of results, for the entire series of four chapters (5, 6, 8 and this chapter).

In the first chapter of this series 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. [26]. For the definition of tree-decompositions see [26].

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 [26]. 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 the second chapter of our series 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 the third chapter of the series, 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

7. Duality theorems for stars and combs III: Undominated combs

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 7.1. *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 [70,71] and Širáň [79], 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 7.1, such graphs G cannot have a normal spanning tree. There are two known constructions of such graphs, by Seymour and Thomas [76] and by Thomassen [83]. Both these constructions are involved.

As a corollary of Theorem 7.1 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 7.2. *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 [32, 52].

Our duality theorem for undominated combs in terms of rayless trees, Theorem 7.1, has two applications, Theorems 7.3 and 7.5 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 [26]).

7. Duality theorems for stars and combs III: Undominated combs

Carmesin [19] 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 7.1 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 7.3. *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 7.4. *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. [12]. Our second application of Theorem 7.1 reads as follows:

Theorem 7.5. *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 [12] for definitions. The space \tilde{G} was used by Bruhn and Diestel [12] to extend Whitney’s theorem [26, 84]—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

7. Duality theorems for stars and combs III: Undominated combs

and Kühn [31, Problem 7.9], they showed that such spanning trees always exist in connected finitely separable graphs. Our Theorem 7.5 provides an alternative proof:

Corollary 7.6. *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 7.7. *Let G be any connected 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 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 7.8. *Let G be any connected 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 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 7.2, we prove our duality theorem for undominated combs in terms of rayless trees, Theorem 7.1. In Section 7.3, we discuss our applications of this duality theorem, i.e., we prove Theorem 7.3 and Theorem 7.5. In Section 7.4, we provide our two full duality theorems for undominated combs: Theorem 7.7 and Theorem 7.8.

Throughout this chapter, $G = (V, E)$ is an arbitrary graph. We use the graph theoretic notation of Diestel’s book [26], and we assume familiarity with the tools and terminology described in the first chapter of this series, Section 5.2.

7.2. Undominated combs and 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 Theorem 5.1 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 7.2.1 (Lemma 5.2.4). *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 [79] 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 [76]. 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 7.2.2 ([76, 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 7.2.2 is dominated as G is infinitely connected, but for $U = V(G)$ the graph does not contain a rayless tree containing U .

A similar counterexample has been obtained independently by Thomassen [83]. Set-theoretic points of view are presented in both [76] and Komjáth's [54]. 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 7.2.3. *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.

7. Duality theorems for stars and combs III: Undominated combs

Proofs are due to Polat [70, 71] and Širáň [79]. Our main result in this section extends Theorem 7.2.3:

Theorem 7.1. *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 7.2.3 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 7.1 (to the knowledge of the authors), is due to Polat [67]: *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 7.1, 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 7.1, 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 7.2.4. *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 7.2. *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 7.2.1 we will prove Theorem 7.1 for normally spanned graphs. Then, in Section 7.2.2, we will deduce Theorem 7.1.

7.2.1. Proof for normally spanned graphs

As a first approximation to Theorem 7.1 we prove the following:

Theorem 7.2.5. *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 .

7. Duality theorems for stars and combs III: Undominated combs

Our proof consists of three key ideas, organised in three lemmas: Lemma 7.2.6, Lemma 7.2.7 and Lemma 7.2.9.

Lemma 7.2.6 (Lemma 5.2.13). *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 7.2.7. *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 (cf. Theorem 5.3.5):

Theorem 7.2.8 (Jung). *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 7.2.9. *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 7.2.9. First we choose a well-ordering of U all whose proper initial segments are dispersed: By Theorem 7.2.8, 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

7. Duality theorems for stars and combs III: Undominated combs

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 7.2.6 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 7.2.6 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 7.2.1.) As u^* dominates R we find an infinite set \mathcal{Q} of u^* - R paths in




Figure 7.2.1.: The situation in the last paragraph of the proof of Lemma 7.2.9.

7. Duality theorems for stars and combs III: Undominated combs

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 $\dot{w}R$. 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 7.2.5. By Lemma 7.2.1 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 Lemma 7.2.7 we may assume that U contains all vertices dominating an end in the closure of U , and by Lemma 7.2.9 there is a rayless tree $T \subseteq G$ that contains U . \square

7.2.2. 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 7.2.5, 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 7.2.5 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 7.2.7, 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 7.2.5 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 7.2.9.

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 7.2.9 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 7.2.10. 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 [32]). 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 [32], 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

7. Duality theorems for stars and combs III: Undominated combs

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 $U \subseteq V(G)$.

Lemma 7.2.11. *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 $W_1 W_2$ 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 7.2.12. *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 7.1. 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 7.2.13 (5.2.10). *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 7.2.14 (Lemma 5.2.11). *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 .*

7. Duality theorems for stars and combs III: Undominated combs

Proof of Theorem 7.1. 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 7.2.1 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 7.2.6 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 7.2.12, the tree $[T]$ is normal in H , and it has vertex set $[U]$ because $V(T) = U$. So, by Lemma 7.2.9, the graph H contains a rayless tree that contains $[U]$. Finally, by Lemma 7.2.11, 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 7.2.13). 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 7.2.12, the tree $[T] \subseteq H$ is normal in H and $V([T]) = [U]$ since $V(T) = U$. And by Lemma 7.2.14 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

7. Duality theorems for stars and combs III: Undominated combs

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 7.2.2.) For every path $P \in \mathcal{P}$ let $P' = P'(P) := v_P P w_P Q_P$, and let $\mathcal{P}' = \mathcal{P}'(\mathcal{P}) := \{P' \mid P \in \mathcal{P}\}$.




Figure 7.2.2.: The final case in the proof of our duality theorem for undominated combs in term of rayless trees.

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 7.2.13, this component is of the form $\llbracket t \rrbracket$ for the successor

7. Duality theorems for stars and combs III: Undominated combs

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

7.3. Spanning trees reflecting the undominated ends

In [47], 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 7.2.2 shows that his conjecture is in general false. Recently, Carmesin [19] amended Halin's conjecture by proving the following:

Theorem 7.3.1 (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 7.2.2 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 7.3 below which characterises the graphs that have spanning trees reflecting their undominated ends. Thereafter, we will characterise in Theorem 7.5 (i) the spanning trees of finitely separable graphs that reflect the undominated ends, and we will establish in Theorem 7.5 (ii) that every connected finitely separable graph has such a tree. Finally, we will deduce Corollary 7.6 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:

7. Duality theorems for stars and combs III: Undominated combs

Theorem 7.3. *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 7.3 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)→(i) yields the following theorem:

Theorem 7.3.2. *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 7.3 can be deduced from Theorem 7.3.2. The implication (i)→(ii) of Theorem 7.3 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 7.3 (ii). Then Theorem 7.3.2 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 in order to prove Theorem 7.3 we may equivalently prove Theorem 7.3.2:

Proof of Theorem 7.3. Employ Theorem 7.3.2 as above. □

Furthermore, the case $U = V(G)$ of Theorem 7.3.2 establishes our second main corollary:

Corollary 7.4. *Every graph that has a normal spanning tree does have a spanning tree reflecting its undominated ends.*

Our proof of Theorem 7.3.2 requires some preparation. First, we need the following strengthening of a structural result by Carmesin. Recall from Chapter 5

7. Duality theorems for stars and combs III: Undominated combs

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 7.3.3 (Theorem 6.2.8). *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 7.3.4. *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 7.3.3, 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 7.3.5. *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 5.2.7), 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 7.3.6 (Lemma 6.3.12). *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 .*

7. Duality theorems for stars and combs III: Undominated combs

Lemma 7.3.7. *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 [44, Lemma 7.2 (b)].

Proof. Without loss of generality both G and H are connected. By Theorem 7.2.8, 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 7.2.8. \square

We need one more lemma for the proof of Theorem 7.3.2:

Lemma 7.3.8. *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 7.3.2. Let G be any graph and let $U \subseteq V(G)$ be normally spanned. Without loss of generality, G is connected. By Theorem 7.3.4 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 7.3.8 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 .

7. Duality theorems for stars and combs III: Undominated combs

As U is normally spanned in G it follows by Lemma 7.3.7 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 7.3.6 every end of H_t in the closure of $[U_t]$ is dominated in H_t . Hence we may apply Theorem 7.1 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 7.2.6 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 7.2.1 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 7.3.2, let us turn to an application.

Theorem 7.5. *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 7.5, we show a corollary for the topological space \tilde{G} (see [12] for definitions regarding \tilde{G}).

Corollary 7.6. *Every connected finitely separable graph G has a spanning tree whose closure in \tilde{G} contains no topological cycle.*

Proof. By Theorem 7.5 (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 \tilde{T} and fix an edge e of T that is contained in C as a topological edge. Let F_e be the fundamental

7. Duality theorems for stars and combs III: Undominated combs

cut of e with respect to T and let us write V_1 and V_2 for the two sides of F_e . Then $C \setminus \dot{e}$ 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 $|G| \setminus \bigcup \{ \overset{\circ}{f} \mid f \in F_e \}$ that is divided into the two closed disjoint sets $\overline{G[V_1]}$ and $\overline{G[V_2]}$ (contradiction). \square

Proof of Theorem 7.5. (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 partial 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 partial 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 [45] 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

7. Duality theorems for stars and combs III: Undominated combs

spanning tree in a connected finitely separable graph directly, as follows. Every 2-connected finitely separable graph G is countable, cf. [82] or [12, 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 5.2.5, 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 5.3.3), 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 5.3.3). 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 7.3.2 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

7.4. 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 7.7. *Let G be any connected 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 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 7.3.3 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 $[W]$ 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

7. Duality theorems for stars and combs III: Undominated combs

The central part of the star-decomposition in Theorem 7.7 (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 7.8. *Let G be any connected 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 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 7.7 (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

7. *Duality theorems for stars and combs III: Undominated combs*

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

8. Duality theorems for stars and combs

IV: Undominating stars

8.1. Introduction

Two properties of infinite graphs are *complementary* in a class of infinite graphs if they partition the class. In a series of four chapters 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. See Chapter 6 for a comprehensive introduction, and a brief overview of results, for the entire series of four chapters (5, 6, 7 and this chapter).

In the first chapter of this series 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. [26]. For the definition of tree-decompositions see [26].

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 the second chapter of our series 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 the third chapter of the series 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

8. Duality theorems for stars and combs IV: Undominating stars

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 the fourth and final chapter of the series, 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 8.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 8.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 [62]. 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. [25, 73] and Chapter 3. For the definitions of 'tame' and 'live', see Section 8.3. 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.

8. Duality theorems for stars and combs IV: Undominating stars

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, Chapter 4, of Robertson and Seymour's tree-of-tangles theorem from their graph-minor series [73].

This chapter is organised as follows. Section 8.2 establishes our duality theorem for undominating stars in terms of end-compactified subgraphs. Section 8.3 proves our duality theorem for undominating stars in terms of star-decompositions. In Section 8.4 we summarise the duality theorems of the complete series.

Throughout this chapter, $G = (V, E)$ is an arbitrary graph. We assume familiarity with the tools and terminology described in the first chapter of this series, Section 5.2.

8.2. Tough subgraphs

In this section, we prove our duality theorem for undominating stars in terms of tough subgraphs:

Theorem 8.1. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 8.2.1, and that (ii) is equivalent to U being tough in G will be shown in Theorem 8.2.2. 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 8.2.1. *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 8.2.2. *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 8.1. Combine Lemma 8.2.1 and Theorem 8.2.2 above. □

While the proof of Theorem 8.2.2 takes the rest of this section, that of Lemma 8.2.1 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

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

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 .

8. Duality theorems for stars and combs IV: Undominating stars

Proof of Lemma 8.2.1. If U is tough in G then no critical vertex set of G lies in the closure of U . We know by Lemma 5.2.9 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 8.2.2, we summarise a few elementary properties that are complementary to containing an undominating star attached to a given vertex set U :

Lemma 8.2.3. *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 8.2.1 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 [62] 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 8.2.2. 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 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 \mathcal{C}_X

8. Duality theorems for stars and combs IV: Undominating stars

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 Chapter 4 for critical vertex sets in graphs, Theorems 8.2.6 and 8.2.7 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. [29]. 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 8.2.4 (Corollary 4.2.10). *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 $C_Y(X)$.

8. Duality theorems for stars and combs IV: Undominating stars

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 8.2.5 (Definition 4.5.9). 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 $\mathcal{K}(X) \subseteq \mathcal{C}_X$ is called *admissable* for \mathcal{Y} if for every two $X, Y \in \mathcal{Y}$ that are incomparable as sets we have either $C_X(Y) \notin \mathcal{K}(X)$ or $C_Y(X) \notin \mathcal{K}(Y)$. If additionally $|\mathcal{C}_X \setminus \mathcal{K}(X)| \leq 1$ for all $X \in \mathcal{Y}$, then \mathcal{K} is *strongly admissable* for \mathcal{Y} .

Theorem 8.2.6 (Theorem 4.5.10). *For every principal collection of vertex sets of a connected graph there is a strongly admissable function.*

Theorem 8.2.7 (Theorem 4.5.11). *Let G be any connected graph, let \mathcal{Y} be any principal collection of vertex sets of G and let \mathcal{K} be any admissable function for \mathcal{Y} . Then for every distinct two $X, Y \in \mathcal{Y}$, after possibly swapping X and Y ,*

$$\text{either } (\mathcal{K}(X), X) \leq (Y, \mathcal{K}(Y)) \text{ or } (\mathcal{K}(X), X) \leq (C_Y(X), Y) \leq (\mathcal{K}(Y), Y).$$

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 admissable 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 8.2.7, 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 8.2.8. *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 \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 \mathcal{C}_X meet U ; we write \mathcal{F}_X for this finite collection. Theorem 8.2.6 yields a strongly admissable 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 admissable for $\text{crit}(G)$, and $\mathcal{K}(X)$ is a cofinite subcollection of $\mathcal{C}_X \setminus \mathcal{F}_X$ for all X . Now Theorem 8.2.7 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 8.2.2. 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

8. Duality theorems for stars and combs IV: Undominating stars

every vertex set $U \subseteq V(G)$ that is tough in G there is a tough subgraph of G containing U . By Corollary 8.2.8 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 8.2.4 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 $\mathcal{K}(X) \subseteq \mathcal{C}_X$ 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,

8. Duality theorems for stars and combs IV: Undominating stars

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 $\ell Q_{e'}$ does meet ℓQ_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 ℓQ_e . But then the graph $x_{\alpha'} P_{e'} \ell \cup \ell Q_{et} P_{e'} y_{\alpha'}$ 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

8.3. Star-decompositions

In this section we prove our second main result, a duality theorem for undominating stars in terms of star-decompositions, Theorem 8.2 below.

Before we state the theorem, let us recall the following definitions from Section 5.3.5. 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 8.3.1 (Lemma 5.3.15). *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 8.3.1, if X is a critical vertex set of G and (T, α) is a tame

8. Duality theorems for stars and combs IV: Undominating stars

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 8.2. *Let G be any connected graph, and let $U \subseteq V(G)$ be any vertex set. 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 8.3.2 below, and then show how it implies our main result, Theorem 8.2 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 8.3.2. *Let G be any graph, and let $(T, \mathcal{Y}, \mathcal{K})$ be any principal tree set so that $O_{\mathcal{K}}$ 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 $O_{\mathcal{K}}$ is included in the part of σ ;
- (ii) for every $\vec{s} \in O_{\mathcal{K}}$ there is some $\vec{r} \in \sigma$ with $\vec{s} \lesssim \vec{r}$.

The technical theorem implies our second main result, Theorem 8.2:

Proof of Theorem 8.2. First, we show that at most one of (i) and (ii) holds. By Lemma 8.2.3 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 8.2.3 we know that U is tough in G . Then, by Corollary 8.2.8, we find a principal tree set $(T, \text{crit}(G), \mathcal{K})$ such that, for every critical vertex set X , no element of $\mathcal{K}(X)$ meets U and the inclusion $\mathcal{K}(X) \subseteq \mathcal{C}_X$ is cofinite. We claim that the star provided by Theorem 8.3.2 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 $(\mathcal{K}(X), X)$ with X critical and $\mathcal{K}(X)$ a cofinite subset of \mathcal{C}_X are tame and thus our star-decomposition is tame. Next, by Theorem 8.3.2 (i), we have that U is contained in the central part of the star-decomposition. Finally, by Theorem 8.3.2 (ii), every critical vertex set of G lives in a leaf's part. \square

8. Duality theorems for stars and combs IV: Undominating stars

Next, we prepare the proof of our technical theorem, Theorem 8.3.2. We will need the following concept of a corridor from Chapter 4. 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. Lemma 4.7.1 and Definition 4.7.2. 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 8.3.3. *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 8.3.4. *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 8.3.5. *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

8. Duality theorems for stars and combs IV: Undominating stars

Before we start with the proof of Theorem 8.3.2 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 8.3.5, 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 8.3.4) 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 $\vec{S}_{\mathbb{N}_0}[\pi(O)]$ that is a subsystem of $\vec{S}_{\mathbb{N}_0}$ whose separations are all nested with each other. Furthermore, $\vec{S}_{\mathbb{N}_0}[\pi(O)]$ and \vec{T} are nested with each other in $\vec{S}_{\mathbb{N}_0}$ by Lemma 8.3.3. Also, the parliament of O is a consistent orientation of $\vec{S}_{\mathbb{N}_0}[\pi(O)]$ 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 8.3.6. *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 $S_{\mathbb{N}_0}$ -trees:

Theorem 8.3.7. *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} such that $\vec{S}_{\mathbb{N}_0}[\pi(O)]$ is regular. Then for every corridor γ of the parliament of O the corresponding regular tree set $\vec{S}_{\mathbb{N}_0}[\gamma]$ is isomorphic to the edge tree set of a tree.*

Proof. Let γ be any corridor of the parliament of O . By Theorem 2.3.1, it suffices to show that $\vec{S}_{\mathbb{N}_0}[\gamma]$ 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 $\vec{S}_{\mathbb{N}_0}[\gamma]$.

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 a number $N < \omega$ 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 8.3.6.

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

8. Duality theorems for stars and combs IV: Undominating stars

separation $\vec{r} \in \gamma$ with $\overleftarrow{s}_\omega \leq \vec{r}$ and $\vec{s}_0 \leq \vec{r}$. For every $n < \omega$, either $\vec{s}_n \leq \vec{r}$ or $\overleftarrow{s}_n \leq \overleftarrow{r}$ or $\overleftarrow{s}_n \leq \vec{r}$ or $\overleftarrow{s}_n \leq \overleftarrow{r}$. We cannot have $\vec{s}_n \leq \overleftarrow{r}$ for any n , since this would imply $\vec{s}_0 < \vec{s}_n \leq \overleftarrow{r} \leq \overleftarrow{s}_0$ contradicting that $\vec{S}_{\aleph_0}[\pi(O)]$ is regular. We cannot have $\overleftarrow{s}_n \leq \vec{r}$ for any n because γ is consistent. And we cannot have $\overleftarrow{s}_n \leq \overleftarrow{r}$, because then $\overleftarrow{s}_\omega \leq \vec{r} \leq \vec{s}_n < \overleftarrow{s}_\omega$ 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 8.3.2. Let $(T_{\mathcal{X}}, \mathcal{Y}, \mathcal{K})$ 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 $\bigcup \mathcal{K}(X)$ 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 8.3.7 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 Section 5.2.7).

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

$$\sigma_\gamma := \{ \vec{s}_\gamma^0 \} \cup \{ \vec{s}_\gamma^n \wedge \overleftarrow{s}_\gamma^{n-1} : n \geq 1 \}. \quad (8.3.1)$$

(See Figure 8.3.1.)

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

$$\vec{s}_\gamma^m \wedge \overleftarrow{s}_\gamma^{m-1} \leq \overleftarrow{s}_\gamma^{m-1} \leq \overleftarrow{s}_\gamma^n \leq \overleftarrow{s}_\gamma^n \vee \vec{s}_\gamma^{n-1} = (\vec{s}_\gamma^n \wedge \overleftarrow{s}_\gamma^{n-1})^*.$$

8. Duality theorems for stars and combs IV: Undominating stars




Figure 8.3.1.: 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 $(A, B) := \vec{s}_\gamma^n \wedge \vec{s}_\gamma^{n-1}$ from the proof of Theorem 8.3.2.

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 [19] 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 (8.3.1). 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 $\vec{r} \in \vec{S}_{\mathbb{N}_0}[\gamma]$ 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 8.3.6, 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 8.3.2 (the ray $R_\gamma = v_\gamma^0 v_\gamma^1 \dots$ was defined right above (8.3.1)).

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

8. Duality theorems for stars and combs IV: Undominating stars




Figure 8.3.2.: The orientation of the image $\vec{S}_{N_0}[\gamma]$ of α_γ and the path P in the second step of the proof of Theorem 8.3.2.

case that \vec{s}_γ^0 instead of \vec{s}_γ^0 lies in the orientation that defines t . In this case we have either $\vec{s}_\gamma^0 \leq (\mathcal{H}(X), X)$ or $(\mathcal{H}(X), X) < \vec{s}_\gamma^0$. But actually, we cannot have $\vec{s}_\gamma^0 \leq (\mathcal{H}(X), X)$ because otherwise $(\mathcal{H}(X), X) \leq \vec{s}_X$ would imply that $\vec{s}_\gamma^0 \leq \vec{s}_X$ meaning that \vec{s}_γ^0 and \vec{s}_X violate the consistency of γ . Therefore, we must have $(\mathcal{H}(X), X) < \vec{s}_\gamma^0$, and then we are done because \vec{s}_γ^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 $\vec{r}_P \leq (\mathcal{H}(X), X) \leq \vec{s}_X$ would follow (here we use that $\vec{S}_{N_0}[\gamma] \subseteq \vec{S}_{N_0}[O]$ is regular to ensure $\vec{r}_P \neq \vec{s}_X$). 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 $\vec{r}_P \leq \vec{s}_\gamma^n$ and $\vec{r}_P \leq \vec{s}_\gamma^{n-1}$. Thus,

$$(\mathcal{H}(X), X) \leq \vec{r}_P \leq \vec{s}_\gamma^n \wedge \vec{s}_\gamma^{n-1} \in \sigma.$$

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 $\vec{s}_\gamma^{n-1} \leq (\mathcal{H}(X), X)$ or $(\mathcal{H}(X), X) < \vec{s}_\gamma^{n-1}$, and we have either $\vec{s}_\gamma^n \leq (\mathcal{H}(X), X)$ or $(\mathcal{H}(X), X) < \vec{s}_\gamma^n$. The case $\vec{s}_\gamma^n \leq (\mathcal{H}(X), X)$ is impossible since otherwise $(\mathcal{H}(X), X) \leq \vec{s}_X \in \gamma$ would imply that $\vec{s}_\gamma^n \leq \vec{s}_X$ meaning that \vec{s}_γ^n and \vec{s}_X violate the consistency of γ . Therefore, we have either $(\mathcal{H}(X), X) \leq \vec{s}_\gamma^n \wedge \vec{s}_\gamma^{n-1} \in \sigma$ as desired, or we have $\vec{s}_\gamma^{n-1} \leq (\mathcal{H}(X), X) < \vec{s}_\gamma^n$. For this latter case, we show that there is a component $C \in \mathcal{H}(X)$ such that $\vec{s}_\gamma^{n-1} \leq (C, X)$ holds. This suffices to complete the proof, because then the inequalities $(\mathcal{H}(X) \setminus \{C\}, X) \leq (X, C) \leq \vec{s}_\gamma^{n-1}$ and $(\mathcal{H}(X) \setminus \{C\}, X) \leq (\mathcal{H}(X), X) < \vec{s}_\gamma^n$

8. Duality theorems for stars and combs IV: Undominating stars

give

$$(\mathcal{K}(X) \setminus \{C\}, X) \leq \vec{s}_\gamma^n \wedge \overleftarrow{s}_\gamma^{n-1} \in \sigma.$$

The separation $\vec{s}_\gamma^{n-1} \in O$ is, by definition, the supremum of some corridor δ of $\{(A, B) \in O_{\mathcal{K}} : |A \cap B| \leq \ell\}$ for some number $\ell \in \mathbb{N}$. Then every separation $(\mathcal{K}(Y), Y) \in \delta$ satisfies $(\mathcal{K}(Y), Y) \leq \vec{s}_\gamma^{n-1} \leq (\mathcal{K}(X), X)$. In particular, as the principal tree set $T_{\mathcal{K}}$ satisfies the conclusions of Theorem 8.2.7, every separation $(\mathcal{K}(Y), Y) \in \delta$ satisfies $(\mathcal{K}(Y), Y) \leq (C_X(Y), X)$. Hence in order to show that $\vec{s}_\gamma^{n-1} \leq (C, X)$ for some component $C \in \mathcal{K}(X)$, 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

8.4. Overview of all duality results

In this section we summarise all duality theorems of this series. 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:

8. Duality theorems for stars and combs IV: Undominating stars

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 $\partial_\Omega U$.

8. Duality theorems for stars and combs IV: Undominating stars

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 .

9. End-faithful spanning trees in graphs without normal spanning trees

9.1. Introduction

Schmidt [26, 78] 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 [13, 26] proved the unfriendly partition conjecture for the class of rayless graphs in this way.

At the turn of the millennium, Halin [44] 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 9.1. *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 [77]. 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 natural 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 [47] conjectured that every connected graph has an end-faithful spanning tree. However, Seymour and Thomas [76] and Thomassen [83] 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. [26]. 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 [67] and directly addresses the counterexamples by Seymour and Thomas and by

9. End-faithful spanning trees in graphs without normal spanning trees

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 [26].

Theorem 9.2. *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 [77], 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 [32]. 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 [32, 66], and every T_{\aleph_1} with tops contains T_{\aleph_1} as a subgraph.

Carmesin [19] 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 [26].

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 $\eta \subseteq \omega$. Recently, it has been shown in Chapter 7 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 9.3. *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 (Chapter 7) or if G does not contain a subdivision of T_{\aleph_1} [67]. The following corollary extends these results, and any T_{\aleph_1} with all tops witnesses that this extension is proper.

Corollary 9.4. *For every normally traceable graph G , having a rayless spanning tree is equivalent to all the ends of G being dominated.* \square

9. End-faithful spanning trees in graphs without normal spanning trees

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 9.5. *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. Section 9.2 provides the tools and terminology that we use throughout this chapter. In Section 9.3 we introduce the κ -rank and prove Theorem 9.1. Then, in Section 9.4 we introduce the normal rank and prove Theorem 5. We prove Theorem 9.2 in Section 9.5 and we prove Theorem 9.3 in Section 9.6.

9.2. Tools and terminology

A 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 $p \geq x$. We say that a rooted tree $T \subseteq G$ contains a set U *cofinally* if $U \subseteq V(T)$ and U is cofinal in the tree-order of T . We remark that the original statement of the following lemma also takes critical vertex sets in the closure of T or U into account.

Lemma 9.2.1 (Lemma 5.2.13). *Let G be any graph. If $T \subseteq G$ is a rooted tree that contains a vertex set U cofinally, then $\partial_\Omega T = \partial_\Omega U$.*

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*.

Lemma 9.2.2 (Lemma 5.2.11). *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 .*

Given any graph G , a set $U \subseteq V(G)$ of vertices is *dispersed* in G if there is no end in the closure of U in G . Equivalently, U is dispersed if and only if G contains no comb with all its teeth in U . In [52], Jung proved that normally spanned sets of vertices can be characterised in terms of dispersed vertex sets:

Theorem 9.2.3 (Jung [52, Satz 6]; Theorem 5.3.5). *Let G be any graph. A vertex set $U \subseteq V(G)$ 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.*

9.3. 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 [26, 78].

The κ -rank behaves quite similar to Schmidt's rank [26, 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 9.1. *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 [77] 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 [30].

9. End-faithful spanning trees in graphs without normal spanning trees

Theorem 9.3.1 ([77, 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 9.1:

Proof of Theorem 9.1. 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 9.3.1 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

9.4. 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 9.5. *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.

9. End-faithful spanning trees in graphs without normal spanning trees

Lemma 9.4.1. *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 9.2.3, 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 9.1, and G has no normal spanning tree as pointed out by Diestel and Leader [32]. \square

Lemma 9.4.2. *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 9.2.3, 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 9.2.3. 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 9.5. 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 .

9. End-faithful spanning trees in graphs without normal spanning trees

(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 9.2.3; 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 9.2.3. 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 9.2.3, and the rank of T is $\leq \alpha$ because every component of $T - t_*$ has rank $< \alpha$. \square

9.5. 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.

Lemma 9.5.1. *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

9. End-faithful spanning trees in graphs without normal spanning trees

Lemma 9.5.2. *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 $\lceil C \rceil_T$ of T . Note that $R_T \subseteq \lceil C \rceil_T$ 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 9.2.1. 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 9.5.1. 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

9. End-faithful spanning trees in graphs without normal spanning trees

Lemma 9.5.3. *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 9.5.4. *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 9.2.3, 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 9.2.3. Therefore $X \cup Y$ witnesses that $G[C \cup X]$ has normal rank $\leq \xi$ in G . Finally, Lemma 9.4.2 (ii) implies that $G[C \cup X]$ has normal rank $\leq \xi$. \square

Theorem 9.2. *Every normally traceable graph has an end-faithful spanning tree.*

Proof. By Theorem 9.5 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 9.4.2 (i). Note that, by Lemma 9.2.2, the tree T_{NT} reflects the ends of G in the closure of X .

By Lemma 9.5.3 (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 9.5.3 (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 9.5.1 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 $\partial_\Omega C \setminus \partial_\Omega X$.

9. End-faithful spanning trees in graphs without normal spanning trees

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 9.5.4 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 9.5.2 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

9.6. 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 9.6.1 (Theorem 7.3.2). *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 9.3. *Every normally traceable graph has a spanning tree that reflects its undominated ends.*

Proof. By Theorem 9.5 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 9.6.1, 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 9.4.2 (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 9.6.1 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 9.5.3 (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

9. *End-faithful spanning trees in graphs without normal spanning trees*

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 9.5.1.

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 9.5.4 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 9.5.2 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

Part III.
End spaces

10. Approximating infinite graphs by normal trees

10.1. Introduction

A rooted tree T contained in a graph G is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T . (In finite graphs, normal spanning trees are their depth-first search trees; see [26] for precise definitions.) Normal spanning trees are perhaps the most useful structural tool in infinite graph theory. Their importance arises from the fact that they capture the separation properties of the graph they span, and so in many situations it suffices to deal with the much simpler tree structure instead of the whole graph. For example, the end space of G coincides, even topologically, with the end space of any normal spanning tree of G . However, not every connected graph has a normal spanning tree, and the structure of graphs without normal spanning trees is still not completely understood [11, 32].

In order to harness and transfer the power of normal spanning trees to arbitrary connected graphs G , one might try to find an ‘approximate normal spanning tree’: a normal tree in G which spans the graph up to some arbitrarily small given error term. To formalize this idea, recall that, as usual, a *neighbourhood* of an end is the component of $G - X$ which contains a tail of every ray of that end, for some (arbitrarily large) finite set of vertices $X \subseteq V(G)$. We say that a graph G can be *approximated by normal trees* if for every selection of arbitrarily small neighbourhoods around its ends there is a normal tree $T \subseteq G$ such that every component of $G - T$ is included in one of the selected neighbourhoods and every end of G has some neighbourhood in G that avoids T .

Our approximation result for normal trees in infinite graphs then reads as follows:

Theorem 10.1. *Every connected graph can be approximated by normal trees.*

Note that the normal trees provided by our theorem will always be rayless.

We indicate the potential of Theorem 10.1 by a number of applications. Our first two applications are of combinatorial nature: we exhibit in Section 10.4 two new existence results for normal spanning trees that Theorem 10.1 implies. One of these, Theorem 10.4.3, says that if every end of a connected graph G has a neighbourhood which has a normal spanning tree then G itself has a normal spanning tree.

Interestingly, Theorem 10.1 may not only be read as a structural result for connected graphs: it also implies and extends a number of previously hard results about topological properties of end spaces [24, 28, 30, 68, 69, 80]. Denote by $\Omega(G)$ the end space of a graph G , and by $|G|$ the space on $G \cup \Omega(G)$ naturally associated with the graph G and its ends; see the next section for precise definitions. When G is locally finite and connected, then $\Omega(G)$ is compact, and $|G|$ is the well-known Freudenthal compactification of G . For arbitrary G , the spaces $\Omega(G)$ and $|G|$ are usually non-compact and far from being completely understood.

10. Approximating infinite graphs by normal trees

Polat has shown that $\Omega(G)$ is ultrametrizable if and only if G contains a topologically end-faithful normal tree [68, Theorem 5.13], and has proved as a crucial auxiliary step that end spaces are always collectionwise normal [68, Lemma 4.14]. Changing focus from $\Omega(G)$ to $|G|$, Sprüssel has shown that $|G|$ is normal [80], and Diestel has characterised when $|G|$ is metrizable or compact [24] in terms of certain normal spanning trees in G . Our combinatorial Theorem 10.1 provides, in just a few lines, new and unified proofs for all these results. Additionally, Theorem 10.1 shows that metrizable end spaces are always ultrametrizable (Theorem 10.4.1), answering an open question by Polat.

Finally, Theorem 10.1 brings new progress to an old problem of Diestel, which asks for a topological characterisation of all end spaces [28, Problem 5.1]. Indeed, note that Theorem 10.1 translates to the topological assertion that every open cover of an end space can be refined to an open partition cover, Corollary 10.3.2. This last property is known in the literature as ultra-paracompactness. It implies that all spaces $|G|$ are paracompact (Corollary 10.3.3), and that all end spaces $\Omega(G)$ are even hereditarily ultra-paracompact (Corollary 10.5.3).

This chapter is organised as follows: The next section contains a recap on end spaces and other technical terms. Section 10.3 contains the proof of our main result, and Section 10.4 derives the consequences outlined above. Section 10.5 indicates a simple argument showing that subspaces of end spaces inherit their property of being ultra-paracompact.

10.2. End spaces of graphs: a reminder

The collection of sets $\Omega(X, C)$ with $X \subseteq V$ finite and C a component of $G - X$ form a basis for a topology on Ω . This topology is Hausdorff, and it is *zero-dimensional* in that it has a basis consisting of closed-and-open sets. Note that when considering end spaces $\Omega(G)$, we may always assume that G is connected; adding one new vertex and choosing a neighbour for it in each component does not affect the end space.

We now describe two common ways to extend this topology on $\Omega(G)$ to a topology on $|G| = G \cup \Omega(G)$, the graph G together with its ends. The first topology, called TOP, has a basis formed by all open sets of G considered as a 1-complex, together with basic open neighbourhoods for ends of the form

$$\hat{C}_*(X, \omega) := C(X, \omega) \cup \Omega(X, \omega) \cup \mathring{E}_*(X, C(X, \omega)),$$

where $\mathring{E}_*(X, C(X, \omega))$ denotes any union of half-open edges from the edge cut $E(X, C(X, \omega))$ with endpoint in $C(X, \omega)$.

As the 1-complex topology on G is not first-countable at vertices of infinite degree, it is sometimes useful to consider a metric topology on G instead: The second topology commonly considered, called MTOP, has a basis formed by all open sets of G considered as a metric length-space (i.e. every edge together with

10. Approximating infinite graphs by normal trees

its endvertices forms a unit interval of length 1, and the distance between two points of the graph is the length of a shortest arc in G between them), together with basic open neighbourhoods for ends of the form

$$\hat{C}_\varepsilon(X, \omega) := C(X, \omega) \cup \Omega(X, \omega) \cup \mathring{E}_\varepsilon(X, C(X, \omega)),$$

where $\mathring{E}_\varepsilon(X, C(X, \omega))$ denotes the open ball around $C(X, \omega)$ in G of radius ε . Note that both topologies TOP and MTOP induce the same subspace topology on $\hat{V}(G) := V(G) \cup \Omega(G)$ and $\Omega(G)$, the last of which coincides with the topology on $\Omega(G)$ described above. Polat observed that $\hat{V}(G)$ is homeomorphic with $\Omega(G^+)$, where G^+ denotes the graph obtained from G by gluing a new ray R_v onto each vertex v of G so that R_v meets G precisely in its first vertex v and R_v is distinct from all other $R_{v'}$, cf. [68, §4.16].

The *tree-order* of a rooted tree (T, r) is defined by setting $u \leq v$ if u lies on the unique path rTv from r to v in T . Given $n \in \mathbb{N}$, the n th level of T is the set of vertices at distance n from r in T . The *down-closure* of a vertex v is the set $[v] := \{u : u \leq v\}$; its *up-closure* is the set $[v] := \{w : v \leq w\}$. The down-closure of v is always a finite chain, the vertex set of the path rTv . A ray $R \subseteq T$ starting at the root is called a *normal ray* of T .

A rooted tree T contained in a graph G is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T . Here, for a given a graph H , a path P is said to be an *H-path* if P is non-trivial and meets H exactly in its endvertices. We remark that for a normal tree $T \subseteq G$ the neighbourhood $N(D)$ of every component D of $G - T$ forms a chain in T . A set U of vertices is *dispersed* in G if for every end ω there is a finite $X \subseteq V$ with $C(X, \omega) \cap U = \emptyset$, or equivalently, if U is a closed subset of $|G|$ (in either TOP or MTOP).

Theorem 10.2.1 (Jung [52], Chapter 5). *A vertex set in a connected graph is dispersed if and only if there is a rayless normal tree including it. Moreover, a connected graph has a normal spanning tree if and only if its vertex set is a countable union of dispersed sets.*

If H is a subgraph of G , then rays equivalent in H remain equivalent in G ; in other words, every end of H can be interpreted as a subset of an end of G , so the natural inclusion map $\iota : \Omega(H) \rightarrow \Omega(G)$ is well-defined. A subgraph $H \subseteq G$ is *end-faithful* if this inclusion map ι is a bijection. The terms *end-injective* and *end-surjective* are defined accordingly. Normal trees are always end-injective; hence, normal trees are end-faithful as soon as they are end-surjective. Given a subgraph $H \subseteq G$, write $\partial_\Omega H \subseteq \Omega(G)$ for the set of ends ω of G which satisfy $C(X, \omega) \cap H \neq \emptyset$ for all finite $X \subseteq V(G)$.

For topological notions we follow the terminology in [36]. All spaces considered in this chapter are Hausdorff, i.e. every two distinct points have disjoint open neighbourhoods. An *ultrametric* space (X, d) is a metric space in which the triangle inequality is strengthened to $d(x, z) \leq \max\{d(x, y), d(y, z)\}$. A topological space X is *ultrametrizable* if there is an ultrametric d on X which induces the topology of X . A topological space is *normal* if for any two disjoint closed sets A_1, A_2 there are disjoint open sets U_1, U_2 with $A_i \subseteq U_i$. A space is *collectionwise normal* if

10. Approximating infinite graphs by normal trees

for every *discrete* family $\{A_s: s \in S\}$ of disjoint closed sets, i.e. a family such that $\bigcup\{A_s: s \in S'\}$ is closed for any $S' \subseteq S$, there is a collection $\{U_s: s \in S\}$ of disjoint open sets with $A_s \subseteq U_s$.

A collection of sets \mathcal{A} is said to *refine* another collection of sets \mathcal{B} if for every $A \in \mathcal{A}$ there is $B \in \mathcal{B}$ with $A \subseteq B$. A cover \mathcal{V} of a topological space X is *locally finite* if every point of X has an open neighbourhood which meets only finitely many elements of \mathcal{V} . A topological space X is *paracompact* if for every open cover \mathcal{U} of X there is a locally finite open cover \mathcal{V} refining \mathcal{U} . All compact Hausdorff spaces and also all metric spaces are paracompact, which in turn are always normal and even collectionwise normal [36, Chapter 5.1]. A space is *ultra-paracompact* if every open cover has a refinement by an open partition.

Lastly, ordinal numbers are identified with the set of all smaller ordinals, i.e. $\alpha = \{\beta: \beta < \alpha\}$ for all ordinals α .

10.3. Proof of the main result

This section is devoted to the proof of our main theorem, which we restate more formally:

Theorem 10.1. *For every collection $\mathcal{C} = \{C(X_\omega, \omega): \omega \in \Omega(G)\}$ in a connected graph G , there is a rayless normal tree T in G such that every component of $G - T$ is included in an element of \mathcal{C} .*

As every rayless normal tree $T \subseteq G$ is dispersed in G by Jung's Theorem 10.2.1, this technical variant of our main result is clearly equivalent to the formulation presented in the introduction.

Let us briefly discuss two other possible notions of 'approximating graphs by normal trees': First, Theorem 10.1 is significantly stronger than just requiring that (every component of) $G - T$ is included in the union $\bigcup \mathcal{C}$ of the selected neighbourhoods; the latter assertion is easily seen to be equivalent to Jung's Theorem 10.2.1. In the other direction, could one strengthen our notion of 'approximating by normal trees' and demand a normal rayless tree T such that for every end ω of G , the component of $G - T$ in which every ray of ω has a tail is included in $C(X_\omega, \omega)$? This notion, however, is too strong and such a T may not exist: Consider the graph $G = K^+$ (see Section 10.2) for an uncountable clique K , and let \mathcal{C} be the collection of all the ray-components of $G - K$ (together with an arbitrary neighbourhood of the end of the clique K). Any normal tree for G satisfying our stronger requirements would restrict to a normal spanning tree of K , an impossibility.

We now turn towards the proof of Theorem 10.1. As a first but crucial step, we prove a result similar to our main theorem, but which is only concerned with the end space of a graph.

Theorem 10.3.1. *For every connected graph G and every open cover \mathcal{U} of its end space $\Omega(G)$ there is a rayless normal tree T in G such that the collection of components of $G - T$ induces an open partition of $\Omega(G)$ refining \mathcal{U} .*

10. Approximating infinite graphs by normal trees

Proof. It suffices to prove the statement for open covers of *basic type*, i.e. open covers \mathcal{U} where each element $U \in \mathcal{U}$ is of the form $U = \Omega(X_U, \mathcal{C}_U)$. The proof proceeds by induction on $|\mathcal{U}| = \kappa$. As the statement clearly holds for finite such covers of basic type, we may assume that κ is infinite and that the assertion holds for any end space $\Omega(G')$ and any open cover of basic type \mathcal{U}' of $\Omega(G')$ with $|\mathcal{U}'| < \kappa$.

Choose an enumeration $\mathcal{U} = \{U_\alpha : \alpha < \kappa\}$ of \mathcal{U} in order type κ , and define a rank function ϱ on $\Omega(G)$ by

$$\varrho: \Omega(G) \rightarrow \kappa, \omega \mapsto \min \{ \alpha : \omega \in U_\alpha \} < \kappa.$$

Call a subset $A \subseteq \Omega(G)$ *bounded* or *unbounded* depending on whether its image $\varrho[A] \subseteq \kappa$ is bounded or unbounded in κ . Similarly, a subgraph $H \subseteq G$ is called *bounded* or *unbounded* if the set of ends in G with a ray in H is bounded or unbounded.

We construct a sequence of rayless normal trees $T_1 \subseteq T_2 \subseteq \dots$ extending each other all with the same root r as follows: Let T_1 be the tree on a single vertex r (for some arbitrarily chosen vertex $r \in G$) and suppose that T_n has already been constructed. For every unbounded component D of $G - T_n$ there exists a finite separator $S_D \subseteq V(D)$ such that $D - S_D$ has either zero or at least two unbounded components: Otherwise, the map d sending each finite vertex set in D to its unique unbounded component is a direction on D and hence defines an end ω of D by Theorem 2.4.1. However, ω has a rank, say $\varrho(\omega) = \alpha$, and since \mathcal{U} consists of open sets, there is a basic open neighbourhood $\Omega(S, \omega) \subseteq U_\alpha$, implying that $d(S \cap D) \subseteq C(S, \omega)$ is bounded, a contradiction. Now for every such unbounded D let S_D be a finite separator of the first kind in D if possible, and otherwise of the second kind. Since G is connected, we may extend T_n simultaneously into every unbounded component D of $G - T_n$ so as to include S_D in an inclusion minimal way preserving normality (using the technique as in [26, Proposition 1.5.6]). Then the extension $T_{n+1} \supseteq T_n$ is a rayless normal tree with root r . This completes the construction.

Now consider the normal tree $T' = \bigcup_{n \in \mathbb{N}} T_n$. We claim that T' is rayless. Indeed, suppose otherwise, that there is a normal ray R in T' belonging to the end $\omega \in \Omega(G)$ say.

Then, for every $n \in \mathbb{N}$, the ray R has a tail in an unbounded component D_n of $G - T_n$, and all finite separators S_{D_n} chosen for these components were of the second kind, since we never extended T_n into a component that was already bounded. In particular R meets each S_{D_n} in at least one vertex, s_n say. Now, fix for every S_{D_n} an unbounded component C_{n+1} of $D_n - S_{D_n}$ different from D_{n+1} . Every C_{n+1} has a neighbour, say u_n , in S_{D_n} . Moreover, the paths $P_n = s_n T u_n$ connecting s_n to u_n in T are pairwise disjoint, as each of them was constructed in the n th step.

From this, we obtain a contradiction as follows. Our end ω has rank $\varrho(\omega) = \alpha$ say. Since κ is infinite and hence a limit ordinal, and since the C_n are all unbounded, we may select for each $n \geq 1$ a ray R_n in C_n belonging to an end ω_n with $\varrho(\omega_n) > \alpha$. We may choose R_n so that its starting vertex sends an edge to u_{n-1} .

10. Approximating infinite graphs by normal trees

However, the union of the ray R with the rays R_n and the paths P_n witnesses that $\omega_n \rightarrow \omega$ in $\Omega(G)$ as $n \rightarrow \infty$. As U_α is an open neighbourhood of ω , we have $\omega_n \in U_\alpha$ eventually, implying in turn that $\varrho(\omega_n) \leq \alpha$, contradicting the choice of ω_n . This shows that ω cannot exist, and hence that T' is rayless.

Next, we claim that every component D of $G - T' = G - \bigcup_{n \in \mathbb{N}} T_n$ is bounded. Since T' is a normal tree, $N(D)$ is a chain in T' , and since T' is rayless, $N(D)$ is finite. Hence, there is $m \in \mathbb{N}$ such that $N(D) \subseteq T_m$, i.e. D is already a component of $G - T_m$. The fact that we have not extended T_m into D means that D is bounded.

In particular, for each component D of $G - T'$, the subcollection

$$\mathcal{U}_D := \{ U_\alpha : \alpha \in \varrho[\Omega(N(D), D)] \} \subseteq \mathcal{U}$$

has size $< \kappa$. Every \mathcal{U}_D restricts to an open cover \mathcal{U}'_D of basic type in $\Omega(D)$ of size $|\mathcal{U}'_D| < \kappa$ as follows. Each $\Omega(X, \mathcal{C}) \in \mathcal{U}_D$ induces an open subset $\Omega(Y, \mathcal{D})$ of $\Omega(D)$ of basic type by letting $Y = X \cap D$ and letting \mathcal{D} be the set of components of $D \cap \bigcup \mathcal{C}$. Indeed, every end of D contained in $\Omega(X, \mathcal{C})$ is also contained in $\Omega(Y, \mathcal{D})$: pick a ray from that end avoiding the finite X ; then the ray lies in $D \cap \bigcup \mathcal{C}$ and, as it is connected, it lies in a component of $D \cap \bigcup \mathcal{C}$.

Hence, by the induction hypothesis, for every component D of $G - T'$ there exists a dispersed set X_D in D for which the components of $D - X_D$ refine the cover \mathcal{U}'_D , which in turn refines \mathcal{U}_D . As T' is normal and rayless, the union of the dispersed sets X_D is dispersed as well. We extend T' to a rayless normal tree T which also includes all these X_D , by Jung's Theorem 10.2.1. Then the collection of the components of $G - T$ induces an open partition of $\Omega(G)$ refining \mathcal{U} as desired. \square

From the observation that $\hat{V}(G) \cong \Omega(G^+)$ (see Section 10.2) we can deduce our main result as a consequence of Theorem 10.3.1.

Proof of Theorem 10.1. Let G be a connected graph. Given a collection $\mathcal{C} = \{ C(X_\omega, \omega) : \omega \in \Omega(G) \}$ in G , we need to find a rayless normal tree T in G such that every component of $G - T$ is included in an element of \mathcal{C} .

Consider the graph G^+ . For $v \in V(G)$ write $\omega_v \in \Omega^+ := \Omega(G^+)$ for the end containing the new ray R_v . The assertion follows by applying Theorem 10.3.1 to the open cover

$$\mathcal{U} = \{ \Omega^+(X_\omega, \omega) : \omega \in \Omega(G) \} \cup \{ \Omega^+(\{v\}, \omega_v) : v \in V(G) \}$$

of the end space Ω^+ of G^+ , and restricting the resulting rayless normal tree T^+ of G^+ to the rayless normal tree $T = T^+ \cap G$ of G . \square

Corollary 10.3.2. *All end spaces $\Omega(G)$ are ultra-paracompact.*

Proof. Since we may assume without loss of generality that G is connected, this follows directly from Theorem 10.3.1. \square

Corollary 10.3.3. *All spaces $|G|$ are paracompact in both TOP and MTOP.*

10. Approximating infinite graphs by normal trees

Proof. First, we consider $|G|$ with MTOP . To show that $|G|$ is paracompact, suppose that any open cover \mathcal{U} of $|G|$ consisting of basic open sets is given. The cover elements come in two types: basic open sets of G , and basic open neighbourhoods of ends. We write $\mathcal{U}_\Omega = \{\hat{C}_{\varepsilon_i}(X_i, \omega_i) : i \in I\}$ for the collection consisting of the latter. As \mathcal{U}_Ω covers the end space of G , applying Theorem 10.1 to the collection $\mathcal{C} := \{C(X_i, \omega_i) : i \in I\}$ yields a rayless normal tree T in G such that $\{C(Y_j, \omega_j) : j \in J\}$, the collection of components of $G - T$ containing a ray, refines \mathcal{C} . For every $j \in J$ we choose $\varepsilon_j := \varepsilon_i$ for some $i \in I$ with $C(Y_j, \omega_j) \subseteq C(X_i, \omega_i)$, ensuring that the disjoint collection $\mathcal{V}_\Omega := \{\hat{C}_{\varepsilon_j}(Y_j, \omega_j) : j \in J\}$ refines \mathcal{U}_Ω .

Next, consider the quotient space H that is obtained from $|G|$ by collapsing every closed subset $C(Y_j, \omega_j) \cup \Omega(Y_j, \omega_j)$ with $j \in J$ to a single point. As the open sets in \mathcal{V}_Ω are disjoint, the quotient is well-defined and we may view H as a rayless multi-graph endowed with MTOP . Now consider the open cover \mathcal{U}_H of H that consists of the quotients of the elements of \mathcal{V}_Ω on the one hand, and on the other hand, for every non-contraction point of H a choice of one basic open neighbourhood in G that is contained in some element of \mathcal{U} . Since metric spaces are paracompact, H admits a locally finite refinement \mathcal{V}_H of \mathcal{U}_H consisting of basic open sets of (H, MTOP) . Then the open cover \mathcal{V} of $|G|$ induced by \mathcal{V}_H gives the desired locally finite refinement of \mathcal{U} .

A similar argument shows that $|G|$ with TOP is paracompact. Here, (H, TOP) is paracompact because all CW-complexes are. \square

Note in particular that paracompactness implies normality and collectionwise normality, and hence we reobtain the previously mentioned results by Polat [68, Lemma 4.14] and Sprüssel [80, Theorems 4.1 & 4.2] as a straightforward consequence of our Corollary 10.3.3.

10.4. Consequences of the approximation result

In [68, Theorem 5.13] Polat characterised the graphs that admit an end-faithful normal tree as the graphs with ultrametrizable end space, and raised the question [69, §10] whether metrizability of the end space is enough to ensure the existence of an end-faithful normal tree. As our first application we show how using Theorem 10.3.1 provides a much simplified proof for Polat's result that simultaneously answers his question about the metrizable case in the affirmative:

Theorem 10.4.1. *For every connected graph G , the following are equivalent:*

- (i) *The end space of G is metrizable,*
- (ii) *the end space of G is ultrametrizable,*
- (iii) *G contains an end-faithful normal tree.*

Proof. The implication (iii) \Rightarrow (ii) is routine, as the end space of any tree is ultrametrizable (see e.g. [50] for a detailed account), and $\Omega(T)$ and $\Omega(G)$ are homeomorphic for every end-faithful normal tree T of G (see e.g. [28, Proposition 5.5]). The implication (ii) \Rightarrow (i) is trivial.

10. Approximating infinite graphs by normal trees

Hence, it remains to prove (i) \Rightarrow (iii). For this, consider the covers \mathcal{U}_n for $n \in \mathbb{N}$ of $\Omega(G)$ given by the open balls with radius $1/n$ around every end; with respect to some fixed metric d inducing the topology of $\Omega(G)$. By applying Theorem 10.3.1 iteratively to the covers $\mathcal{U}_1, \mathcal{U}_2, \dots$, it is straightforward to construct a sequence of rayless normal trees $T_1 \subseteq T_2 \subseteq \dots$ all rooted at the same vertex such that the partition of $\Omega(G)$ given by the components of $G - T_n$ refines \mathcal{U}_n . Observe that any two ends $\omega \neq \eta$ of G are separated by any T_n with $2/n < d(\omega, \eta)$. Consider the normal tree $T' = \bigcup_{n \in \mathbb{N}} T_n$. We claim that each end $\omega \in \Omega(G) \setminus \partial_\Omega T'$ belongs to a component C of $G - T'$ such that $N(C)$ is finite. Otherwise $N(C)$ lies on a unique normal ray R of T belonging to some end $\eta \in \partial_\Omega T'$, but then clearly, none of the T_n would separate ω from η , a contradiction. Hence, $N(C)$ is finite, and since C contains at most one end, T' extends to an end-faithful normal tree of G . \square

From the new implication (i) \Rightarrow (iii) in Theorem 10.4.1 one also obtains a simple proof of Diestel's characterisation from [24] when $|G|$ is metrizable.

Corollary 10.4.2. *For every connected graph G , the following are equivalent:*

- (i) $|G|$ with MTOP is metrizable,
- (ii) the space $\hat{V}(G)$ is metrizable,
- (iii) G has a normal spanning tree.

Proof. The first implication (iii) \Rightarrow (i) is routine, see e.g. [24]. The implication (i) \Rightarrow (ii) is trivial. For (ii) \Rightarrow (iii) apply Theorem 10.4.1 to the space $\Omega(G^+) \cong \hat{V}(G)$, noting that every end-faithful normal tree of G^+ is automatically spanning. \square

To motivate our next applications, suppose that a given graph G admits a normal spanning tree. Let us call such graphs *normally spanned*. If G is normally spanned, then every component of $G - X$ is normally spanned, too, for any finite $X \subseteq V(G)$. Conversely, the question arises whether a graph admits a normal spanning tree as soon as every end ω has some basic neighbourhood $C(X, \omega)$ that is normally spanned. It turns out that the answer is yes:

Theorem 10.4.3. *If every end of a connected graph G has a normally spanned neighbourhood, then G itself is normally spanned.*

Proof. Let $\mathcal{C} = \{C(X_\omega, \omega) : \omega \in \Omega(G)\}$ be a selection of normally spanned neighbourhoods for all ends of G , and apply Theorem 10.1 to \mathcal{C} to find a rayless normal tree T such that the collection of components of $G - T$ refines \mathcal{C} . By Jung's Theorem 10.2.1, each such component C of $G - T$ is the union of countably many dispersed sets, say $V(C) = \bigcup_{n \geq 1} V_n^C$. But then $V_0 = V(T)$ together with all the sets $V_n := \bigcup \{V_n^C : C \text{ a component of } G - T\}$, for $n \geq 1$, witnesses that $V(G)$ is a countable union of dispersed sets. Hence, G has a normal spanning tree by Jung's theorem. \square

There is also a more topological viewpoint of the above result: The assumptions of Theorem 10.4.3 are by Corollary 10.4.2 equivalent to the assertion that $\hat{V}(G)$ is locally metrizable. But locally metrizable paracompact spaces are metrizable,

[36, Exercise 5.4.A]. Hence, applying Corollary 10.4.2 once again to $\hat{V}(G)$ yields the desired normal spanning tree of G .

Continuing along these lines, we now address the question whether the existence of some *local end-faithful normal tree* for every end of a graph already ensures the existence of an end-faithful normal tree of the entire graph. For a graph G and an end ω , we say that ω has a *local end-faithful normal tree* if there is a normal tree T in G such that $\partial_\Omega T$ is a neighbourhood of ω in $\Omega(G)$.

Theorem 10.4.4. *If every end of a connected graph G has a local end-faithful normal tree, then G has an end-faithful normal tree.*

Proof. By Theorem 10.4.1 every end in $\Omega(G)$ has a metrizable neighborhood. But (ultra-)paracompact spaces which are locally metrizable are metrizable, [36, Exercise 5.4.A]. Consequently, we have by Corollary 10.3.2 that $\Omega(G)$ is metrizable. Applying again Theorem 10.4.1 yields the desired end-faithful normal tree of G . \square

10.5. Paracompactness in subspaces of end spaces

We conclude this chapter with an observation concerning the following fundamental problem on the structure of end spaces raised by Diestel in 1992 [28, Problem 5.1]:

Problem 10.5.1. *Which topological spaces can be represented as an end space $\Omega(G)$ for some graph G ?*

In Corollary 10.3.2 we established that end spaces are always ultra-paracompact. In this section we show that also all subspaces of end spaces inherit the property of being ultra-paracompact, i.e. that end spaces are hereditarily ultra-paracompact. This significantly reduces the number of topological candidates for a solution of Problem 10.5.1, and for example shows that certain compact spaces cannot occur as end space, which Corollary 10.3.2 wouldn't do on its own.

It is known that paracompactness and ultra-paracompactness, along with a number of other properties which are not per se hereditary such as normality and collectionwise normality, have the property that they are inherited by *all* subspaces as soon as they are inherited by all *open* subspaces. For the easy proof in case of paracompactness see e.g. Didonné's original paper [33, p. 68]. Hence, our assertion follows at once from Corollary 10.3.2 given the following observation:

Lemma 10.5.2. *Open subsets of end spaces are again end spaces.*

Proof. Let G be any graph, and consider some open, non-empty set $\Gamma \subseteq \Omega(G)$. Write Γ^c for its complement in $\Omega(G)$. Using Zorn's lemma, pick a maximal collection \mathcal{R} of disjoint rays all belonging to ends in Γ^c , and let W be the union $\bigcup \{V(R) : R \in \mathcal{R}\}$ of their vertex sets. We claim that Γ is homeomorphic to the end space of the graph $G' := G - W$.

In order to find a homeomorphism $\varphi: \Omega(G') \rightarrow \Gamma$, note first that, due to the maximality of \mathcal{R} , every ray in G' is (as a ray of G) contained in an end of Γ .

10. Approximating infinite graphs by normal trees

Consequently, every end ω' of G' is contained in a unique end ω of Γ and we define φ via this correspondence.

To see that φ is surjective, consider an open neighbourhood $\Omega(X, \omega) \subseteq \Gamma$, for a given $\omega \in \Gamma$. Then W has only finite intersection with $C(X, \omega)$, as only finitely many rays from \mathcal{R} can intersect $C(X, \omega)$, but do not have a tail in $C(X, \omega)$. So we may assume that $C(X, \omega)$ is contained in G' , by extending X . Now, every ray of ω contained in $C(X, \omega)$ gives an end in G' that is mapped to ω .

To see that φ is injective, suppose there are two rays R_1, R_2 in G' that are not equivalent in G' but equivalent in G . Then, there are infinitely many pairwise disjoint R_1 - R_2 paths in G and all but finitely many of these paths hit W . Then the end ω of G containing R_1 and R_2 is an end in Γ which lies in the closure of Γ^c , contradicting that Γ^c is closed.

Finally, let us show that φ is continuous and open. For the continuity of φ remember that for any open set $\Omega(X, \omega) \subseteq \Gamma$ we may assume that $C(X, \omega)$ is contained in G' . In particular the preimage of $\Omega(X, \omega)$ is open in G' .

For φ being open, consider an open set $\Omega(X, \omega') \subseteq \Omega(G')$. Now, $C(X, \omega') \subseteq G' - X$ might not be a component of $G - X$. However, the set of vertices in $C(X, \omega')$ having a neighbour in W is dispersed. Again by extending X , we may assume that $C(X, \omega')$ is a component of $G - X$. Consequently, its image is open in $\Omega(G)$. □

Corollary 10.5.3. *All end spaces are hereditarily ultra-paracompact.* □

Interestingly, a careful reading of Sprüssel's proof that spaces $|G|$ are normal from [80] establishes that every end space $\Omega(G)$ is in fact *completely normal*, i.e. that subsets with $\overline{A} \cap B = \emptyset = A \cap \overline{B}$ can be separated by disjoint open sets – a property which is equivalent to hereditary normality, see [36, Theorem 2.1.7]. In any case, also this stronger result of hereditary normality is implied by our paracompactness result in Corollary 10.5.3.

11. Uniquely determined ends and graphs

11.1. Introduction

Halin [47] defined the ends of an infinite graph ‘from below’ as equivalence classes of rays in the graph, where two rays are equivalent if no finite set of vertices separates them. As a complementary description of Halin’s ends, Diestel and Kühn [30] introduced the notion of directions of infinite graphs. These are defined ‘from above’: A *direction* of a graph G is a map f , with domain the collection $\mathcal{X} = \mathcal{X}(G)$ of all finite vertex sets of G , that assigns to every $X \in \mathcal{X}$ a component $f(X)$ of $G - X$ such that $f(X) \supseteq f(X')$ whenever $X \subseteq X'$. Every end ω of G defines a direction f_ω of G by letting $f_\omega(X)$ be the component of $G - X$ that contains a subray of every ray in ω . Diestel and Kühn showed that the natural map $\omega \mapsto f_\omega$ is in fact a bijection between the ends of G and its directions. This correspondence is now well known and has become a standard tool in the study of infinite graphs. See [24, 25, 62] and Chapters 3, 5 and 10 for examples.

The domain of the directions of G might be arbitrarily large as its size is equal to the order of G . This contrasts with the fact that every direction of G is induced by a ray of G and rays have countable order. Hence the question arises whether every direction of G is ‘countably determined’ in G also by a countable subset of its choices. A *directional choice* in G is a pair (X, C) of a finite vertex set $X \in \mathcal{X}$ and a component C of $G - X$. We say that a directional choice (X, C) in G *distinguishes* a direction f from another direction h if $f(X) = C$ and $h(X) \neq C$. A direction f of G is *countably determined* in G if there is a countable set of directional choices in G that distinguish f from every other direction of G .

Curiously, the answer to this question is in the negative: Consider the graph G that arises from the uncountable complete graph K^{\aleph_1} by adding a new ray R_v for every vertex $v \in K^{\aleph_1}$ so that R_v meets K^{\aleph_1} precisely in its first vertex v and R_v is disjoint from all the other new rays $R_{v'}$. Then $K^{\aleph_1} \subseteq G$ induces a direction of G that is not countably determined in G .

This example raises the question of which directions of a given graph G are countably determined. In the first half of our chapter we answer this question: we characterise for every graph G , by unavoidable substructures, both the countably determined directions of G and its directions that are not countably determined.

If $R \subseteq G$ is any ray, then every finite initial segment X of R naturally defines a directional choice in G , namely (X, C) for the component C that contains $R - X$. Let us call R *directional* in G if its induced direction is distinguished from every other direction of G by the directional choices that are defined by R . By definition, every direction of G that is induced by a directional ray is countably determined in G . Surprisingly, our characterisation implies that the converse holds as well: if a direction of G is distinguished from every other direction by countably many directional choices (X, C) , then no matter how the vertex sets X lie in G we can always assume that the sets X are the finite initial segments of a directional ray:

11. Uniquely determined ends and graphs

Theorem 11.1. *For every graph G and every direction f of G the following assertions are equivalent:*

- (i) *The direction f is countably determined in G .*
- (ii) *The direction f is induced by a directional ray of G .*

As our second main result we characterise by unavoidable substructures the directions of any given graph that are not countably determined in that graph, and thereby complement our first characterisation. Our theorem is phrased in terms of substructures that are uncountable star-like combinations either of rays or of double rays. 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 . An end of G is *dominated* if one (equivalently: each) of its rays is dominated, see [26]. Given a direction f of G we write ω_f for the unique end ω of G whose rays induce f , i.e., which satisfies $f_\omega = f$.

Theorem 11.2. *For every graph G and every direction f of G the following assertions are equivalent:*

- (i) *The direction f is not countably determined in G .*
- (ii) *The graph G contains either*
 - *uncountably many disjoint pairwise inequivalent rays all of which start at vertices that dominate ω_f , or*
 - *uncountably many disjoint double rays, all having one tail in ω_f and another not in ω_f , so that the latter tails are inequivalent for distinct double rays.*

Note that (ii) clearly implies (i).

Does the local property that every direction of G is countably determined in G imply the stronger global property that there is one countable set of directional choices that distinguish every two directions of G from each other? We answer this question in the negative; see Lemma 11.4.1. Let us call a graph G *countably determined* if there is a countable set of directional choices in G that distinguish every two directions of G from each other.

In the second half of our chapter we structurally characterise both the graphs that are countably determined and the graphs that are not countably determined. 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. [26]. (A T -path in G is a non-trivial path that meets T exactly in its endvertices.)

Theorem 11.3. *For every connected graph G the following assertions are equivalent:*

- (i) *G is countably determined.*
- (ii) *G contains a countable normal tree that contains a ray from every end of G .*

Complementing this characterisation we structurally characterise, as our fourth main result, the graphs that are not countably determined. If G is a graph and (T, \mathcal{V}) is a tree-decomposition of G that has finite adhesion, then every direction of

11. Uniquely determined ends and graphs

G either corresponds to a direction of T or lives in a part of (T, \mathcal{V}) ; see Section 11.2.2. An *uncountable star-decomposition* is a tree-decomposition whose decomposition tree is a star $K_{1,\kappa}$ for some uncountable cardinal κ .

Theorem 11.4. *For every connected graph G the following assertions are equivalent:*

- (i) G is not countably determined.
- (ii) G has an uncountable star-decomposition of finite adhesion such that in every leaf part there lives a direction of G .

Interestingly, countably determined directions and countably determined graphs admit natural topological interpretations. Over the course of the last two decades, the topological properties of end spaces have been extensively investigated, see e.g. [24, 30, 68, 69, 80]. However, not much is known about such fundamental properties as countability axioms. Recall that a topological space is *first countable* at a given point if it has a countable neighbourhood base at that point. A direction of a graph G is countably determined in G if and only if it is defined by an end that has a countable neighbourhood base in the end space of G (Theorem 11.3.2). Thus, Theorems 11.1 and 11.2 characterise combinatorially when the end space of a graph is first countable or not first countable at a given end, respectively. Similarly, a graph is countably determined if and only if its end space is *second countable* in that its entire topology has a countable base (Theorem 11.4.7). Therefore, Theorems 11.3 and 11.4 characterise combinatorially the infinite graphs whose end spaces are second countable or not second countable, respectively.

This chapter is organised as follows: In the next section we give a reminder on end spaces and recall all the results from graph theory and general topology that we need. We prove Theorems 11.1 and 11.2 in Section 11.3 and we prove Theorems 11.3 and 11.4 in Section 11.4.

11.2. Preliminaries

For topological notions we follow the terminology in [36].

11.2.1. Normal trees

The *tree-order* of a rooted tree $T = (T, r)$ is defined by setting $u \leq v$ if u lies on the unique path rTv from r to v in T . Given $n \in \mathbb{N}$, the *n th level* of T is the set of vertices at distance n from r in T . The *down-closure* of a vertex v is the set $[v] := \{u : u \leq v\}$; its *up-closure* is the set $[v] := \{w : v \leq w\}$. The down-closure of v is always a finite chain, the vertex set of the path rTv . A ray $R \subseteq T$ starting at the root is called a *normal ray* of T .

A rooted tree T contained in a graph G is *normal* in G if the endvertices of every T -path in G are comparable in the tree-order of T . Here, for a given a graph H , a path P is said to be an *H -path* if P is non-trivial and meets H exactly in its endvertices. We remark that for a normal tree $T \subseteq G$ the neighbourhood $N(C)$ of every component C of $G - T$ forms a chain in T .

11. Uniquely determined ends and graphs

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 $[x]$. Every graph G reflects the separation properties of each normal tree $T \subseteq G$:

Lemma 11.2.1 (Lemma 5.2.10). *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 $[x] \cap [y]$.*
- (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 $T - W$.*

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 , [26, Lemma 8.2.3]. More generally,

Lemma 11.2.2 (Lemma 5.2.11). *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 .*

Not every connected graph has a normal spanning tree. However, every countable connected graph does; more generally:

Lemma 11.2.3 (Jung [52], Corollary 5.3.3). *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 .*

If H is a subgraph of G , then rays equivalent in H remain equivalent in G ; in other words, every end of H can be interpreted as a subset of an end of G , so the natural inclusion map $\iota: \Omega(H) \rightarrow \Omega(G)$ is well-defined. A subgraph $H \subseteq G$ is *end-faithful* if this inclusion map ι is a bijection. The terms *end-injective* and *end-surjective* are defined accordingly. Normal trees are always end-injective; hence, normal trees are end-faithful as soon as they are end-surjective.

11.2.2. Tree-decompositions, S -trees and ends

We assume familiarity with [26, 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 [26, 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.

Next, 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.

We write $S_{\aleph_0} = S_{\aleph_0}(G)$ for the set of all finite-order separations of G . Let (T, α) be any S_{\aleph_0} -tree. If ω is an end of G , then ω orients every finite-order separation

11. Uniquely determined ends and graphs

$\{A, B\} \in S_{\aleph_0}$ 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 $\vec{F}_t := \{(s, t) \in \vec{E}(T) \mid st \in T\}$ 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}) are contained in O . The direction f_ω *lives* at the node of T or *corresponds* to the end of T that ω lives at or corresponds to, respectively. When (T, α) corresponds to a tree-decomposition (T, \mathcal{V}) and f_ω and ω live at t , then we also say that f_ω and ω *live* in the part V_t at t .

Let G be any graph and let ω_1, ω_2 be any two ends of G . A finite-order separation $\{A, B\}$ of G *distinguishes* ω_1 and ω_2 if it satisfies $C(A \cap B, \omega_1) \subseteq G[A \setminus B]$ and $C(A \cap B, \omega_2) \subseteq G[B \setminus A]$ or vice versa. If $\{A, B\}$ distinguishes ω_1 and ω_2 and has minimal order $|A \cap B|$ among all the separations of G that distinguish ω_1 and ω_2 , then $\{A, B\}$ distinguishes ω_1 and ω_2 *efficiently*. The ends ω_1 and ω_2 are said to be *k-distinguishable* for an integer $k \geq 0$ if there is a separation of G of order at most k that distinguishes ω_1 and ω_2 . An S_{\aleph_0} -tree (T, α) *distinguishes* ω_1 and ω_2 if some separation $\{A, B\} \in S_{\aleph_0}$ that α associates with an edge of T distinguishes ω_1 and ω_2 ; it distinguishes ω_1 and ω_2 *efficiently* if $\{A, B\}$ can be chosen to distinguish ω_1 and ω_2 efficiently.

The following theorem is a consequence of [20, Theorem 6.2] and its proof:

Theorem 11.2.4 ([20]). *Every connected graph G has for every number $k \in \mathbb{N}$ a tree-decomposition that efficiently distinguishes all the k -distinguishable ends of G .*

(This chapter is based on an article [60] in preparation; my co-author and I will ask the authors of [20] to include this theorem in their paper as a corollary.)

11.3. Countably determined directions and the first axiom of countability

In this section we characterise for every graph G , by unavoidable substructures, both the countably determined directions of G and its directions that are not countably determined.

Given a graph G we call a ray $R \subseteq G$ *topological* in G if the end ω of G that contains R has a countable neighbourhood base $\{\Omega(X_n, \omega) : n \in \mathbb{N}\}$ in $\Omega(G)$ where each vertex set X_n consists of the first n vertices of R . I remark that my co-author and I will likely change the name ‘topological’ in the final version of the paper [60] which this chapter is based on. Our first lemma shows that rays are topological if and only if they are directional:

Lemma 11.3.1. *For every graph G and every ray $R \subseteq G$ the following assertions are equivalent:*

- (i) R is directional in G .
- (ii) R is topological in G .

11. Uniquely determined ends and graphs

Proof. Let us write ω for the end of G that is represented by R , and let us denote by X_n the set of the first n vertices of R .

(ii) \Rightarrow (i) By assumption, $\{\Omega(X_n, \omega) : n \in \mathbb{N}\}$ is a countable neighbourhood base for $\omega \in \Omega(G)$. Then f_ω is countably determined by the directional choices $(X_n, f_\omega(X_n))$ because the end space is Hausdorff.

(i) \Rightarrow (ii) We claim that $\{\Omega(X_n, \omega) : n \in \mathbb{N}\}$ is a neighbourhood base for $\omega \in \Omega(G)$. Now, suppose for a contradiction that there is a basic open neighbourhood $\Omega(X, \omega)$ of ω in $\Omega(G)$ that contains none of the sets $\Omega(X_n, \omega)$. We recursively construct a sequence of pairwise disjoint rays R_n all having precisely their first vertex on R and belonging to ends not in $\Omega(X, \omega)$. Having these rays at hand will give the desired contradiction; as X is finite one of these rays lies in $C(X, \omega)$ contradicting that its end is not in $\Omega(X, \omega)$.

So suppose we have found R_0, \dots, R_{n-1} . In order to define R_n , choose k large enough that $(X_k, f_\omega(X_k))$ distinguishes f_ω from all directions induced by the rays R_0, \dots, R_{n-1} (if $n = 0$ pick $k = 0$). Such k exists because R is directional. Since the rays R_0, \dots, R_{n-1} have precisely their first vertex on R and X_k consists of vertices of R , none of the rays R_0, \dots, R_{n-1} meets the component $f_\omega(X_k) = C(X_k, \omega)$. By our assumption there is an end, η say, that is contained in $\Omega(X_k, \omega)$ but not in $\Omega(X, \omega)$. We choose any ray of η in $C(X_k, \omega)$ having precisely its first vertex on R to be the n th ray R_n . \square

Theorem 11.3.2. *For every graph G and every end ω of G the following assertions are equivalent:*

- (i) *The end space of G is first countable at ω .*
- (ii) *The direction f_ω is countably determined in G .*
- (iii) *The end ω is represented by a directional ray.*
- (iv) *The end ω is represented by a topological ray.*

This theorem clearly implies Theorem 11.1:

Proof of Theorem 11.1. Theorem 11.3.2 (ii) \Leftrightarrow (iii) is the statement of Theorem 11.1. \square

Proof. (i) \Rightarrow (ii) Let $\{\Omega(X_n, \omega) : n \in \mathbb{N}\}$ be a countable neighbourhood base of basic open sets for $\omega \in \Omega(G)$. Then f_ω is countably determined by its countably many directional choices $(X_n, f_\omega(X_n))$ because the end space is Hausdorff.

(ii) \Rightarrow (iii) Let $\{\Omega(X_n, f_\omega(X_n)) : n \in \mathbb{N}\}$ be a countable set of directional choices that distinguish f_ω from every other direction. Fix any ray $R \in \omega$ and denote by U the union of $V(R)$ and all X_n . By Lemma 11.2.3 there is a normal tree $T \subseteq G$ that contains U cofinally. As $V(R) \subseteq T$ we have that $\omega \in \partial_\Omega T$. Now, by Lemma 11.2.2, there is a normal ray R_ω in T belonging to ω . We claim that R_ω is directional in G . For this, it suffices to show that for every other end $\eta \neq \omega$ of G there is a finite initial segment of R_ω separating ω from η . By assumption, there is a number $n \in \mathbb{N}$ such that X_n separates ω from η . Let v be any vertex of the ray $R_\omega - [X_n]$ where the down-closure is taken in T . Since T is normal in G , we have $C(\lceil v \rceil, \omega) \subseteq C(X_n, \omega)$ by Lemma 11.2.1. In particular, the initial segment $\lceil v \rceil$ of R_ω separates ω from η .

11. Uniquely determined ends and graphs




Figure 11.3.1.: The black rays form a sun centred at ω

(iii) \Rightarrow (iv) This is Lemma 11.3.1 (i) \Rightarrow (ii).

(iv) \Rightarrow (i) holds by the definition of a topological ray. □

Our second main result, the characterisation by unavoidable substructures of the directions of any given graph that are not countably determined in that graph, needs some preparation.

Definition 11.3.3 (Generalised paths). Let G be any graph. A *generalised path* in G with *endpoints* $\omega_1 \neq \omega_2 \in \Omega(G)$ is an ordered pair $(P, \{\omega_1, \omega_2\})$ where $P \subseteq G$ is one of the following:

- a double ray with one tail in the end ω_1 and another tail in the end ω_2 ;
- a finite path $v_0 \dots v_k$ such that v_0 dominates the end ω_1 and v_k dominates the end ω_2 ;
- a ray in ω_1 whose first vertex dominates ω_2 .

Two generalised paths (P, Ψ) and (P', Ψ') are *vertex-disjoint* if P and P' are disjoint. Two generalised paths (P, Ψ) and (P', Ψ') are *disjoint* if they are vertex-disjoint and $\Psi \cap \Psi' = \emptyset$.

Definition 11.3.4 (Generalised star and sun). Let G be any graph. A *generalised star* in G with *centre* $\omega \in \Omega(G)$ is a collection of pairwise vertex-disjoint generalised paths $\{(P^i, \{\omega, \omega^i\}) : i \in I\}$ such that each end ω^i is distinct from all other ends ω^j with $j \neq i \in I$. Then the ends ω^i with $i \in I$ are the *leaves* of the generalised star.

A generalised star $\{(P^i, \{\omega, \omega^i\}) : i \in I\}$ is *proper* if either every path P^i is a double ray or every path P^i is a ray in ω^i whose first vertex dominates ω . A proper generalised star in G with centre ω is also called a *sun* in G with centre ω .

In Figures 11.3.1 and 11.3.2 we see two examples of suns of size eight centred at an end ω . If we increase their size from eight to \aleph_1 in the obvious way, then the direction f_ω is no longer countably determined.

11. Uniquely determined ends and graphs




Figure 11.3.2.: The black double rays form a sun centred at ω

Theorem 11.3.5. *For every graph G and every end ω of G the following assertions are equivalent:*

- (i) *The end space of G is not first countable at ω .*
- (ii) *There is an uncountable sun in G centred at ω .*

This theorem clearly implies Theorem 11.2:

Proof of Theorem 11.2. Combine Theorem 11.3.5 with Theorem 11.3.2 (i) \Leftrightarrow (ii). □

Proof of Theorem 11.3.5. (ii) \Rightarrow (i) Suppose for a contradiction that (ii) and \neg (i) hold, i.e., let $\{ (P^i, \{\omega, \omega^i\}) : i \in I \}$ be an uncountable sun in G centred at ω , and let $\{ \Omega(X_n, \omega) : n \in \mathbb{N} \}$ be a countable open neighbourhood base for ω in $\Omega(G)$. As all the vertex sets X_n are finite and all P^i are pairwise disjoint, there is an index $i \in I$ such that P^i misses all of the vertex sets X_n . Hence, the leaf ω^i is contained in all of the neighbourhoods $\Omega(X_n, \omega)$ contradicting that $\Omega(G)$ is Hausdorff and that $\{ \Omega(X_n, \omega) : n \in \mathbb{N} \}$ is a neighbourhood base for $\omega \in \Omega(G)$.

(i) \Rightarrow (ii) Suppose that the end space $\Omega(G)$ is not first countable at ω . Our aim is to construct an uncountable sun in G centred at ω . Let Δ be the set of vertices dominating ω . By Zorn's lemma there is an inclusionwise maximal set \mathcal{R} of pairwise disjoint rays all belonging to ω . Denote by $V[\mathcal{R}]$ the union of the vertex sets of the rays in \mathcal{R} , and let $A := \Delta \cup V[\mathcal{R}]$. Then $\partial_\Omega A = \{\omega\}$. By Zorn's lemma there is an inclusionwise maximal set \mathcal{P} of pairwise disjoint rays all starting at A and belonging to ends of G other than ω .

First, note that \mathcal{P} yields the desired sun if \mathcal{P} is uncountable: Since $\partial_\Omega A = \{\omega\}$ only finitely many rays in \mathcal{P} belong to the same end and every ray in \mathcal{P} has a tail avoiding A . Hence, if \mathcal{P} is uncountable, we pass to an uncountable subset

11. Uniquely determined ends and graphs

$\mathcal{P}' \subseteq \mathcal{P}$ such that all rays in \mathcal{P}' belong to pairwise distinct ends and have precisely their vertex in A . If uncountably many rays in \mathcal{P}' start at Δ we are done. So we may assume that uncountably many rays in \mathcal{P}' start at $V[\mathcal{R}]$. As only countably many rays in \mathcal{P}' start at the same ray in \mathcal{R} , we may pass to an uncountable subset $\mathcal{P}'' \subseteq \mathcal{P}'$ such that all rays in \mathcal{P}'' start at distinct rays in \mathcal{R} . Extending every ray in \mathcal{P}'' by a tail of the unique ray in \mathcal{R} it hits yields again the desired uncountable sun.

Therefore, we may assume that \mathcal{P} is countable. In the remainder of this proof, we show that this is impossible: we deduce that then there is a countable neighbourhood base for ω in Ω , contradicting our assumption.

We claim that if \mathcal{P} is finite, then ω is an isolated point in Ω , that is, there is a finite vertex set separating ω from all other ends of G simultaneously. Let $\omega_0, \dots, \omega_n$ be the ends of the rays in \mathcal{P} . As these are only finitely many, there is a finite vertex set $X \subseteq V(G)$ separating ω from all of the ω_i simultaneously. Now, $C(X, \omega)$ contains only finitely many vertices of the rays in \mathcal{P} ; by possibly extending X we may assume that $C(X, \omega)$ contains no vertex from any ray in \mathcal{P} . Then no end of G other than ω lies in $\Omega(X, \omega)$, because any such end has a ray in $C(X, \omega)$ starting at A and avoiding all rays in \mathcal{P} , contradicting the maximality of \mathcal{P} .

Thus, \mathcal{P} must be countably infinite. Then the vertex set $V[\mathcal{P}]$ is countable as well. By Lemma 11.2.3 there is a normal tree $T \subseteq G$ that contains $V[\mathcal{P}]$ cofinally. Moreover, as \mathcal{P} is infinite we have $\omega \in \partial_\Omega T$, and so there is a normal ray $R_\omega \subseteq T$ belonging to ω by Lemma 11.2.2.

We claim that for any end $\eta \neq \omega$ of G there is a finite initial segment of R_ω separating ω from η in G . This suffices to derive the desired contradiction, because then Lemma 11.3.1 shows that the finite initial segments of R_ω define a countable open neighbourhood base for ω in Ω .

First, suppose $\eta \in \partial_\Omega T$. Then η has a normal ray R_η in T by Lemma 11.2.2. As T is normal in G , the initial segment $R_\omega \cap R_\eta$ of R_ω separates ω from η in G .

Second, suppose $\eta \notin \partial_\Omega T$. Then there is a unique component C of $G - T$ that contains a tail of every ray in η . The neighbourhood $N(C)$ of C in T is a chain. If the neighbourhood $N(C)$ of C is not cofinal in R_ω , then any finite initial segment of R_ω containing $N(C) \cap R_\omega$ separates ω from η in G . So suppose that $N(C)$ is cofinal in R_ω and denote by U the set of all the vertices in C having a neighbour in R_ω .

If a vertex $u \in U$ dominates R_ω , then we pick a ray in η that is contained in C and starts at $u \in \Delta \subseteq A$, contradicting the maximality of \mathcal{P} . Hence we may assume that every vertex in U sends only finitely many edges to R_ω ; in particular, U is infinite. Then we choose infinitely many distinct vertices u_0, u_1, \dots in U and $u_n - R_\omega$ edges e_n ($n \in \mathbb{N}$) such that the edge set $\{e_n : n \in \mathbb{N}\}$ is independent. Next, we apply the star-comb lemma in C to $U' := \{u_n : n \in \mathbb{N}\}$. This cannot return a star for the same reason the vertices $u \in U$ cannot dominate R_ω . Thus, we obtain a comb in C attached to U' . Then its spine, R say, is a ray belonging to ω . Thus, by the maximality of \mathcal{R} there is a vertex of $V[\mathcal{R}]$ on R and in particular in C . Consequently, there is a ray in η that is contained in C and starts at $V[\mathcal{R}] \subseteq A$, contradicting the maximality of \mathcal{P} . This completes the proof. \square

11. Uniquely determined ends and graphs

For the interested reader we remark that, even though end spaces of graphs are in general not first countable, it is straightforward to show that every end space is strong Fréchet–Urysohn (which is a generalisation of the first axiom of countability): A topological space X is called a *strong Fréchet–Urysohn space* if for any sequence of subsets A_0, A_1, \dots of X and every $x \in \bigcap_{n \in \mathbb{N}} \overline{A_n}$ there is a sequence of points x_0, x_1, \dots converging to x such that $x_n \in A_n$ for all $n \in \mathbb{N}$.

Lemma 11.3.6. *End spaces of graphs are strong Fréchet–Urysohn.* □

11.4. Countably determined graphs and the second axiom of countability

In this section we structurally characterise both the graphs that are countably determined and the graphs that are not countably determined.

Clearly, a graph G is countably determined if and only if every component of G is countably determined and only countably many components of G have directions. Similarly, the end space of a graph G is second countable if and only if every component of G has a second countable end space and only countably many components of G have ends. Thus, to structurally characterise the countable determined graphs and the graphs that are not countably determined, and to link this to whether or not the end space is second countable, it suffices to consider only connected graphs.

The local property that every direction of G is countably determined in G does not imply the stronger global property that G is countably determined, and curiously we find a counterexample with a compact end space:

Lemma 11.4.1. *There exists a connected graph G all whose directions are countably determined in it but which is itself not countably determined. The graph G can be chosen such that its end space is compact and first countable at every end, but neither metrisable nor second countable nor separable.*

Recall that a topological space is called *separable* if it admits a countable dense subset. Every second countable space is separable, but the converse is generally false. For end spaces, however, we shall see in Theorem 11.4.2 that the converse is true: the end space of any graph is second countable if and only if it is separable.

Proof of Lemma 11.4.1. Let T be the rooted infinite binary tree. The graph G arises from T by disjointly adding a new ray R' for every rooted ray $R \subseteq T$ such that R'_1 and R'_2 are disjoint for distinct rooted rays $R_1, R_2 \subseteq T$, and joining the first vertex $v_{R'}$ of each R' to all the vertices of R . Then for every rooted ray $R \subseteq T$ the two rays R' and $v_{R'}R$ are directional in G . Since every direction of G is induced by precisely one of these directional rays, all the directions of G are countably determined.

The graph G , however, is not countably determined: If $\{(X_n, C_n) : n \in \mathbb{N}\}$ is any countable collection of directional choices in G , then there is a rooted ray $R \subseteq T$ such that R' avoids all X_n (because T contains continuum many rooted rays

11. Uniquely determined ends and graphs

and $\bigcup_n X_n$ is countable). But then, for all $n \in \mathbb{N}$, the subgraph $G - X_n$ contains a double ray formed by R' and a subray of R avoiding X_n (that is connected to R' by one of the infinitely many $v_{R'}-R$ edges). These double rays then witness that none of the directional choices (X_n, C_n) distinguishes the direction induced by R from the direction induced by R' or vice versa.

The end space of G is first countable at every end because every direction of G is countably determined (Theorem 11.3.2). It is compact because the deletion of any finite set of vertices of G leaves only finitely many components, cf. [24, Theorem 4.1] or Lemma 5.2.8. However, the end space of G is not separable, because every dense subset of $\Omega(G)$ must contain all the continuum many ends represented by the rays R' . Thus, it is neither second countable nor metrisable. \square

When we showed this lemma and its graph construction to Max Pitz, he said that it reminds him of the [Alexandroff double circle](#).

Now we structurally characterise the countably determined graphs and the graphs that are not countably determined, and structurally characterise the graphs whose end spaces are second countable or not. Our introduction suggests that this is the order in which we prove these results, but we will prove them in a different order: First, we shall structurally characterise the graphs whose end spaces are second countable or not. Then, we shall prove that the end space of a graph is second countable if and only if the graph is countably determined. Finally, we shall use this equivalence to immediately obtain structural characterisations of the countably determined graphs and the graphs that are not countably determined. Here, then is our structural characterisation of the graphs whose end spaces are second countable:

Theorem 11.4.2. *For every connected graph G the following assertions are equivalent:*

- (i) *The end space of G is second countable.*
- (ii) *The end space of G is separable.*
- (iii) *There is a countable end-faithful normal tree $T \subseteq G$.*
- (iv) *The end space of G has a countable base that consists of basic open sets.*

Proof. (i) \Rightarrow (ii) Every second countable space is separable.

(ii) \Rightarrow (iii) Let $\Psi \subseteq \Omega$ be any countable and dense subset. Pick a ray $R_\omega \in \omega$ for every end $\omega \in \Psi$ and let $U := \bigcup \{V(R_\omega) : \omega \in \Psi\}$. By Lemma 11.2.3 there is a countable normal tree $T \subseteq G$ that contains U cofinally. We have to show that T is end-faithful. As mentioned in Section 11.2, normal trees are always end-injective. To show that T is end-surjective note that $\partial_\Omega T$ is closed in Ω . Hence we have $\Omega = \overline{\Psi} \subseteq \overline{\partial_\Omega T} = \partial_\Omega T$. So by Lemma 11.2.2 the normal tree T contains a normal ray of every end of G .

(iii) \Rightarrow (iv) Let $T \subseteq G$ be any countable end-faithful normal tree. We claim that the collection $\mathcal{B} := \{\Omega([t], \omega) : t \in T, \omega \in \Omega\}$ is a countable base of the topology on Ω . Note first that \mathcal{B} is indeed countable: Since T is end-faithful and countable, the deletion of finitely many vertices of T from G results in only countable many components containing an end. Consequently, for every $t \in T$ there are only countable many distinct sets of the form $\Omega([t], \omega)$.

11. Uniquely determined ends and graphs

Now, given a basic open set of Ω , say $\Omega(X, \omega)$, our goal is to find a vertex $t \in T$ such that $\Omega([t], \omega) \subseteq \Omega(X, \omega)$. By Lemma 11.2.2, every end η of G in the closure of T contains a normal ray $R_\eta \subseteq T$. By the normality of T and Lemma 11.2.1, every end $\eta \neq \omega$ of G is separated from ω in G by the finite initial segment $R_\eta \cap R_\omega$ of R_ω . In particular, R_ω is directional in G . Hence, by the implication (i) \Rightarrow (ii) of Lemma 11.3.1 the ray R_ω is topological. Thus, there is a vertex $t \in R_\omega$ such that $\Omega([t], \omega) \subseteq \Omega(X, \omega)$ holds.

(iv) \Rightarrow (i) is clear. □

Next, we structurally characterise the graphs whose end spaces are not second countable. The characterising structure is not the star-decomposition in Theorem 11.4 that one would expect; that is because this result is an auxiliary result that we will use in a second step to prove a second structural characterisation, Theorem 11.4.4, which then is phrased in terms of the desired star-decomposition.

Theorem 11.4.3. *For every connected graph G the following assertions are equivalent:*

- (i) *The end space of G is not second countable.*
- (ii) *The graph G contains either*
 - *an uncountable sun,*
 - *uncountably many disjoint generalised paths, or*
 - *a finite vertex set that separates uncountably many ends of G simultaneously.*

Proof. Recall that, by Theorem 11.4.2, the end space of G is second countable if and only if it has a countable base that consists of basic open sets. Then clearly (ii) \Rightarrow (i).

(i) \Rightarrow (ii) For this, suppose that G is given such that the end space of G is not second countable. We have to find one of the three substructures for G listed in (ii). By Zorn's lemma we find an inclusionwise maximal collection \mathcal{P} of pairwise vertex-disjoint generalised paths in G . Our proof consists of two halves. In the first half we show that if \mathcal{P} is uncountable, then we find either an uncountable sun in G or uncountably many disjoint generalised paths in G . In the second half we show that if \mathcal{P} is countable, then we find a finite vertex set of G that separates uncountably many ends of G simultaneously.

First, we assume that \mathcal{P} is uncountable. In this case, we consider the auxiliary multigraph that is defined on the set of ends of G by declaring every generalised path $(P, \{\omega_1, \omega_2\}) \in \mathcal{P}$ to be an edge between ω_1 and ω_2 . Note that the auxiliary multigraph contains only finitely many parallel edges between any two vertices. Thus, by replacing \mathcal{P} with a suitable uncountable subset we may assume that the auxiliary multigraph is in fact a graph.

If that auxiliary graph has a vertex ω of uncountable degree, then its incident edges correspond to uncountably many generalised paths that form an uncountable generalised star in G with centre ω . This uncountable generalised star need not be proper in general. However, it shows that ω has no countable neighbourhood base in Ω , so Theorem 11.3.5 yields an uncountable sun in G with centre ω .

11. Uniquely determined ends and graphs

Otherwise, every vertex of the auxiliary graph has countable degree. Then we greedily find an uncountable independent edge set, and this edge set corresponds to an uncountable collection of disjoint generalised paths in G .

Second, we assume that \mathcal{P} is countable. Then our goal is to find a finite vertex set $X \subseteq V(G)$ that separates uncountably many ends of G simultaneously. By Lemma 11.2.3 there is a countable normal tree $T \subseteq G$ that cofinally contains the union of the vertex sets of the generalised paths in \mathcal{P} . Then $\Omega \setminus \partial_\Omega T$ is uncountable, since otherwise applying Lemma 11.2.3 to the union of $V(T)$ with the vertex set of a ray from every end in $\Omega \setminus \partial_\Omega T$ gives a countable end-faithful normal tree in G , contradicting Theorem 11.4.2.

Every ray from an end in $\Omega \setminus \partial_\Omega T$ has a tail in one of the components of $G - T$ and this component is the same for any two rays in the same end. We say that an end in $\omega \in \Omega \setminus \partial_\Omega T$ *lives* in the unique component of $G - T$ in which every ray in ω has a tail. By the maximality of \mathcal{P} , distinct ends in $\Omega \setminus \partial_\Omega T$ live in distinct components of $G - T$. As $\Omega(G) \setminus \partial_\Omega T$ is uncountable, we conclude that there are uncountably many components of $G - T$ in which an end of $\Omega(G) \setminus \partial_\Omega T$ lives; we call these components *good*.

We claim that every good component of $G - T$ has finite neighbourhood. For this, assume for a contradiction that there is a good component C of $G - T$ whose neighbourhood $N(C) \subseteq T$ is infinite. We write ω for the end in $\Omega \setminus \partial_\Omega T$ that lives in C . The down-closure of $N(C)$ in T forms a ray and we denote by η the end in $\partial_\Omega T$ represented by this ray. Consider the set U of all the vertices in C sending an edge to T . If some vertex $u \in U$ sends infinitely many edges to T , then u dominates η ; in particular, there is a generalised path $(P, \{\omega, \eta\})$ in G where P is a ray in ω that is contained in C and starts at u , contradicting the maximality of \mathcal{P} . Therefore, we may assume that every vertex in U sends only finitely many edges to T ; in particular, U is infinite. Thus, we find an independent set M of infinitely many U - T edges in G ; we denote by U' the set of the endvertices that these edges have in U .

Applying the star-comb lemma in C to U' gives either a star attached to U' or a comb attached to U' . The centre of a star attached to U' would dominate η , yielding the same contradiction that would be caused by a vertex in U sending infinitely many edges to T . Hence we obtain a comb attached to U' . The comb's spine represents η , because of the edges in M . Consequently, there is a double ray $D \subseteq C$ defining a generalised path $(D, \{\omega, \eta\})$ vertex-disjoint from all generalised paths in \mathcal{P} , contradicting the maximality of \mathcal{P} . This completes the proof of the claim that every good component of $G - T$ has finite neighbourhood.

Finally, as all of the uncountably many good components of $G - T$ have a finite neighbourhood in T and T is countable, there are uncountably many such components having the same finite neighbourhood $X \subseteq V(T)$. Then X is a finite vertex set of G that separates uncountably many ends of G simultaneously, as desired. \square

Next, we will prove the structural characterisation of the graphs whose end spaces are not second countable, in terms of the desired star-decomposition:

11. Uniquely determined ends and graphs

Theorem 11.4.4. *For every connected graph G the following assertions are equivalent:*

- (i) *The end space of G is not second countable.*
- (ii) *G has an uncountable star-decomposition of finite adhesion such that in every leaf part there lives an end of G .*

Our proof requires some preparation. Oriented separations of the form $(C, N(C))$ with $C = C(X, \omega)$ for some finite vertex set $X \in \mathcal{X}$ and an end ω of G are called *golden*. A star σ of finite-order separations is *golden* if every separation in σ is golden.

Lemma 11.4.5. *Let G be any connected graph. If there is an uncountable sun in G , then G admits an uncountable golden star of separations.*

This lemma will be superseded by Lemma 11.4.6, but the cost for that will be a huge increase in proof complexity that is hidden in the proof of the advanced Theorem 11.2.4 which we will use. That is why we included Lemma 11.4.5 above and its proof below nonetheless: to offer a glimpse of intuition which the proof of Lemma 11.4.6 cannot offer.

Proof of Lemma 11.4.5. Suppose that we are given an uncountable sun $S = \{(P^i, \{\omega, \omega^i\}) : i \in I\}$ in G with centre ω . An oriented finite-order separation (A, B) of G is *S -separating* if (A, B) separates the centre of S from some leaf ω^i of S in that $C(A \cap B, \omega) \subseteq G[B]$ while $C(A \cap B, \omega^i) \subseteq G[A]$.

Consider the set Σ of all the golden stars that are formed by S -separating separations of G , partially ordered by inclusion, and apply Zorn's lemma to (Σ, \subseteq) to obtain a maximal element $\sigma \in \Sigma$. If σ is uncountable, we are done. We claim that σ must be uncountable, and assume for a contradiction that σ is countable. Let us write U for the union of the separators of the separations in σ . As U is countable, some path P^j avoids U . We consider the two cases that the end ω^j lies in the closure of U or not.

First, suppose that the end ω^j does not lie in the closure of U . It is straightforward to find an S -separating golden separation (C, X) with ω^j living in C and C avoiding U . Note that (C, X) is not contained in σ . We claim that $\sigma' := \sigma \cup \{(C, X)\}$ is again a star contained in Σ . Since all the elements of σ' are S -separating and golden, it remains to show that the separations in σ' form indeed a star. As $\sigma \subseteq \sigma'$ already is a star, it suffices to show $(C, X) \leq (Y, D)$ for all separations $(D, Y) \in \sigma$. For this, let any separation $(D, Y) \in \sigma$ be given. To establish $(C, X) \leq (Y, D)$ it suffices to show that C avoids $Y \cup D$, because X is the neighbourhood of C and Y is the neighbourhood of D . The component C avoids Y because it avoids U which contains Y as a subset. Therefore, the component C is contained in some component of $G - Y$. Now suppose for a contradiction that C and D meet. Then C must be contained in D . Since P^j avoids Y and contains a ray that lies C , we deduce $P^j \subseteq D$. But then ω must live in D , contradicting that (D, Y) is S -separating. Thus, σ' is again an element of Σ , contradicting the maximal choice of σ .

11. Uniquely determined ends and graphs

Second, suppose that the end ω^j lies in the closure of U . We show that this implies $\omega^j = \omega$, a contradiction. For this, let any finite vertex set $X \subseteq V(G)$ be given; we show $C(X, \omega^j) = C(X, \omega)$. To get started, we observe that all but finitely many of the paths P^i avoid X . Also, all but finitely many of the separations $(D, Y) \in \sigma$ have their component D avoid X (the components are disjoint for distinct separations in σ because σ is a star of separations). As U meets $C(X, \omega^j)$ infinitely, this allows us to find a separation $(C(Y, \omega^i), Y) \in \sigma$ that has its separator Y meet the component $C(X, \omega^j)$ while both $C(Y, \omega^i)$ and P^i avoid the finite separator X . To show $C(X, \omega^j) = C(X, \omega)$, it suffices to find a P^i – P^j path in G that avoids X . We find such a path as follows: In $C(Y, \omega^i)$ we find a path from P^i to a vertex that sends an edge to a vertex v in the non-empty intersection $Y \cap C(X, \omega^j)$. And in $C(X, \omega^j)$ we find a path from P^j to v . Then both paths avoid X , and their union contains the desired P^i – P^j path avoiding X . \square

Lemma 11.4.6. *Let G be any connected graph. If there exist uncountably many pairwise vertex-disjoint generalised paths in G , then G admits an uncountable golden star of separations.*

Proof. Let $\{(P^i, \{\omega_1^i, \omega_2^i\}) : i \in I\}$ be any uncountable collection of pairwise vertex-disjoint generalised paths in G . By the pigeonhole principle there exists a number $k \in \mathbb{N}$ and an uncountable subset $J \subseteq I$ such that for all $j \in J$ the ends ω_1^j and ω_2^j are k -distinguishable. Without loss of generality $J = I$. By Theorem 11.2.4 we find a tree-decomposition (T, \mathcal{V}) of G that efficiently distinguishes all the k -distinguishable ends of G . We write (T, α) for the S_{\aleph_0} -tree that corresponds to the tree-decomposition (T, \mathcal{V}) .

Fix an arbitrary root $r \in T$ and write F for the collection of all the edges $e \in T$ whose induced separation $\alpha(e)$ distinguishes two ends ω_1^i and ω_2^i . Then let $T' \subseteq T$ be the subtree that is induced by the down-closure of the endvertices of the edges in F in the rooted tree T , and put $\alpha' := \alpha \upharpoonright \vec{E}(T')$. Note that (T', α') is again an S_{\aleph_0} -tree. If T' has a vertex t of uncountable degree, then we are done. We claim that T' must have a vertex of uncountable degree. Otherwise, T' is countable. Then the union U of the separators of (T', α') is a countable vertex set. In order to obtain a contradiction note that every P^i meets U in at least one vertex and these vertices are distinct for distinct P^i . \square

Proof of Theorem 11.4.4. Recall that, by Theorem 11.4.2, the end space of G is second countable if and only if it has a countable base that consists of basic open sets. Then clearly (ii) \Rightarrow (i).

(i) \Rightarrow (ii) Suppose that the end space of G is not second countable. We are done if there is a finite vertex set separating uncountably many ends of G simultaneously. Therefore, we may assume by Theorem 11.4.3 that either there is an uncountable sun in G or G contains uncountably many disjoint generalised paths. In the first case we are done by Lemma 11.4.5 and in the second case we are done by Lemma 11.4.6. \square

Theorems 11.4.2 and 11.4.4 structurally characterise the graphs whose end spaces are second countable or not, by the structures in terms of which Theorems 11.3

11. Uniquely determined ends and graphs

and 11.4 are phrased. The next theorem allows us to deduce Theorems 11.3 and 11.4 from Theorems 11.4.2 and 11.4.4 immediately.

Theorem 11.4.7. *The end space of a graph is second countable if and only if the graph is countably determined.*

Proof. Let G be any graph. For the forward implication, suppose that the end space of G is second countable. Then, by Theorem 11.4.2, the end space of G has a countable base $\{\Omega(X_n, C_n) : n \in \mathbb{N}\}$ that consists of basic open sets $\Omega(X_n, C_n) \subseteq \Omega(G)$. We claim that the directional choices (X_n, C_n) distinguish every two directions of G from each other. For this, let any two distinct directions f, h of G be given. Since the end space of G is Hausdorff, there is a number $n \in \mathbb{N}$ such that $\Omega(X_n, C_n)$ contains ω_f but not ω_h ; in particular, $f(X_n) = C_n$ and $h(X_n) \neq C_n$ as desired.

For the backward implication suppose that G is countably determined, and let $\{(X_n, C_n) : n \in \mathbb{N}\}$ be any countable set of directional choices (X_n, C_n) in G that distinguish every two directions of G from each other. Let us assume for a contradiction that the end space of G is not second countable. Then, by Theorem 11.4.3, the graph G contains either

- an uncountable sun,
- uncountably many disjoint generalised paths, or
- a finite vertex set that separates uncountably many ends of G simultaneously.

If G contains an uncountable sun or uncountably many disjoint generalised paths, then in either case G contains a generalised path $(P, \{\omega_1, \omega_2\})$ such that P avoids all the countably many finite vertex sets X_n . But then no directional choice (X_n, C_n) distinguishes f_{ω_1} from f_{ω_2} or vice versa, a contradiction. Thus, there must be a finite vertex set $X \subseteq V(G)$ that separates uncountably many ends ω_i ($i \in I$) of G simultaneously. We abbreviate $C(X, \omega_i)$ as D_i . By the pigeonhole principle we may assume that $X = N(D_i)$ for all $i \in I$.

Let us consider the subset $N \subseteq \mathbb{N}$ of all indices $n \in \mathbb{N}$ whose directional choice (X_n, C_n) distinguishes some f_{ω_i} from some f_{ω_j} . Every component C_n with $n \in N$ meets some component D_i because there lives some end ω_i in C_n . Then, for all $n \in N$, either C_n is contained in some component D_i entirely, or C_n meets X and contains all of the components D_i except possibly for the finitely many components D_i that meet X_n . This means that every directional choice (X_n, C_n) with $n \in N$ either distinguishes finitely many directions f_{ω_i} from uncountably many directions f_{ω_j} or vice versa. Thus, for every $n \in N$ there is a cofinite subset $I_n \subseteq I$ such that no f_{ω_i} with $i \in I_n$ is distinguished by (X_n, C_n) from any other f_{ω_j} with $j \in I_n$. But then the uncountably many directions f_{ω_i} with $i \in \bigcap_{n \in N} I_n$ are not distinguished from each other by any directional choices (X_n, C_n) , a contradiction. \square

Proof of Theorem 11.3. Theorem 11.4.2 and Theorem 11.4.7 together imply Theorem 11.3. \square

Proof of Theorem 11.4. Theorem 11.4.4 and Theorem 11.4.7 together imply Theorem 11.4. \square

Part IV.

The Farey graph

12. Every infinitely edge-connected graph contains the Farey graph or $T_{\aleph_0} * t$ as a minor




Figure 12.0.1.: The Farey graph




Figure 12.0.2.: The graph $T_{\aleph_0} * t$

12.1. Introduction

The Farey graph, shown in Figure 12.0.1 and surveyed in [22, 49], plays a role in a number of mathematical fields ranging from group theory and number theory to geometry and dynamics [22]. Curiously, graph theory is not among these. In this chapter we show that the Farey graph plays a central role in graph theory too: it is one of two infinitely edge-connected graphs that must occur as a minor in every infinitely edge-connected graph. Previously it was not known that there was any set of graphs determining infinite edge-connectivity by forming a minor-minimal list in this way, let alone a finite set.

Ramsey theory and the study of connectivity intersect in the problem of finding for any given connectivity k a small set of k -connected subgraphs that occur in every k -connected graph, and thereby characterise k -connectedness. To keep these unavoidable sets small for $k \geq 3$, the subgraph relation referred to above is usually relaxed to the graph minor relation. Here, a graph is a *minor* of a graph G if it can be obtained from a subgraph of G by contracting connected (possibly infinite) induced disjoint subgraphs [26]. We refer to [26, §9.4] or the introduction of [40] for surveys on the known results for this problem and its variations [26, 38, 40, 45, 51, 65]. Such sets of minor-minimal k -connected graphs are known only for $k \leq 4$, and only for finite graphs [65]. These results of Oporowski, Oxley and Thomas were generalised to $k > 4$ by Geelen and Joeris [38] for finite graphs, and by Gollin and Heuer [40] for infinite graphs, but with a different notion of connectivity.

For infinite connectivity, the problem asks for a small selection of infinitely connected graphs such that every infinitely connected graph contains at least one

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

of the selected graphs as a minor. Here, ‘infinitely connected’ can be understood in two ways. When it is understood as ‘infinitely vertex-connected’, the answer is already known: Every infinitely connected graph contains the countably infinite complete graph K^{\aleph_0} as a minor [26, §8.1]. But when ‘infinitely connected’ is understood as ‘infinitely edge-connected’ then, as we shall see, K^{\aleph_0} is not the answer, and in fact no answer has been known. Indeed it is not even clear a priori that there exists a finite set of unavoidable infinitely edge-connected minors. Any such unavoidable infinitely edge-connected minors will be countable, because in every infinitely edge-connected graph we can greedily find a countable infinitely edge-connected subgraph. But the countable graphs are not known to be well-quasi-ordered by the minor-relation. It is therefore not clear that any minor-minimal set of infinitely edge-connected graphs must be finite, nor even that such a minimal set exists.

In this chapter we find a pair of infinitely edge-connected graphs that occur unavoidably as minors in any infinitely edge-connected graph, and which are unique with this property up to minor-equivalence: the Farey graph F , and the graph $T_{\aleph_0} * t$ obtained from the infinitely-branching tree T_{\aleph_0} by joining an additional vertex t to all its vertices (Figure 12.0.2).

Theorem 12.1. *Every infinitely edge-connected graph contains either the Farey graph or $T_{\aleph_0} * t$ as a minor.*

The uniqueness of the pair $\{F, T_{\aleph_0} * t\}$, up to minor-equivalence, follows from the fact that they are not minors of each other (Lemmas 12.3.1 and 12.3.2):

Theorem 12.2. *Let M be any set of infinitely edge-connected graphs such that every infinitely edge-connected graph has a minor in M and no element of M is a minor of another. Then M consists of two graphs, of which one is minor-equivalent to the Farey graph and the other is minor-equivalent to $T_{\aleph_0} * t$.*

Theorem 12.1 is best possible also in the sense that one cannot replace ‘minor’ with ‘topological minor’ in its wording (Theorem 12.3.3).

Since both the Farey graph and $T_{\aleph_0} * t$ are planar, our result implies that every infinitely edge-connected graph contains a planar infinitely edge-connected graph as a minor. Thus, in this sense, infinite edge-connectivity is an inherently planar property.

This chapter is organised as follows. Section 12.2 formally introduces the Farey graph. In Section 12.3 we show that the Farey graph and $T_{\aleph_0} * t$ are not minors of each other, and deduce Theorem 12.2. Theorem 12.3.3 above is proved there as well. We outline the overall strategy of the proof of Theorem 12.1 in Section 12.4. The proof itself consists of two halves. The first half of the proof is carried out in Section 12.5, and the second half is carried out in Section 12.6.

Acknowledgement. I am grateful to Konstantinos Stavropoulos for stimulating conversations.

12.2. Preliminaries

Two graphs are *minor-equivalent* if they are minors of each other. If G is any graph and $X \subseteq V(G)$ is any vertex set, then we denote by $\partial X = \partial_G X$ the subset of X formed by the vertices in X that send an edge in G to $V(G) \setminus X$.

The *Farey graph* F is the graph on $\mathbb{Q} \cup \{\infty\}$ in which two rational numbers a/b and c/d in lowest terms (allowing also $\infty = (\pm 1)/0$) form an edge if and only if $\det \begin{pmatrix} a & c \\ b & d \end{pmatrix} = \pm 1$, cf. [22]. In this chapter we do not distinguish between the Farey graph and the graphs that are isomorphic to it. For our graph-theoretic proofs it will be more convenient to work with the following purely combinatorial definition of the Farey graph that is indicated in [22] and [49].

The *halved Farey graph* \check{F}_0 of order 0 is a K^2 with its sole edge coloured blue. Inductively, the *halved Farey graph* \check{F}_{n+1} of order $n + 1$ is the edge-coloured graph that is obtained from \check{F}_n by adding a new vertex v_e for every blue edge e of \check{F}_n , joining each v_e precisely to the endvertices of e by two blue edges, and colouring all the edges of $\check{F}_n \subseteq \check{F}_{n+1}$ black. The *halved Farey graph* $\check{F} := \bigcup_{n \in \mathbb{N}} \check{F}_n$ is the union of all \check{F}_n without their edge-colourings, and the *Farey graph* is the union $F = G_1 \cup G_2$ of two copies G_1, G_2 of the halved Farey graph such that $G_1 \cap G_2 = \check{F}_0$.

Lemma 12.2.1. *The halved Farey graph and the Farey graph are minor-equivalent.*

Proof. The halved Farey graph is a subgraph of the Farey graph. Conversely, the Farey graph is a minor of the halved Farey graph: if e is a blue edge of \check{F}_1 , then the Farey graph is the contraction minor of $\check{F} - e$ whose sole non-trivial branch set is $V(\check{F}_0)$, i.e., $(\check{F} - e)/V(\check{F}_0) \cong F$. \square

12.3. Uniqueness and topological minors

12.3.1. Uniqueness

In this section we show that the pair $\{F, T_{\aleph_0} * t\}$ is unique up to minor-equivalence:

Theorem 12.2. *Let \mathcal{H} be any set of infinitely edge-connected graphs such that every infinitely edge-connected graph has a minor in \mathcal{H} and no element of \mathcal{H} is a minor of another. Then \mathcal{H} consists of two graphs, of which one is minor-equivalent to the Farey graph and the other is minor-equivalent to $T_{\aleph_0} * t$.*

This will follow easily from the following two lemmas:

Lemma 12.3.1. *The Farey graph is not a minor of $T_{\aleph_0} * t$.*

Proof. The Farey graph contains two disjoint cycles, but $T_{\aleph_0} * t$ does not. \square

Lemma 12.3.2. *The graph $T_{\aleph_0} * t$ is not a minor of the Farey graph.*

Proof of Theorem 12.2. We write $\mathcal{G} = \{F, T_{\aleph_0} * t\}$ and note that neither element of \mathcal{G} is a minor of another by Lemmas 12.3.1 and 12.3.2. Every graph $H \in \mathcal{H}$ contains a graph $G \in \mathcal{G}$ as a minor (Theorem 12.1) which in turn contains a graph

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

$H' \in \mathcal{H}$ as a minor, and then $H \succcurlyeq G \succcurlyeq H'$ implies $H = H'$ because no element of \mathcal{H} is a minor of another. Thus, every graph in \mathcal{H} is minor-equivalent to some graph in \mathcal{G} and, conversely, every graph in \mathcal{G} is minor-equivalent to some graph in \mathcal{H} by symmetry. Since no two graphs in \mathcal{H} or in \mathcal{G} are comparable with regard to the minor-relation, we deduce that minor-equivalence induces a bijection between \mathcal{H} and \mathcal{G} . \square

Showing that $T_{\aleph_0} * t$ is not a minor of the Farey graph requires more effort:

Proof of Lemma 12.3.2. Since K_{2,\aleph_0} is a subgraph of $T_{\aleph_0} * t$, it suffices to show that the Farey graph does not contain K_{2,\aleph_0} as a minor. So let us assume for a contradiction that the Farey graph contains a K_{2,\aleph_0} minor. By applying the star-comb lemma inside the branch sets of the two infinite-degree vertices of K_{2,\aleph_0} if necessary, and using that the Farey graph does not contain infinitely many independent paths between any two of its vertices, we find that our model of K_{2,\aleph_0} contains a subdivision G of one of the following two graphs G_1 and G_2 . The graph G_1 is the ladder with every rung subdivided exactly once, i.e., it is the disjoint union of two rays $R = v_1v_2\dots$ and $R' = v'_1v'_2\dots$ with infinitely many disjoint R - R' paths $v_nz_nv'_n$ ($n \in \mathbb{N}$). And the graph G_2 is obtained from G_1 by contracting R' to a single vertex that we call d .

In either case, the sole end of $G \subseteq F$ is included in a unique end ω of F . The end ω chooses, for every $n \in \mathbb{N}$, a blue edge $e_n \in F_n$ with vertex set X_n for which it lives in the component C_n of $F - X_n$ avoiding F_n . Then C_n has neighbourhood X_n , and so does the other component D_n of $F - X_n$. We remark that, by the construction of the Farey graph, for every vertex of F there is a number n such that the vertex is not contained in C_n . For all n the two vertex sets X_n and X_{n+1} together induce a triangle Δ_n in F . We write x_n for the vertex in which X_n and X_{n+1} meet, and we write Y_n for vertex set consisting of the other two vertices of the triangle Δ_n . The graph $F - \Delta_n$ has precisely three components, namely D_n and C_{n+1} and a third component with neighbourhood Y_n which we denote by H_n .

First, we consider the case that $G \subseteq F$ is a subdivision of G_1 , and we write \hat{R} and \hat{R}' for the subdivisions of R and R' in G . Then there cannot be a number N such that $x_n = x_N$ for all $n \geq N$: Indeed, for every $k \in \mathbb{N}$ there is a number $f(k) \geq k$ such that both v_k and v'_k are not contained in $C_{f(k)}$ and, as a consequence, $x_{f(k)}$ must be contained in $v_k\hat{R} \cup v'_k\hat{R}'$. Thus, every vertex of F lies in a component D_n eventually (and ω is undominated). Let N be the least number for which the first vertices of \hat{R} and \hat{R}' lie in D_N . To derive a contradiction from $G \subseteq F$, let us consider any \hat{R} - \hat{R}' path $P \subseteq G$ that lies entirely in the component C_N , and consider the maximal number n for which P avoids D_n , noting $n \geq N$. Since the two rays \hat{R} and \hat{R}' induce a bipartition of the 2-set X_{n+1} , the path P cannot meet C_{n+1} without contradicting the maximality of n . Therefore, the path P is contained entirely in $F[H_n \sqcup \Delta_n]$. Without loss of generality we have $x_n \in \hat{R}$. Then $Y_n \subseteq \hat{R}'$ follows. But now the non-empty subpath \hat{P} must be contained in H_n , contradicting that H_n has neighbourhood $Y_n \subseteq \hat{R}'$.

Second, we consider the case that $G \subseteq F$ is a subdivision of G_2 , and again we write \hat{R} for the subdivision of R in G . Since $d \in G_2$ dominates the end of G_2 ,

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

the end ω is dominated in F by d . Let N be the least number such that both d and the first vertex of \hat{R} are not contained in C_N . Then $d = x_N = x_n$ for all $n \geq N$ because d dominates ω . Thus, $Y_n \subseteq \hat{R}$ for all $n \geq N$. Now consider any d - \hat{R} path $P \subseteq G$ with $\dot{d}P \subseteq C_N$ and choose n maximal with the property that the non-empty subpath \dot{P} avoids D_n , noting $n \geq N$. Then $\dot{P} \subseteq H_n$ follows because of $Y_n \subseteq \hat{R}$, contradicting that $d = x_n$ does not lie in the neighbourhood Y_n of H_n . \square

12.3.2. Minor versus topological minor

Theorem 12.1 is best possible in the sense that one cannot replace ‘minor’ with ‘topological minor’ in its wording:

Theorem 12.3.3. *There exists an infinitely edge-connected graph that contains neither the Farey graph nor $T_{\aleph_0} * t$ as a topological minor.*

Proof. As we will show in Chapter 14, there exists an infinitely edge-connected graph G which does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices. Here, two u - v paths are *order-compatible* if they traverse their common vertices in the same order. Then the graph G does not contain a subdivision of the Farey graph or of $T_{\aleph_0} * t$ because both the Farey graph and $T_{\aleph_0} * t$ have pairs of vertices with infinitely many edge-disjoint pairwise order-compatible paths between them. \square

12.4. Overall proof strategy

Our aim for the remainder of this chapter is to show that every infinitely edge-connected graph contains either the Farey graph or $T_{\aleph_0} * t$ as a minor (Theorem 12.1). The proof consists of two halves. In the first half (Section 12.5) we show that every infinitely edge-connected graph without a $T_{\aleph_0} * t$ minor is ‘robust’ (Theorem 12.5.13), explained below. Then, in the second half (Section 12.6), we employ Theorem 12.5.13 to prove that every infinitely edge-connected graph without a $T_{\aleph_0} * t$ minor must contain a Farey graph minor, completing the proof of Theorem 12.1.

The Farey graph and $T_{\aleph_0} * t$ are both infinitely edge-connected, but in different ways. The infinite edge-connectivity of the Farey graph, on the one hand, is robust in that deleting the two endvertices of an edge always leaves only infinitely edge-connected components. The infinite edge-connectivity of $T_{\aleph_0} * t$, on the other hand, is fragile in that deleting t results in a tree. In the first half of the proof of Theorem 12.1 we show that every infinitely edge-connected graph without a $T_{\aleph_0} * t$ minor is essentially robust, not fragile (Theorem 12.5.13).

In the second half of the proof of Theorem 12.1 we construct a Farey graph minor in an arbitrary infinitely edge-connected $T_{\aleph_0} * t$ free graph G . By Lemma 12.2.1 it suffices to construct a halved Farey graph minor. Using that G is robust by Theorem 12.5.13, we shall essentially prove the following assertion:

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

For every two vertices u and v of G there exist two induced subgraphs $H_u, H_v \subseteq G$ containing u and v respectively and which satisfy the following conditions:

- (i) $X := V(H_u) \cap V(H_v)$ is finite, non-empty and connected in G ;
- (ii) both H_u/X and H_v/X are infinitely edge-connected;
- (iii) X avoids u and v ;
- (iv) uX is an edge of H_u/X and vX is an edge of H_v/X .

If we choose u and v to form an edge of G , then the three vertices u, v and X span a triangle \tilde{F}_1 in $(H_u \cup H_v)/X$. And since H_u/X and H_v/X are both infinitely edge-connected and robust again, we can reapply the assertion in $(H_u/X) - uX$ to u and X , and in $(H_v/X) - vX$ to v and X . By iterating this process, we obtain a halved Farey graph minor in the original graph G at the limit, and this will complete the proof of Theorem 12.1.

12.5. Robustness

The aim of this section is to prove Theorem 12.5.13 which has been outlined in the previous section. Our proof proceeds in three steps. First, we provide some tools that will help us to (i) identify infinitely edge-connected ‘parts’ of arbitrary graphs and (ii) allow us to distinguish all these ‘parts’ at once in a tree-like way. In the second step, then we employ these tools to analyse the components of $G - u - v$ for infinitely edge-connected graphs G and vertices u, v of G . In the third step, we proceed to prove Theorem 12.5.13.

12.5.1. Finitely separating spanning trees

Let G be any graph. Two vertices of G are said to be *finitely separable* in G if there is a finite set of edges of G separating them in G . If every two distinct vertices of G are finitely separable, then G itself is said to be *finitely separable*. An equivalence relation $\sim = \sim_G$ is declared on the vertex set of G by letting $x \sim y$ whenever x and y are not finitely separable. The graph \tilde{G} is defined on $V(G)/\sim$ by declaring XY an edge whenever $X \neq Y$ and there is an X - Y edge in G . Note that the graph \tilde{G} is always finitely separable. A spanning tree T of G is *finitely separating* if all its fundamental cuts are finite. The following theorem is Theorem 6.3 in [12] and Theorem 7.5 in Chapter 7.

Theorem 12.5.1. *Every connected finitely separable graph has a finitely separating spanning tree.*

Usually, we will employ Theorem 12.5.1 to find a finitely separating spanning tree T of \tilde{G} that we will then use to analyse the overall structure of G with regard to infinite edge-connectivity. In this context, the nodes of $T \subseteq \tilde{G}$ will also be viewed as the vertex sets of G that they formally are. When we view a node of T as a vertex set of G we will refer to it as *part* for clarity.

Every finitely separating spanning tree $T \subseteq \tilde{G}$ defines an S -tree (T, α) for the set $S = \mathcal{B}_{\aleph_0}(G)$ of all the separations of the vertex set $V(G)$ that are bipartitions

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

induced by finite bonds of G : Let the map α send every oriented edge $(t_1, t_2) \in \vec{E}(T)$ to the ordered pair $(\bigcup V(T_1), \bigcup V(T_2))$ for the two components T_1 and T_2 of $T - t_1 t_2$ containing t_1 and t_2 respectively. Then $\alpha(t_1, t_2)$ clearly is an oriented bipartition of $V(G)$. Moreover, we have $\alpha(\vec{e}) \leq \alpha(\vec{f})$ whenever $\vec{e} \leq \vec{f} \in \vec{E}(T)$ and $(\alpha(\vec{e}))^* = \alpha(\vec{e})$ for all $\vec{e} \in \vec{E}(T)$. It remains to show that $\alpha(\vec{e})$ always stems from a finite bond of G . For this, it suffices to show that if $\{A, B\} \in \mathcal{B}_{\aleph_0}(\tilde{G})$ then $\{\bigcup A, \bigcup B\} \in \mathcal{B}_{\aleph_0}(G)$, because all the fundamental cuts of T are finite bonds. Between every two \sim -classes U and W of G there are only finitely many edges, because $u \in U$ is separated from $w \in W$ by a finite cut of G and then U and W must respect this finite cut. Hence the finitely many A - B edges in \tilde{G} give rise to only finitely many $(\bigcup A)$ - $(\bigcup B)$ edges in G , and these are all $(\bigcup A)$ - $(\bigcup B)$ edges in G . Using that G contains for all \sim -equivalent vertices x and y an x - y path avoiding the finitely many $(\bigcup A)$ - $(\bigcup B)$ edges, it is straightforward to show that both $G[\bigcup A]$ and $G[\bigcup B]$ are connected.

The *part* of a star $\{(A_i, B_i) \mid i \in I\}$ of separations of a given set is the intersection $\bigcap_{i \in I} B_i$. If (T, α) is a $\mathcal{B}_{\aleph_0}(G)$ -tree that is defined by a finitely separating spanning tree T of \tilde{G} , then for every node $t \in T$ the part of the star $\alpha[\vec{F}_t] \subseteq \vec{\mathcal{B}}_{\aleph_0}(G)$ associated with t is equal to the part $t \subseteq V(G)$. And the parts $t \subseteq V(G)$ in turn are precisely the \sim -classes of G . Thus, in this sense, by Theorem 12.5.1 every connected graph admits a tree structure that displays all its \sim -classes.

Parts of infinite stars in $\vec{\mathcal{B}}_{\aleph_0}(G)$ can be made connected for a reasonable price:

Lemma 12.5.2. *Suppose that G is a connected graph, that $\sigma = \{(A_i, B_i) \mid i \in I\}$ is an infinite star in $\vec{\mathcal{B}}_{\aleph_0}(G)$ and that $i_* \in I$ is given. Then there is an infinite subset $J \subseteq I$ containing i_* such that the part of the infinite substar $\{(A_j, B_j) \mid j \in J\} \subseteq \sigma$ is connected in G .*

Proof. For each $i \in I$ we write F_i for the finite bond $E(A_i, B_i)$ of G .

Inductively, we construct an ascending sequence $T_0 \subseteq T_1 \subseteq \dots$ of finite trees in G together with a sequence of distinct indices i_0, i_1, \dots in $I \setminus \{i_*\}$ such that, for all $n \in \mathbb{N}$ and $J_n := \{i_*\} \sqcup \{i_0, \dots, i_{n-1}\}$, the tree T_n is a subgraph of $G_n := G[\bigcap_{j \in J_n} B_j]$ containing all ∂B_j with $j \in J_n$. Then the tree $T := \bigcup_{n \in \mathbb{N}} T_n$ will ensure that $G_\infty := G[\bigcap_{j \in J} B_j]$ is connected for $J := \bigcup_{n \in \mathbb{N}} J_n$. (For whenever a path in G connecting two given vertices in G_∞ uses vertices that are not in G_∞ , then the path crosses one of the bonds F_j , and the number of bonds crossed can be decreased by replacing path segments with detours in $T \supseteq \partial B_j$ because $T \subseteq G_\infty$. Therefore, choosing a path that crosses as few bonds F_j as possible will suffice to find a path that lies entirely in G_∞ .)

To start the construction, let T_0 be any finite tree in $G[B_{i_*}]$ that contains ∂B_{i_*} . At step $n + 1$ of the construction, suppose that we have already constructed T_n and J_n . As T_n is finite, we find an index $i_n \in I \setminus J_n$ for which A_{i_n} avoids T_n , ensuring $T_n \subseteq G_{n+1}$. To ensure that T_n can be extended in G_{n+1} to a finite tree T_{n+1} that contains ∂B_{i_n} , it suffices to show that G_{n+1} is connected. Given any two vertices in G_{n+1} , consider any path between them in $G[B_{i_n}]$, chosen to cross as few of the finite bonds F_j with $j \in J_n$ as possible. Then the path avoids all these F_j , for otherwise the number of bonds crossed could be decreased by replacing path segments with detours in $T_n \supseteq \bigcup_{j \in J_n} \partial B_j$. Therefore, G_{n+1} is connected. \square

12.5.2. Analysing the components

Now we analyse the components of $G - u - v$ for infinitely edge-connected graphs G and vertices u, v of G . The main results here are the two Lemmas 12.5.3 and 12.5.8. Here is the first main lemma:

Lemma 12.5.3. *Suppose that G is an infinitely edge-connected graph, that u, v are two distinct vertices of G , and that C is a component of $G - u - v$. If \tilde{C} has a finitely separating spanning tree that contains a subdivision of the infinite binary tree, then $G[C + u + v]$ contains $T_{\aleph_0} * t$ as a minor.*

Proof. Consider any finitely separating spanning tree of \tilde{C} that contains a subdivision of the infinite binary tree. Then this spanning tree also contains T_{\aleph_0} as a contraction minor which gives rise to a $\mathcal{B}_{\aleph_0}(C)$ -tree (T, α) . Next, we fix any root $r \in T$, and for every edge $e \in T$ we fix \vec{e} as its orientation pointing away from the root r (the orientation $\vec{e} = (x, y)$ of $e = \{x, y\}$ satisfying $x \in rTy$). Let $O := \{\vec{e} \mid e \in E(T)\}$. Since G is infinitely edge-connected, O is equal to the union $O_u \cup O_v$ where $\vec{e} \in O_w$ (for $w = u, v$) if and only if w sends an edge in G to B for $\alpha(\vec{e}) = (A, B)$. Now O_u is cofinal¹ in $O \subseteq \vec{E}(T)$ or there is an oriented edge $\vec{e} \in O$ with O_v cofinal in $[\vec{e}]_O := \{\vec{f} \in O \mid \vec{e} \leq \vec{f}\}$. In either case, there is $\vec{e} \in O$ with O_u or O_v cofinal in $[\vec{e}]_O$. Without loss of generality O_u is cofinal in $[\vec{e}]_O$ for some $\vec{e} \in O$. By replacing T with one of its subtrees and restricting α accordingly, we may even assume that O_u is cofinal in O . In fact, then $O_u = O$ follows as O_u is down-closed in O . We will use this to show $T_{\aleph_0} * t \preceq G[C + u]$.

For this, we enumerate the vertices of T_{\aleph_0} as x_0, x_1, \dots such that every x_n is neighbour to some earlier x^k ($k < n$). Inductively, we construct a sequence W_0, W_1, \dots of disjoint connected vertex sets $W_n \subseteq V(C)$, a sequence w_0, w_1, \dots of vertices $w_n \in W_n$, and a sequence t_0, t_1, \dots of distinct nodes $t_n \in T$ such that, for all $n \in \mathbb{N}$:

- (i) $uw_n \in G$;
- (ii) C contains a W_i - W_j edge ($i, j \leq n$) whenever $x_i x_j \in T_{\aleph_0}$;
- (iii) w_n is contained in the part of the star $\alpha[\vec{F}_{t_n}]$;
- (iv) for all $k \leq n$ there are infinitely many oriented edges $\vec{e} \in O \cap (\vec{F}_{t_k})^*$ such that, for $\alpha(\vec{e}) = (B, A)$, the vertex set W_k contains $\partial_C B$ while A is avoided by all W_i with $i \leq n$.

Once the construction is completed, the sets W_n and $\{u\}$ will give rise to a model of $T_{\aleph_0} * t$ in $G[C + u]$ by (i) and (ii).

At the construction start, we choose any neighbour w_0 of u in C (which exists as $O_u = O$ and T is infinite), guaranteeing (i). Then t_0 is defined by (iii). Applying Lemma 12.5.2 in C to the infinite star $\alpha[\vec{F}_{t_0}]$ yields an infinite substar whose connected part $W_0 \subseteq V(C)$ contains w_0 and satisfies both (ii) and (iv) trivially.

At step $n > 0$ of the construction, consider the $k < n$ for which $x_k x_n$ is an edge of T_{\aleph_0} , and pick an edge $\vec{e} \in O \cap (\vec{F}_{t_k})^*$ that (iv) provides for $k \leq n - 1$. If we

¹A 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 $p \geq x$.

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

write $\alpha(\vec{e}) = (B, A)$, then the vertex set W_k contains $\partial_C B$ while A is avoided by all W_i with $i \leq n - 1$. Using $O_u = O$ we find a neighbour w_n of u in A giving (i), and w_n defines t_n by (iii). Then we apply Lemma 12.5.2 in C to the infinite star

$$\{(A_i, B_i) \mid i \in I\} := \alpha[(\vec{F}_{t_n} \setminus O) \cup \{\vec{e}\}]$$

where we take $i_* \in I$ to be the index of the separation $\alpha(\vec{e})$. This yields an infinite substar whose connected part $W_n \subseteq V(C)$ contains w_n and satisfies (ii) because W_n contains $\partial_C A$ while W_k contains $\partial_C B$. Using the infinite substar and the choice of \vec{e} it is straightforward to verify (iv) for all $k \leq n$. \square

Our second main lemma, Lemma 12.5.8, requires some preparation.

Definition 12.5.4 (Arrow). Suppose that u and v are two distinct vertices.

An *arrow from u to v* is a graph G that arises from the two vertices u and v by disjointly adding an infinitely edge-connected graph H , adding a u - H edge uh , and adding infinitely many v - $(H - h)$ edges. Then H is the arrow's *payload*, u is its *nock* and v is its *head*.

An *arrow barrage from u to v* is a countably infinite union $\bigcup_{n \in \mathbb{N}} A_n$ of arrows A_n from u to v such that A_n and A_m do not meet in any vertices other than u and v for all $n \neq m$. Then u and v are the *nock* and *head* of the arrow barrage.

When we say that some graph contains an arrow (barrage) minor *from x to y* for two vertices x and y , we mean that the graph contains an arrow (barrage) minor such that the branch set corresponding to the arrow (barrage)'s nock contains x while the branch set corresponding to the arrow (barrage)'s head contains y .

The next definition captures the concept of recursive pruning that Diestel describes in his book [26] as follows:

Definition 12.5.5 (Recursive pruning). Let T be any tree, equipped with a root and the corresponding tree-order on its vertices. We recursively label the vertices of T by ordinals, as follows. Given an ordinal α , assume that we have decided for every $\beta < \alpha$ which of the vertices of T to label β , and let T_α be the subgraph of T induced by the vertices that are still unlabelled. Assign label α to every vertex t of T_α whose up-closure $[t]_{T_\alpha} = [t]_T \cap T_\alpha$ in T_α is a chain. The recursion terminates at the first α not used to label any vertex; for this α we put $T_\alpha =: T^*$. We call T *recursively prunable* if every vertex of T gets labelled in this way, i.e., if $T^* = \emptyset$.

Proposition 12.5.6 ([26, Proposition 8.5.1]). *A rooted tree is recursively prunable if and only if it contains no subdivision of the infinite binary tree.*

The next lemma is an observation that we will use often:

Lemma 12.5.7. *Suppose that G is an infinitely edge-connected graph, that u, v are two distinct vertices of G , and that C is a component of $G - u - v$. If T is a finitely separating spanning tree of \tilde{C} and $t \in T$ has finite degree in T , then $C[t]$ is infinitely edge-connected and either u or v sends infinitely many edges in G to the part $t \subseteq V(C)$.*

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

Proof. As t has finite degree in T , the finite fundamental cuts of the edges of T incident with t together give rise to a finite cut of C with the part t as one of its sides. Thus, in the graph G every vertex in the part t can send only finitely many edges to $C - t$, at most one edge to each of u and v , and some edges to the rest of the part t . As every vertex of the infinitely edge-connected graph G has infinite degree, it follows that the part t must be infinite. And since no two vertices in t are finitely separable in C while t is separated from the rest of C by a single finite cut, it follows that $C[t]$ is infinitely edge-connected. Finally, at least one of u and v sends infinitely many edges to the part t , because otherwise t is separated from the rest of G by a finite cut, contradicting that G is infinitely edge-connected. \square

Here is the second main lemma of this section:

Lemma 12.5.8. *Suppose that G is an infinitely edge-connected graph, that u, v are two distinct vertices of G , and that C is a component of $G - u - v$ such that u sends at least one edge to C . If \tilde{C} has a recursively prunable finitely separating rooted spanning tree T such that u sends no edges to parts $t \in T$ that are finite-degree nodes of T , then $G[C + u + v]$ contains an arrow barrage minor from u to v .*

Proof. Given T , we let $X \subseteq V(T)$ consist of the 0-labelled nodes of T that are minimal in the tree-order. Then the nodes in X form a maximal antichain in the tree-order, giving $T = \lfloor X \rfloor \cup \lceil X \rceil$, as T is recursively prunable. Note that all the nodes in $\lfloor X \rfloor$ have degree at most two in T . We claim that X must be infinite. Indeed, if X is finite, then so is $\lceil X \rceil$, and in particular T is locally finite. But then u may send no edges to C by assumption, contradicting our other assumption that u does send an edge to C . Therefore, X must be infinite.

Recall that the finitely separating spanning tree $T \subseteq \tilde{C}$ gives rise to a $\mathcal{B}_{\aleph_0}(C)$ -tree (T, α) . For every $x \in X$ let us write $(A_x, B_x) := \alpha(x, p_x)$ for the predecessor p_x of x in T . As u sends some edges to C , but none to the parts in $\lfloor X \rfloor$, there is a neighbour w of u in the part $\bigcap_{x \in X} B_x$ of the star $\sigma := \{ (A_x, B_x) \mid x \in X \}$. By Lemma 12.5.2 we find an infinite subset $Y \subseteq X$ such that the part of the infinite substar $\sigma' := \{ (A_y, B_y) \mid y \in Y \} \subseteq \sigma$ is connected. Note that w is contained in the part of σ' because the part of σ is included in the part of σ' . We now find an arrow barrage minor from u to v in $G[C + u + v]$ as follows. For the branch set of the nock we take the part of σ' plus the vertex u . For the branch set of the head we take $\{v\}$. The payloads we let be modelled by the subgraphs $C[y]$, one for every $y \in Y$ (here, each $C[y]$ is infinitely edge-connected and sends infinitely many edges in G to v by Lemma 12.5.7 and $Y \subseteq X$). \square

12.5.3. Football minors

We are almost ready now to prove Theorem 12.5.13. But first, we prove an intermediate proposition, which requires the following lemma and definition:

Lemma 12.5.9. *If G is an infinitely edge-connected graph and G' is obtained from G by contracting disjoint finite vertex sets that are possibly disconnected, then G' is infinitely edge-connected as well.*

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

Proof. To show that G' is infinitely edge-connected, consider any two distinct vertices x and y of G' , and choose vertices $\tilde{x} \in x$ and $\tilde{y} \in y$ of G . Now, in the infinitely edge-connected graph G we choose infinitely many pairwise edge-disjoint \tilde{x} – \tilde{y} paths P_0, P_1, \dots as follows. To get started, choose P_0 arbitrarily. At step $n > 0$, consider all the branch sets that are met by some P_k with $k < n$, and let X_n be their union. Then X_n is finite, and we let P_n be an \tilde{x} – \tilde{y} path in G that avoids all the finitely many edges of G running inside X_n .

Now every \tilde{x} – \tilde{y} path $P_n \subseteq G$ gives rise to some x – y path $P'_n \subseteq G'$ satisfying $E(P'_n) \subseteq E(P_n)$ by a slight abuse of notation. We claim that the paths P'_0, P'_1, \dots are all edge-disjoint. For this, consider any two paths P'_n and P'_m with $n < m$, and let e be any edge of P'_n . Then e , viewed as an edge of G , runs between two branch sets that P_n meets because it uses e . Hence these two branch sets are both included in X_m , and so P_m does not use any of the edges running between them. In particular, P'_m does not use e . \square

Definition 12.5.10 (Football, Muscle). Suppose that u and v are two distinct vertices.

A *football* with *endvertices* u and v is an infinitely edge-connected graph G containing u and v such that $G - u - v$ is again infinitely edge-connected.

When we say that some graph contains a football minor *connecting* two vertices x and y we mean that the graph contains a football minor with some endvertices u and v such that the branch set corresponding to u contains x and the branch set corresponding to v contains y (or vice versa).

A *muscle* with *endvertices* u and v is a graph G that is obtained from the vertices u and v by disjointly adding an infinitely edge-connected graph H and adding one u – H edge ux and one v – H edge vy such that $x \neq y$.

A *muscle barrage* with *endvertices* u and v is a countably infinite union $\bigcup_{n \in \mathbb{N}} G_n$ of muscles G_n with endvertices u and v such that G_n and G_m do not meet in any vertices other than u and v for all $n \neq m$.

Muscle (barrage) minors *connecting* two vertices are defined like for footballs.

Proposition 12.5.11. *Suppose that G is an infinitely edge-connected graph, that u, v are two distinct vertices of G , and that C is a component of $G - u - v$ to which both u and v do send some edges. Then at least one of the following assertions holds:*

- (i) $G[C + u + v]$ contains a $T_{\aleph_0} * t$ minor;
- (ii) $G[C + u + v]$ contains a football minor connecting u and v ;
- (iii) $G[C + u + v]$ contains an arrow barrage minor either from u to v or from v to u ; in particular, $G[C + u + v]$ contains a muscle barrage minor connecting u and v .

Proof. We may assume that both u and v send infinitely many edges to C . Indeed, if—say— u sends only finitely many edges to C , then consider the infinitely edge-connected graph $G' := G[C + v]$ and let u' be one of the neighbours of u in C . If there is a component C' of $G' - u' - v$ to which both u' and v send infinitely many edges, then we may replace G, u, v, C with G', u', v, C' . Hence we may assume that

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

there are infinitely many components C'_0, C'_1, \dots of $G' - u' - v$ such that, without loss of generality, u' sends only finitely many but at least one edge to each C'_n while v sends infinitely many edges to each C'_n .

By Theorem 12.5.1, all \tilde{C}'_n have finitely separating spanning trees. If one \tilde{C}'_n has a finitely separating spanning tree that contains a subdivision of the infinite binary tree, then Lemma 12.5.3 provides a $T_{\aleph_0} * t$ minor witnessing (i). Otherwise, by Proposition 12.5.6, every \tilde{C}'_n has a rooted finitely separating spanning tree T_n that is recursively prunable. Then we pick for every n a finite-degree node $t_n \in T_n$, and we let P_n be a path in C'_n that links a neighbour of u' to the subgraph $C'_n[t_n]$ such that P_n has only its endvertex x_n in $C'_n[t_n]$. Now we obtain an arrow barrage minor in $G[C + u + v]$ from u to v that is sought in (iii), as follows. For the branch set of the arrow barrage's nock we take $\{u, u'\} \cup \bigcup_{n \in \mathbb{N}} V(P_n \dot{x}_n)$. The arrows' payloads we let be modelled by the infinitely edge-connected subgraphs $C''[t_n]$ (see Lemma 12.5.7). And for the branch set of the arrow barrage's head we take $\{v\}$ (that v sends infinitely many edges to each part t_n is ensured by Lemma 12.5.7 and the assumption that u' sends only finitely many edges to each C'_n).

Therefore, we may assume that both u and v send infinitely many edges to C . By Theorem 12.5.1 we may let T be a finitely separating spanning tree of \tilde{C} , rooted arbitrarily. We make the following two observations.

If T contains a subdivision of the infinite binary tree, then Lemma 12.5.3 yields a $T_{\aleph_0} * t$ minor giving (i).

If T has finite-degree nodes t_u and t_v (possibly $t_u = t_v$) such that u sends infinitely many edges to the part $t_u \subseteq V(C)$ in G and v sends infinitely many edges to the part $t_v \subseteq V(C)$ in G , then we deduce (ii), as follows. By Lemma 12.5.7 both $C[t_u]$ and $C[t_v]$ are infinitely edge-connected. If $t_u = t_v$, then $G[t_u + u] \cup G[t_v + v]$ is a football subgraph connecting u and v . Otherwise t_u and t_v are distinct. Then we let P be any t_u - t_v path in C , and $(G[t_u + u] \cup G[t_v + v] \cup P)/V(P)$ is a football minor connecting u and v .

By these two observations and Proposition 12.5.6, we may assume that T is recursively prunable and that, without loss of generality, whenever $t \in T$ has finite degree then v does send infinitely many edges to the part $t \subseteq V(C)$ in G while u may send only finitely many edges to it.

If u sends edges in G to infinitely many parts $t \in T$ that have finite degree in T , then we find an arrow barrage minor from u to v giving (iii), because v sends infinitely many edges to all of the infinitely edge-connected subgraphs $C[t]$ (cf. Lemma 12.5.7) by our assumption above. Otherwise u sends, in total, only finitely many edges in G to the parts $t \in T$ that have finite degree in T . Since u sends infinitely many edges in G to C , this means that we may assume without loss of generality that u sends no edges to the parts $t \in T$ that have finite degree in T . Then Lemma 12.5.8 yields an arrow barrage minor from u to v giving (iii). \square

Now we have all we need to prove the main result of the section, Theorem 12.5.13. In its proof, we will face the construction of a minor in countably many steps. The following notation and lemma will help us to keep the technical side of this construction to the minimum.

Suppose that G and H are two graphs with H a minor of G . Then there are

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

a vertex set $U \subseteq V(G)$ and a surjection $f: U \rightarrow V(H)$ such that the preimages $f^{-1}(x) \subseteq U$ form the branch sets of a model of H in G . A *minor-map* $\varphi: G \succcurlyeq H$ formally is such a pair (U, f) . Given $\varphi = (U, f)$ we address U as $V(\varphi)$ and we write $\varphi = f$ by abuse of notation. Usually, we will abbreviate ‘minor-map’ as ‘map’. If we are given two maps $\varphi: G \succcurlyeq H$ and $\varphi': H \succcurlyeq H'$, then these give rise to another map $\psi: G \succcurlyeq H'$ by letting $V(\psi) := \varphi^{-1}(\varphi'^{-1}(V(H')))$ and $\psi := \varphi' \circ (\varphi \upharpoonright V(\psi))$. On the notational side we write $\varphi' \diamond \varphi = \psi$.

Lemma 12.5.12. *If G_0, G_1, \dots and $H_0 \subseteq H_1 \subseteq \dots$ are sequences of graphs $H_n \subseteq G_n$ with maps $\varphi_n: G_n \succcurlyeq G_{n+1}$ that restrict to the identity on H_n , then $G \succcurlyeq \bigcup_{n \in \mathbb{N}} H_n$.*

Proof. Recursively, each map $\varphi_n: G_n \succcurlyeq G_{n+1}$ gives rise to a map $\hat{\varphi}_n: G_0 \succcurlyeq G_{n+1}$ via $\hat{\varphi}_0 := \varphi_0$ and $\hat{\varphi}_{n+1} := \varphi_{n+1} \diamond \hat{\varphi}_n$. For every $n \in \mathbb{N}$ we write $V_x^n = \hat{\varphi}_n^{-1}(x)$ for all vertices $x \in H_{n+1}$. For every vertex $x \in H := \bigcup_{n \in \mathbb{N}} H_n$ we denote by $N(x)$ the least number n with $x \in H_n$. As the maps φ_n restrict to the identity on H_n , for every vertex $x \in H$ the vertex sets V_x^n form an ascending sequence $V_x^{N(x)} \subseteq V_x^{N(x)+1} \subseteq \dots$ whose overall union we denote by V_x . We claim that the vertex sets V_x form the branch sets of an H minor in G .

Indeed, every branch set V_x is non-empty and connected in G because all V_x^n are. If xy is an edge of H , then G contains a $V_x^n - V_y^n$ edge as soon as $xy \in H_n$, and this edge is a $V_x - V_y$ edge due to the inclusions $V_x^n \subseteq V_x$ and $V_y^n \subseteq V_y$. It remains to show that V_x and V_y are disjoint for distinct vertices $x, y \in H$. This follows at once from the vertex sets V_x^n and V_y^n being disjoint for all n and the definition of V_x and V_y as ascending unions of these vertex sets. \square

Finally, we prove the main result of the section:

Theorem 12.5.13. *Suppose that G is any infinitely edge-connected graph, that u, v are two distinct vertices of G , and that C is a component of $G - u - v$ to which both u and v do send some edges. Then at least one of the following assertions holds:*

- (i) $G[C + u + v]$ contains a $T_{\aleph_0} * t$ minor;
- (ii) $G[C + u + v]$ contains a football minor connecting u and v .

Proof. Assume for a contradiction that both (i) and (ii) fail. We will use Proposition 12.5.11 to find the following graph H as a minor in $G' := G[C + u + v]$. Let T_u be an \aleph_0 -regular tree with root r_u , and let T_v be a copy of T_u that is disjoint from T_u . We write r_v for the root of T_v . The graph H is obtained from the disjoint union of the two trees T_u and T_v by adding the perfect matching between their vertex sets that joins every vertex of T_u to its copy in T_v . For every number $n \in \mathbb{N}$ we write H_n for the subgraph of H that is induced by the first n levels of T_u together with the first n levels of T_v . Thus, $H = \bigcup_{n \in \mathbb{N}} H_n$. Finding an H minor in G' completes the proof, because H/T_u is isomorphic to $T_{\aleph_0} * t$.

A *foresighted* H_n is a graph that is obtained from H_n by adding for every edge $xy \in H_n$ that runs between the two n th levels of T_u and T_v a muscle barrage B_{xy}

having endvertices x and y such that B_{xy} contains no vertices from H_n other than x and y , and all muscle barrages added are pairwise disjoint.

By Lemma 12.5.12 it suffices to find a sequence $G' \succ \hat{H}_0 \succ \hat{H}_1 \succ \dots$ of graphs \hat{H}_n that are foresighted H_n with maps $\varphi_n: \hat{H}_n \succ \hat{H}_{n+1}$ that restrict to the identity on $H_n \subseteq \hat{H}_n$ in order to find an H minor in $\hat{H}_0 \preceq G'$. To get started, we apply Proposition 12.5.11 to G, u, v, C to obtain in G' a muscle barrage minor connecting u and v . By turning one of the muscles into an edge we obtain $\hat{H}_0 \preceq G'$.

At step $n > 0$, consider the muscle barrages B_{xy} that turn H_n into \hat{H}_n . For every muscle M_{xy}^k of each of these muscle barrages $B_{xy} = \bigcup_{k \in \mathbb{N}} M_{xy}^k$ we apply Proposition 12.5.11 in $M := M_{xy}^k - x - y$ to the neighbours x' and y' of x and y in M_{xy}^k and some component of $M - x' - y'$ to which both x' and y' send some edges to find a muscle barrage minor connecting x' and y' . By turning one muscle of each new barrage into an edge, we find $\varphi_n: \hat{H}_n \succ \hat{H}_{n+1}$. \square

12.6. Proof of the main result

In this section we employ the main result of the previous section (Theorem 12.5.13) to prove the main result of this chapter (Theorem 12.1).

Lemma 12.6.1. *If A and B are two infinite vertex sets in a graph G that does not contain a subdivision of K^{\aleph_0} , then there are vertices $a \in A$ and $b \in B$ plus a finite vertex set $S \subseteq V(G) \setminus \{a, b\}$ such that S separates a and b in $G - ab$.*

Proof. The absence of such an S for a pair $a \neq b$ means that, inductively, we can find infinitely many independent a - b paths in G . So if there is no S for every pair $a \neq b$, then inductively we find a TK_{\aleph_0, \aleph_0} in G , and $TK^{\aleph_0} \subseteq TK_{\aleph_0, \aleph_0}$ (contradiction). \square

Lemma 12.6.2. *Suppose that G is a football with endvertices u and v . If G does not contain a subdivision of K^{\aleph_0} , then G contains an infinitely edge-connected graph H as a minor with branch sets V_h ($h \in H$) such that u and v are contained in distinct branch sets V_x and V_y , respectively, and there is a finite vertex set $S \subseteq V(H) \setminus \{x, y\}$ separating x and y in H .*

Proof. Write C for the infinitely edge-connected graph $G - u - v$. We apply Lemma 12.6.1 in C to the infinite neighbourhoods $N(u)$ and $N(v)$ of u and v in G to obtain vertices $a \in N(u)$ and $b \in N(v)$ plus a finite vertex set $S \subseteq V(C) \setminus \{a, b\}$ that separates a and b in $C - ab$. Then H can be obtained from the infinitely edge-connected graph $G - ab$ as follows. We discard all the edges that are incident with u or v , except for the two edges ua and vb each of which we contract. Then H is infinitely edge-connected because it is isomorphic to the infinitely edge-connected graph $C - ab$. And the way we treated the edges at u and v ensures that S separates the two vertices $\{u, a\}$ and $\{v, b\}$ in H as desired. \square

Lemma 12.6.3. *Suppose that G is an infinitely edge-connected graph and that u, v are two distinct vertices of G that are separated in G by some finite vertex set $S \subseteq V(G) \setminus \{u, v\}$. Then there exist induced subgraphs $H_u, H_v \subseteq G$ containing u and v respectively, such that the following assertions hold:*

- (i) $X := V(H_u) \cap V(H_v)$ is finite, non-empty and connected in G ;
- (ii) both H_u/X and H_v/X are infinitely edge-connected;
- (iii) X avoids u and v ;
- (iv) uX is an edge of H_u/X and vX is an edge of H_v/X .

Proof. Given G, u, v, S let us write C_u and C_v for the distinct components of $G - S$ that contain u and v respectively. For both $w \in \{u, v\}$ we abbreviate $\sim_{G[C_w \cup S]}$ as \sim_w . As G is infinitely edge-connected, we infer that every \sim_w -class meets S . In particular, there are only finitely many \sim_w -classes in total, which means that each of the non-singleton classes induces an infinitely edge-connected subgraph of G . Let us write K_u and K_v for the infinitely edge-connected subgraphs induced by the classes containing u and v respectively, i.e., $K_u := G[[u]_{\sim_u}]$ and $K_v := G[[v]_{\sim_v}]$. To find H_u and H_v , we distinguish two cases.

In the first case, K_u and K_v are disjoint. For both $w \in \{u, v\}$, the finite partition of $V(C_w) \cup S$ induced by \sim_w has only finitely many cross-edges. Since G is infinitely edge-connected, this means that we can find a $(K_u \cap S) - (K_v \cap S)$ path P in G avoiding all these finitely many edges. Then P , as it may not use these edges, is a $K_u - K_v$ path with endvertices in S . We let P_w be a $w - P$ path in K_w for both $w \in \{u, v\}$. Letting $H_u := G[K_u \cup P \cup \dot{v}P_v]$ and $H_v := G[K_v \cup P \cup \dot{u}P_u]$ completes this case with $X = V(P_u \cup P \cup P_v) \setminus \{u, v\}$ because the graph H_w/X contains the spanning subgraph $K_w/V(\dot{w}P_w)$, and $K_w/V(\dot{w}P_w)$ is infinitely edge-connected by Lemma 12.5.9 and because K_w is infinitely edge-connected.

In the second case, K_u and K_v meet in a vertex $s \in S$. We write D_u for the component of $K_u - u$ containing s . In D_u we pick a finite tree T that contains the finite intersection $V(D_u) \cap V(K_v) \subseteq S$ and contains a neighbour of u . Then T contains s but neither u nor v . We let P_v be any $v - s$ path in K_v . Letting $H_u := G[D_u \cup \dot{v}P_v + u]$ and $H_v := G[K_v \cup T]$ completes this case with $X = V(T \cup \dot{v}P_v)$: On the one hand, the graph H_u/X is infinitely edge-connected because it contains the spanning subgraph $G[D_u + u]/V(T)$ which is infinitely edge-connected by Lemma 12.5.9 and the fact that $G[D_u + u]$ is an infinitely edge-connected subgraph of K_u . On the other hand, the graph H_v/X contains the spanning subgraph K_v/Y for $Y := (V(K_v) \cap V(D_u)) \cup V(\dot{v}P_v)$, and K_v/Y is infinitely edge-connected by Lemma 12.5.9 and because K_v is infinitely edge-connected. \square

Definition 12.6.4 (Plows). Suppose that u and v are two distinct vertices. A *half-plow* with *endvertices* u and v is an infinitely edge-connected graph containing the edge uv . A *plow* with *endvertices* u and v and *head* h is a union of two half-plows with end-vertices u, h and h, v that do not meet in any vertex other than h . Plow minors *connecting* some two vertices are defined like for footballs and muscles.

12. Infinite edge-connectivity, the Farey graph and $T_{\aleph_0} * t$

Theorem 12.6.5. *If G is an infinitely edge-connected graph and u, v are two distinct vertices of G , then at least one of the following two assertions holds:*

- (i) G contains a $T_{\aleph_0} * t$ minor;
- (ii) G contains a plow minor connecting u and v .

Proof. Let G, u, v be given, we show $\neg(\text{i}) \rightarrow (\text{ii})$. For this, suppose that G does not contain a $T_{\aleph_0} * t$ minor. By Theorem 12.5.13 and Lemma 12.6.2 we may assume that there is a finite vertex set $S \subseteq V(G) \setminus \{u, v\}$ that separates u and v in G . Then applying Lemma 12.6.3 provides induced subgraphs $H_u, H_v \subseteq G$ containing u and v respectively, such that the following assertions hold:

- $X := V(H_u) \cap V(H_v)$ is finite, non-empty and connected in G ;
- both H_u/X and H_v/X are infinitely edge-connected;
- X avoids u and v ;
- uX is an edge of H_u/X and vX is an edge of H_v/X .

Then $(H_u \cup H_v)/X$ is a plow minor connecting u and v . □

Theorem 12.1. *Every infinitely edge-connected graph contains either the Farey graph or $T_{\aleph_0} * t$ as a minor.*

Proof. If G contains $T_{\aleph_0} * t$ as a minor, then we are done. So let us suppose that G does not contain a $T_{\aleph_0} * t$ minor. Our task then is to find a Farey graph minor in G . By Lemma 12.2.1 it suffices to find a halved Farey graph minor.

Call a graph a *foresighted* halved Farey graph of order $n \in \mathbb{N}$ if it is the union of \check{F}_n with infinitely edge-connected graphs A_{xy} , one for every blue edge $xy \in \check{F}_n$, such that:

- (i) each A_{xy} meets \check{F}_n precisely in x and y but $xy \notin A_{xy}$;
- (ii) every two distinct A_e and $A_{e'}$ meet precisely in the intersection $e \cap e'$ of their corresponding edges (viewed as vertex sets).

By Lemma 12.5.12 it suffices to find a sequence H_0, H_1, \dots of foresighted halved Farey graphs of orders $0, 1, \dots$ with maps $\varphi_n: H_n \succcurlyeq H_{n+1}$ that restrict to the identity on $\check{F}_n \subseteq H_n$ to yield a halved Farey graph minor in $G =: H_0$.

To get started, pick any edge e of G , and note that $G = H_0$ is a foresighted halved Farey graph of order 0 when we rename e to the edge of which $\check{F}_0 = K^2$ consists. At step $n + 1$, suppose that we have already constructed $H_n \supseteq \check{F}_n$, and consider the infinitely edge-connected graphs A_{xy} that were added to \check{F}_n to form H_n . Theorem 12.6.5 yields in each A_{xy} a plow minor with head h_{xy} that connects x and y . These plow-minors combine with \check{F}_n and with each other to give a map $\varphi_n: H_n \succcurlyeq H_{n+1} \supseteq \check{F}_{n+1}$ that sends the branch set of every head h_{xy} to the vertex $v_{xy} \in \check{F}_{n+1} - \check{F}_n$ that arises from the blue edge $xy \in \check{F}_n$ in the recursive definition of \check{F}_{n+1} . □

12.7. Outlook

Here are two open problems that came to my mind.

Problem 12.7.1. *Can Theorem 12.1 be strengthened to always find one of the two minors with finite branch sets?*

Seymour and Thomas [77], together with Robertson [74, 75], have characterised the graphs without K^κ or T_κ minors in terms of tree-decompositions and, alternatively, in terms of various other structures. Can their list be extended to include the Farey graph? Tree-decompositions might not be the right complementary structures for infinitely edge-connected substructures, but there might be other structures (e.g. $\mathcal{B}_{\aleph_0}(G)$ -trees):

Problem 12.7.2. *Characterise the graphs without a Farey graph minor in terms of tree-decompositions or in terms of other structures.*

13. The Farey graph is uniquely determined by its connectivity




Figure 13.0.1.: The Farey graph

13.1. Introduction

The Farey graph, shown in Figure 13.0.1 and surveyed in [22, 49], plays a role in a number of mathematical fields ranging from group theory and number theory to geometry and dynamics [22]. Curiously, graph theory has not been among these until very recently, when it was shown that the Farey graph plays a central role in graph theory too: it is one of two infinitely edge-connected graphs that must occur as a minor in every infinitely edge-connected graph; see Chapter 12. Infinite edge-connectivity, however, is only one aspect of the connectivity of the Farey graph, and it contrasts with a second aspect: the Farey graph does not contain infinitely many independent paths between any two of its vertices. In this chapter we show that the Farey graph is uniquely determined by these two contrasting aspects of its connectivity: up to minor-equivalence, the Farey graph is the unique minor-minimal graph that is infinitely edge-connected but such that every two vertices can be finitely separated. This is the first graph-theoretic characterisation of the Farey graph.

A Π -graph is an infinitely edge-connected graph that does not contain infinitely many independent paths between any two of its vertices. A Π -graph is *typical* if it occurs as a minor in every Π -graph. Note that any two typical Π -graphs are minors of each other; we call such graphs *minor-equivalent*. Our main result reads as follows:

Theorem 13.1. *Up to minor-equivalence, the Farey graph is the unique typical Π -graph.*

We shall see that there exist Π -graphs that contain the Farey graph as a minor but are not minors of the Farey graph (Theorem 13.3.1).

13. The Farey graph is uniquely determined by its connectivity

Theorem 13.1 continues to hold if we require all minors to have finite branch sets; see Section 13.3.2 and Theorem 13.3.2. This is best possible in the sense that one cannot replace ‘minors with finite branch sets’ with ‘topological minors’ (Theorem 13.3.3).

This chapter is organised as follows. Section 13.2 formally introduces the Farey graph. In Section 13.3 we prove Theorems 13.3.1–13.3.3. We outline the overall strategy of the proof of Theorem 13.1 in Section 13.4. We prepare the proof of Theorem 13.1 in Section 13.5 and we prove Theorem 13.1 in Section 13.6.

13.2. Preliminaries

We defined the Farey graph in Chapter 12. Two u – v paths are *order-compatible* if they traverse their common vertices in the same order.

Lemma 13.2.1. *The halved Farey graph contains the Farey graph as a minor with finite branch sets.*

Proof. If e is a blue edge of \check{F}_1 , then the Farey graph is the contraction minor of $\check{F} - e$ whose sole non-trivial branch set is $V(\check{F}_0)$, i.e., $(\check{F} - e)/V(\check{F}_0) \cong F$. \square

13.3. Atypical Π -graphs and variations of the main result

In this section we provide details on and prove the three Theorems 13.3.1–13.3.3 that we briefly mentioned in the introduction.

13.3.1. Atypical Π -graphs

Even though every Π -graph contains the Farey graph as a minor by Theorem 13.1, the converse is generally false:

Theorem 13.3.1. *There exist Π -graphs that contain the Farey graph as a minor but are not minors of the Farey graph.*

Proof. Let the graph G be obtained from some union of uncountably many disjoint copies of the Farey graph by selecting one vertex in every copy and identifying all the selected vertices. Then G is an uncountable Π -graph that contains the Farey graph as a subgraph. However, G is not a minor of the Farey graph, because every minor of the Farey graph must be countable. \square

13.3.2. Variations of the main result

To prove Theorem 13.1 it suffices to show the following theorem. A *tight* minor is a minor with finite branch sets.

Theorem 13.2. *Every Π -graph contains the Farey graph as a tight minor.*

13. The Farey graph is uniquely determined by its connectivity

Theorem 13.2 also implies the following variation of Theorem 13.1 where all minors are required to have finite branch sets. Two graphs are *tightly* minor-equivalent if they are tight minors of each other. A Π -graph is *tightly* typical if it occurs as a tight minor in every Π -graph.

Theorem 13.3.2. *Up to tight minor-equivalence, the Farey graph is the unique tightly typical Π -graph.* \square

This raises the question whether Theorem 13.1 continues to hold if we require all minors to be topological minors. We answer this question in the negative:

Theorem 13.3.3. *There is a Π -graph that contains the Farey graph as a tight minor but not as a topological minor.*

Proof. As we will see in Chapter 14, there exists an infinitely edge-connected graph G that does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices; in particular, G is a Π -graph. By Theorem 13.2, the graph G contains the Farey graph as a tight minor. However, G does not contain a subdivision of the Farey graph because the Farey graph contains infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices. \square

13.4. Overall proof strategy

Our aim for the remainder of this chapter is to prove Theorem 13.1. As we discussed in the previous section, to prove Theorem 13.1 it suffices to show that every Π -graph contains the Farey graph as a minor with finite branch sets (Theorem 13.2). And by Lemma 13.2.1 in turn it suffices to find a halved Farey graph minor with finite branch sets in any given Π -graph. The key idea of the proof is summarised in Theorem 13.6.1 which states:

Suppose that G is any subdivided Π -graph and that u, v are two distinct branch vertices of G . Then there exist subgraphs $H_u, H_v \subseteq G$ that satisfy the following conditions:

- (i) $H_u[X] = H_v[X]$ is finite and connected for $X := V(H_u) \cap V(H_v) \neq \emptyset$;
- (ii) X avoids u and v ;
- (iii) both H_u/X and H_v/X are subdivided Π -graphs in which u, X and v, X are branch vertices, respectively;
- (iv) uX is an edge of H_u/X and vX is an edge of H_v/X .

With this theorem at hand, it is straightforward to construct a halved Farey graph minor with finite branch sets in any given Π -graph G : Consider any edge uv of G and apply the theorem in G to u and v to obtain subgraphs H_u, H_v and a non-empty finite connected vertex set $X \subseteq V(G)$. Then the three vertices u, v and X span a triangle \tilde{F}_1 in $(H_u \cup H_v)/X$. And since both H_u/X and H_v/X are subdivided Π -graphs, we can reapply the theorem in H_u/X to u and X , and in H_v/X to v and X . By iterating this process, we obtain a halved Farey graph minor

13. The Farey graph is uniquely determined by its connectivity

with finite branch sets in the original graph G at the limit, and this will complete the proof. Therefore, it remains to prove Theorem 13.6.1 on the one hand, and to use it to formally construct a halved Farey graph minor on the other hand. In the next section, we prepare the proof of Theorem 13.6.1, and in the section after next we prove Theorem 13.6.1 which we then use to prove Theorems 13.1 and 13.2.

13.5. Grain lines

It is possible to prove Theorem 13.6.1 from first principles. In this chapter, however, I favour a more methodic proof. The advantage of this proof is that it introduces a new tool, an x - y grain line, that allows one to control infinite systems of edge-disjoint x - y paths even when no two paths in the system are pairwise order-compatible. In this section we introduce the concept of an x - y grain line, we show that these exist whenever it matters (Theorem 13.5.4) and we show two lemmas that will help us prove Theorem 13.6.1 using grain lines at the beginning of the next section.

Informally, we may think of an x - y grain line as a pair (L, \mathcal{P}) where \mathcal{P} is a sequence of pairwise edge-disjoint x - y paths P_0, P_1, \dots that need not be pairwise order-compatible but solve all incompatibilities at their linearly ordered ‘limit’ L . The limit L will not be a graph-theoretic path but will be a linearly ordered set of vertices. We remark, however, that it is possible to use the limit L to define a topological x - y path in a topological extension of any graph containing the grain line, see [57, §6.3].

Here is the formal definition of an x - y grain-line:

Definition 13.5.1. An x - y grain line between two distinct vertices x and y is an ordered pair (L, \mathcal{P}) where $L = (L, \leq_L)$ is a linearly ordered countable set of vertices with least element x and greatest element y , and $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ is a sequence of pairwise edge-disjoint x - y paths P_n , such that the following three conditions are satisfied:

- (GL1) $L = \left\{ v \mid \{ n \in \mathbb{N} : v \in V(P_n) \} \text{ is a final segment of } \mathbb{N} \right\}$;
- (GL2) if a vertex of a path P_n is not contained in L , then it is not a vertex of any other path P_m ($m \neq n$);
- (GL3) for all $n \in \mathbb{N}$, the x - y path P_n and the linearly ordered vertex set L induce the same linear ordering on the vertex set $L_{<n} := L \cap \bigcup_{k < n} V(P_k)$.

We remark that (GL3) allows P_n and L to induce distinct linear orderings on the vertex set $V(P_n) \cap L$ if the inclusion $L_{<n} \subseteq V(P_n) \cap L$ is proper; in particular, P_n and P_{n+1} need not be order-compatible. Allowing this becomes necessary, for example, if an infinitely edge-connected graph does not contain infinitely many edge-disjoint pairwise order compatible paths between x and y , see Example 13.5.3.

Clearly, $L = \bigcup_n L_{<n}$. Note that if $(L, (P_n)_{n \in \mathbb{N}})$ is a grain line, then a vertex v lies in L if and only if it lies on all paths P_n with $n \geq N$ for N the first number with

13. The Farey graph is uniquely determined by its connectivity

$v \in P_N$ if and only if it lies on at least two paths P_n, P_m ($n \neq m$). In particular,

$$V(P_n) \cap \bigcup_{k < n} V(P_k) = L_{<n} \text{ for all } n \in \mathbb{N}.$$

We speak of an x - y grain line $(L, (P_n)_{n \in \mathbb{N}})$ in a graph G if $\bigcup_{n \in \mathbb{N}} P_n \subseteq G$ (and hence $L \subseteq V(G)$). Whenever a grain line is introduced as (L, \mathcal{P}) , we tacitly assume $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$. In general, however, we also allow sequences $\mathcal{P} = (P_n)_{n \geq N}$ whose indexing starts at an arbitrary number $N > 0$ in which case the definition of a grain line adapts in the obvious way. We use the interval notation for L as usual, i.e., we write $[\ell_1, \ell_2]_L = \{\ell \in L \mid \ell_1 \leq_L \ell \leq_L \ell_2\}$ and so on.

Example 13.5.2. The blue Hamilton paths $P_n \subseteq \check{F}_n$ are pairwise edge-disjoint and order-compatible, and hence give rise to an x - y grain line in \check{F} for x and y the two vertices of \check{F}_0 . In this case, $L = V(\check{F})$ is order-isomorphic to $\mathbb{Q} \cap [0, 1]$.

Example 13.5.3. We will show in Chapter 14 that there exists an infinitely edge-connected graph G that does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices; in particular, G is a Π -graph. We shall see that the graph G contains a grain line between any two of its vertices because it is infinitely edge-connected; see Theorem 13.5.4 below. However, since G does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices, every grain line (L, \mathcal{P}) in G has two paths P_n and P_{n+1} that are not order-compatible; in particular, P_n induces the same linear ordering on $L_{<n} \subsetneq V(P_n) \cap L$ as L does, but disagrees with L on $V(P_n) \cap L = V(P_n) \cap V(P_{n+1})$ because L induces the ordering of P_{n+1} on $V(P_n) \cap V(P_{n+1}) = L_{<n+1}$. This is why we do not strengthen (GL3) to require that P_n and L induce the same linear ordering on $V(P_n) \cap L \supseteq L_{<n}$.

Our first result on grain lines shows that they exist whenever it matters:

Theorem 13.5.4. *Let x and y be any two distinct vertices of a graph G . Then there exists an x - y grain line in G if and only if G contains infinitely many edge-disjoint x - y paths.*

Proof of Theorem 13.5.4. Every x - y grain line comes with a system of infinitely many edge-disjoint x - y paths. For the backward implication let x and y be given, and let \mathcal{Q} be any countably infinite collection of edge-disjoint x - y paths in G . Moreover, we let \mathcal{X} be the collection of all finite subsets of the vertex set of the subgraph $\bigcup \mathcal{Q} \subseteq G$, directed by inclusion.

Given $X \in \mathcal{X}$ we write $\text{lin}(X)$ for the finite collection of all linearly ordered subsets of X . Letting, for all $X \subseteq X' \in \mathcal{X}$, the map $\varphi_{X',X}: \text{lin}(X') \rightarrow \text{lin}(X)$ take every linearly ordered subset of X' to its restriction with respect to X turns the finite sets $\text{lin}(X)$ into an inverse system $\{\text{lin}(X), \varphi_{X',X}, \mathcal{X}\}$.

Every x - y path $P \in \mathcal{Q}$ naturally induces a linear ordering \leq_P on its vertex set with $x <_P y$, and for every $X \in \mathcal{X}$ we denote by \leq_P^X the linear ordering on $V(P) \cap X$ induced by \leq_P . Then for every $X \in \mathcal{X}$ we define a map $\psi_X: \mathcal{Q} \rightarrow \text{lin}(X)$ by letting

$$\psi_X(P) := (V(P) \cap X, \leq_P^X)$$

13. The Farey graph is uniquely determined by its connectivity

for all $P \in \mathcal{Q}$, and we put

$$\mathcal{L}_X := \{ \xi \in \text{lin}(X) \mid \psi_X^{-1}(\xi) \subseteq \mathcal{Q} \text{ is infinite} \}$$

noting that $\mathcal{L}_X \subseteq \text{lin}(X)$ is non-empty by the pigeonhole principle. Since the maps ψ_X commute with the bonding maps $\varphi_{X',X}$ as pictured in the diagram below,

$$\begin{array}{ccc} & \mathcal{Q} & \\ \psi_X \swarrow & & \searrow \psi_{X'} \\ \text{lin}(X) & \xleftarrow{\varphi_{X',X}} & \text{lin}(X') \end{array}$$

the restrictions of these bonding maps to the sets \mathcal{L}_X yield another inverse system, namely $\{ \mathcal{L}_X, \varphi_{X',X} \upharpoonright \mathcal{L}_{X'}, \mathcal{X} \}$. And as the finite sets \mathcal{L}_X are all non-empty, this inverse system has an element $((L_X, \leq_X) \mid X \in \mathcal{X})$ in its limit.

Finally, we define an x - y grain line (L, \mathcal{P}) , as follows. We let $L := \bigcup_{X \in \mathcal{X}} L_X$ and $\leq_L := \bigcup_{X \in \mathcal{X}} \leq_X$. To obtain $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ we choose pairwise edge-disjoint x - y paths P_0, P_1, \dots from \mathcal{Q} inductively, as follows. Choose an enumeration x_0, x_1, \dots of the countable vertex set $\bigcup \mathcal{X}$ of $\bigcup \mathcal{Q}$. At step 0, we let $X_0 := \{x_0\}$ and choose $P_0 \in \psi_{X_0}^{-1}(L_{X_0})$ arbitrarily (we abbreviate $L_X = (L_X, \leq_X)$). At step $n+1$, we let $X_{n+1} := X_n \cup V(P_n) \cup \{x_{n+1}\}$ and we pick from the infinite preimage $\psi_{X_{n+1}}^{-1}(L_{X_{n+1}})$ a path P_{n+1} other than the previously chosen paths P_0, \dots, P_n . It is straightforward to check that (L, \mathcal{P}) is an x - y grain line in G . \square

A grain line (L, \mathcal{P}) is *wild* if L is order-isomorphic to $\mathbb{Q} \cap [0, 1]$. We call a grain line (L, \mathcal{P}) *wildly presented* if, for every $n \in \mathbb{N}$, whenever $\ell_1 <_L \ell_2$ are elements of $L_{<n} \subseteq L$ then $\ell_1 P_n \ell_2$ has a vertex in $(\ell_1, \ell_2)_L$. The grain line in Example 13.5.2 is both wild and wildly presented. Wildly presented grain lines are wild. Conversely, if a grain line (L, \mathcal{P}) is wild, then $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ has a subsequence $(P_{n_k})_{k \in \mathbb{N}}$ such that $(L, (P_{n_k})_{k \in \mathbb{N}})$ is wildly presented.

Lemma 13.5.5. *Every grain line in a subdivided Π -graph is wild; in particular, in a subdivided Π -graph every grain line can be chosen to be wildly presented.*

In the proof we use the following properties of grain lines. Given a grain line (L, \mathcal{P}) we say that a path P_n does (L, \mathcal{P}) -*grain* a set U of vertices if, for all $m \geq n$, we have $V(P_m) \cap U = L \cap U$ and the path P_m induces the same linear ordering on this intersection as L does. If (L, \mathcal{P}) is clear from context, we also say that P_n *grains* U . Every path P_n grains the union $\bigcup_{k < n} V(P_k)$ by (GL3). And for every finite vertex set X there is a number $n \in \mathbb{N}$ such that P_n grains X . We will use this latter property frequently in the proofs to come.

Proof of Lemma 13.5.5. Suppose that (L, \mathcal{P}) is any grain line in some given subdivided Π -graph G . It suffices to show that (L, \mathcal{P}) is wild. For this, consider any two elements $\ell_1, \ell_2 \in L$ with $\ell_1 <_L \ell_2$. Then ℓ_1 and ℓ_2 must have infinite degree in G ; in particular, ℓ_1 and ℓ_2 must be branch vertices of G . Since G is a subdivided Π -graph, we find a finite vertex set $S \subseteq V(G) \setminus \{\ell_1, \ell_2\}$ that separates ℓ_1 and ℓ_2 in $G - \ell_1 \ell_2$. Then we pick $N \in \mathbb{N}$ such that P_N avoids the edge $\ell_1 \ell_2$ and grains the finite vertex set $S \cup \{\ell_1, \ell_2\}$. Now $\ell_1 P_N \ell_2$ must meet S in a vertex s , and then P_N graining $S \cup \{\ell_1, \ell_2\}$ implies $s \in L$ with $\ell_1 <_L s <_L \ell_2$ as desired. \square

13. The Farey graph is uniquely determined by its connectivity

Grain lines can be restricted such that the restriction is again a grain line, and restricting a grain line preserves wild presentations:

Lemma 13.5.6. *If (L, \mathcal{P}) is a grain line with $\ell_1 <_L \ell_2$ and $N \in \mathbb{N}$ is such that P_N grains $\{\ell_1, \ell_2\}$, then $([\ell_1, \ell_2]_L, (\ell_1 P_n \ell_2)_{n \geq N})$ is an ℓ_1 - ℓ_2 grain line that is wildly presented if (L, \mathcal{P}) is.*

Proof. First, we show that $([\ell_1, \ell_2]_L, (\ell_1 P_n \ell_2)_{n \geq N})$ is an ℓ_1 - ℓ_2 grain line.

(GL1) We have to show the equality

$$[\ell_1, \ell_2]_L = \left\{ v \mid \{ n \in \mathbb{N}_{\geq N} : v \in V(\ell_1 P_n \ell_2) \} \text{ is a final segment of } \mathbb{N}_{\geq N} \right\}.$$

We start with the backward inclusion. If a vertex v lies on $\ell_1 P_n \ell_2$ for all n in some final segment of $\mathbb{N}_{\geq N}$ then it lies in L by (GL2) for (L, \mathcal{P}) , and in particular it also lies on $\ell_1 P_n \ell_2$ when P_n does (L, \mathcal{P}) -grain $\{\ell_1, v, \ell_2\}$ so $v \in [\ell_1, \ell_2]_L$ follows. Conversely, if v is a vertex in $[\ell_1, \ell_2]_L$ and $k \geq N$ is minimal with $v \in \ell_1 P_k \ell_2$, then P_{k+1} does (L, \mathcal{P}) -grain $\{\ell_1, v, \ell_2\}$. Therefore, v is contained in $\ell_1 P_n \ell_2$ for all $n \geq k$, and hence $\mathbb{N}_{\geq k}$ witnesses that v is contained in the right hand side of the equation.

(GL2) Consider any vertex $v \in (\bigcup_{n \geq N} \ell_1 P_n \ell_2) - [\ell_1, \ell_2]_L$ and let $k \geq N$ be minimal such that $\ell_1 P_k \ell_2$ contains v . If v is not contained in L , then P_k is the only path from \mathcal{P} containing v , and hence $\ell_1 P_k \ell_2$ is the only path from $(\ell_1 P_n \ell_2)_{n \geq N}$ containing v . Otherwise v is contained in $L \setminus [\ell_1, \ell_2]_L$ so, say, $\ell_2 <_L v$. Then, as P_n with $n > k$ does (L, \mathcal{P}) -grain $V(P_k)$, the vertex ℓ_2 precedes v on P_n , giving $v \notin \ell_1 P_n \ell_2$ as desired.

(GL3) Consider any $n \geq N$ and write $L'_{<n} := [\ell_1, \ell_2]_L \cap \bigcup_{k=N}^{n-1} V(\ell_1 P_k \ell_2)$. By the already shown (GL1) we have $L'_{<n} \subseteq V(\ell_1 P_n \ell_2)$, so $\ell_1 P_n \ell_2$ does induce a linear ordering on $L'_{<n}$, and it coincides with the linear ordering induced by $[\ell_1, \ell_2]_L$ by (GL3) for (L, \mathcal{P}) .

Therefore, $([\ell_1, \ell_2]_L, (\ell_1 P_n \ell_2)_{n \geq N})$ is an ℓ_1 - ℓ_2 grain line; now we show that it is wildly presented if (L, \mathcal{P}) is. For this consider any $n \geq N$ with some two elements $\ell <_L \ell'$ of $L'_{<n}$. Then, as (L, \mathcal{P}) is wildly presented and $L'_{<n} \subseteq L_{<n}$, the subpath $\ell P_n \ell'$ of $\ell_1 P_n \ell_2$ has a vertex in $(\ell, \ell')_L$. \square

13.6. Proof of the main result

In this section, we employ our results on grain lines to prove Theorem 13.6.1, which we then use to prove Theorems 13.1 and 13.2.

Theorem 13.6.1. *Suppose that G is any subdivided Π -graph and that u, v are two distinct branch vertices of G . Then there exist subgraphs $H_u, H_v \subseteq G$ that satisfy the following conditions:*

- (i) $H_u[X] = H_v[X]$ is finite and connected for $X := V(H_u) \cap V(H_v) \neq \emptyset$;
- (ii) X avoids u and v ;
- (iii) both H_u/X and H_v/X are subdivided Π -graphs in which u, X and v, X are branch vertices, respectively;
- (iv) uX is an edge of H_u/X and vX is an edge of H_v/X .

Proof. Without loss of generality we may assume that uv is not an edge of G . Using that G is a subdivided Π -graph we find a finite vertex set $S \subseteq V(G) \setminus \{u, v\}$ that separates u and v in G . We write C_u and C_v for the two distinct components of $G - S$ that contain u and v respectively. Next, we use Theorem 13.5.4 and Lemma 13.5.5 to find a wildly presented u - v grain line (L, \mathcal{P}) in G . Without loss of generality we may assume that P_0 grains the finite vertex set S . We let s_u be the first vertex of the u - v path P_0 in S , and we let s_v be the last vertex of P_0 in S . That is to say that s_u and s_v are the least and greatest vertex of L in S . Then, for all $n \in \mathbb{N}$, the paths $uP_n s_u$ and $s_v P_n v$ are contained in $G[C_u + s_u]$ and $G[s_v + C_v]$ respectively.

Next, we let x_u and x_v be the least and greatest vertex of L in $V(\mathring{P}_0)$. Moreover, we let $L_u := [u, x_u]_L$ and $\mathcal{P}_u := (uP_n x_u)_{n \geq 1}$, and we let $L_v := [x_v, v]_L$ and $\mathcal{P}_v := (x_v P_n v)_{n \geq 1}$. Then (L_u, \mathcal{P}_u) and (L_v, \mathcal{P}_v) are wildly presented u - x_u and x_v - v grain lines in G by Lemma 13.5.6. We claim that $H_u := P_0 \overset{\circ}{\cup} \bigcup \mathcal{P}_u$ and $H_v := \overset{\circ}{P}_0 \cup \bigcup \mathcal{P}_v$ are the desired subgraphs.

First, we show that $X = V(\mathring{P}_0)$ and that X satisfies (i), (ii) and (iv). For this, it suffices to show that for every $n \geq 1$ the paths $uP_n x_u$ and $x_v P_n v$ are u - \mathring{P}_0 and \mathring{P}_0 - v paths in $G[C_u + s_u]$ and $G[s_v + C_v]$, respectively. The vertex $s_u \in L \cap S \subseteq L \cap V(\mathring{P}_0)$ was a candidate for x_u , implying $x_u \leq_L s_u$, and then for all $n \geq 1$ the path P_n graining $V(P_0)$ gives $uP_n x_u \subseteq uP_n s_u \subseteq G[C_u + s_u]$ on the one hand and that x_u is the first vertex of P_n in \mathring{P}_0 on the other hand; for the paths $x_v P_n v$ we employ symmetry.

(iii) follows from the facts that (L_u, \mathcal{P}_u) and (L_v, \mathcal{P}_v) are wildly presented and that all paths $uP_n x_u$ and $x_v P_n v$ ($n \geq 1$) are u - \mathring{P}_0 and \mathring{P}_0 - v paths respectively. \square

Now we have almost all we need to prove Theorems 13.1 and 13.2. In the proof of Theorem 13.2, we will face the construction of a minor with finite branch sets in countably many steps. The following notation and lemma will help us to keep the technical side of this construction to the minimum.

Suppose that G and H are two graphs with H a minor of G . Then there are a vertex set $U \subseteq V(G)$ and a surjection $f: U \rightarrow V(H)$ such that the preimages $f^{-1}(x) \subseteq U$ form the branch sets of a model of H in G . A *minor-map* $\varphi: G \succcurlyeq H$

13. The Farey graph is uniquely determined by its connectivity

formally is such a pair (U, f) . Given $\varphi = (U, f)$ we address U as $V(\varphi)$ and we write $\varphi = f$ by abuse of notation. Usually, we will abbreviate ‘minor-map’ as ‘map’. The proof of Lemma 12.5.12 shows:

Lemma 13.6.2. *Let G_0, G_1, \dots and $H_0 \subseteq H_1 \subseteq \dots$ be two sequences of graphs $H_n \subseteq G_n$ with maps $\varphi_n: G_n \succcurlyeq G_{n+1}$ such that for every vertex $x \in G_{n+1}$ the preimage $\varphi_n^{-1}(x)$ is finite if $x \notin H_n$ and equal to $\{x\}$ if $x \in H_n$. Then G contains $\bigcup_{n \in \mathbb{N}} H_n$ as a minor with finite branch sets.* \square

Proof of Theorem 13.2. Let G be any Π -graph. We have to find a Farey graph minor in G with finite branch sets. By Lemma 13.2.1 it suffices to find a halved Farey graph minor with finite branch sets in G .

Call a graph a *foresighted* halved Farey graph of order $n \in \mathbb{N}$ if it is the edge-disjoint union of \check{F}_n with subdivided Π -graphs A_{uv} , one for every blue edge $uv \in \check{F}_n$, such that:

- each A_{uv} meets \check{F}_n precisely in u and v but $uv \notin A_{uv}$;
- u and v are branch vertices of A_{uv} ;
- every two distinct A_e and $A_{e'}$ meet precisely in the intersection $e \cap e'$ of their corresponding edges (viewed as vertex sets).

To find a halved Farey graph minor with finite branch sets in G , it suffices by Lemma 13.6.2 to find a sequence $G =: H_0, H_1, \dots$ of foresighted halved Farey graphs of orders $0, 1, \dots$ with maps $\varphi_n: H_n \succcurlyeq H_{n+1}$ such that $\varphi_n^{-1}(x)$ is finite for all $x \in H_{n+1} - \check{F}_n$ and $\varphi_n^{-1}(x) = \{x\}$ for all $x \in \check{F}_n$.

To get started, pick any edge e of G , and note that $G = H_0$ is a foresighted halved Farey graph of order 0 with $A_e = G - e$ when we rename e to the edge of which $\check{F}_0 = K^2$ consists.

At step $n + 1$ suppose that we have already constructed $H_n \supseteq \check{F}_n$ and consider the subdivided Π -graphs A_e that were added to \check{F}_n to form H_n . Theorem 13.6.1 yields in each A_e two subgraphs H_u^e, H_v^e for $e = uv$ that satisfy the following conditions:

- (i) $H_u^e[X^e] = H_v^e[X^e]$ is finite and connected for $X^e := V(H_u^e) \cap V(H_v^e) \neq \emptyset$;
- (ii) X^e avoids u and v ;
- (iii) both H_u^e/X^e and H_v^e/X^e are subdivided Π -graphs in which u, X^e and v, X^e are branch vertices, respectively;
- (iv) uX^e is an edge of H_u^e/X^e and vX^e is an edge of H_v^e/X^e .

Then we let $A_{uv_e} := H_u^e/X^e$ and $A_{v_e v} := H_v^e/X^e$ for every blue edge $uv \in \check{F}_n$, where we recall that v_e is the vertex $v_e \in \check{F}_{n+1} - \check{F}_n$ that arises from $uv \in \check{F}_n$ in the recursive definition of \check{F}_{n+1} . After renaming the vertex X^e to v_e in both A_{uv_e} and $A_{v_e v}$, we let

$$H_{n+1} := \check{F}_{n+1} \cup \bigcup \{ A_f \mid f \in \check{F}_{n+1} \text{ is a blue edge} \}$$

$$V(\varphi_n) := V(\check{F}_n) \cup \bigcup \{ V(H_u^e) \cup V(H_v^e) \mid e = uv \in \check{F}_n \text{ is a blue edge} \}$$

and we let $\varphi_n: V(\varphi_n) \rightarrow V(H_{n+1})$ send w to v_e if $w \in X^e$ for some blue edge $e \in \check{F}_n$ and $\varphi_n(w) := w$ otherwise. This completes the proof. \square

Proof of Theorem 13.1. Theorem 13.2 implies Theorem 13.1. \square

14. Ubiquity and the Farey graph




Figure 14.0.1.: The whirl graph, coloured

14.1. Introduction

“One of the most basic problems in an infinite setting that has no finite equivalent is whether or not ‘arbitrarily many’, in some context, implies ‘infinitely many’.” (Diestel [26]). For example, Halin [26, 46] proved that if a graph contains k disjoint rays for every integer k , then it contains infinitely many disjoint rays. Substructures of a given type—subgraphs, minors, rooted minors or whatever—of which there must exist infinitely many disjoint copies (for some notion of disjointness) in a given graph as soon as there are arbitrarily (finitely) many such copies are called *ubiquitous* [26]. Examples of ubiquity results can be found in [2–5, 7–10, 26, 42, 46, 64, 85].

Usually, ubiquity problems are trivial as soon as the substructures considered are finite. For example, if a graph G contains k disjoint u – v paths for every integer k and some fixed vertices u and v , we can greedily find infinitely many disjoint u – v paths in G . Similarly, edge-disjoint paths between two fixed vertices are clearly ubiquitous. Interestingly, this changes as soon as we require our edge-disjoint paths to traverse their common vertices in the same order.

Let us call two u – v paths *order-compatible* if they traverse their common vertices in the same order. Our first aim in this chapter is to show that edge-disjoint order-compatible paths between two given vertices are not ubiquitous: we shall construct a graph G , the whirl graph shown in Figure 14.0.1, that has two vertices u and v such that G contains k edge-disjoint order-compatible u – v paths for every integer k , but not infinitely many. In fact, the whirl graph G will have this property for *all* pairs of vertices:

Theorem 14.1. *The whirl graph is a countable planar graph that contains k edge-disjoint pairwise order-compatible paths between every two of its vertices for every $k \in \mathbb{N}$, but which does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices.*

Applications

Our result has two applications.




Figure 14.1.1.: The Farey graph




Figure 14.1.2.: The graph $T_{\aleph_0} * t$

The Farey graph, shown in Figure 14.1.1 and surveyed in [22, 49], plays a role in a number of mathematical fields ranging from group theory and number theory to geometry and dynamics [22]. Curiously, graph theory has not been among these until very recently, when it was shown that the Farey graph plays a central role in graph theory too: it is one of two infinitely edge-connected graphs that must occur unavoidably as a minor in every infinitely edge-connected graph. The second graph is $T_{\aleph_0} * t$, the graph obtained from the infinitely-branching tree T_{\aleph_0} by joining an additional vertex t to all its vertices; see Figure 14.1.2.

Theorem 12.1. *Every infinitely edge-connected graph contains either the Farey graph or $T_{\aleph_0} * t$ as a minor.*

This result lies in the intersection of Ramsey theory and the study of connectivity; see the introduction of Chapter 12. Related results can be found in [26, 38, 40, 45, 51, 65]; see [26, §9.4] or the introduction of [40] for surveys.

The obvious question this theorem raises is whether it is best possible in the sense that one cannot replace ‘minor’ with ‘topological minor’ in its wording. The whirl graph and Theorem 14.1 are needed in Chapter 12 to answer this question in the affirmative:

Theorem 12.3.3. *The whirl graph is infinitely edge-connected but contains neither the Farey graph nor $T_{\aleph_0} * t$ as a topological minor.*

The second application of the whirl graph and Theorem 14.1 concerns the first graph-theoretic characterisation of the Farey graph. Very recently it was shown that the Farey graph is uniquely determined by its connectivity: up to minor-equivalence, the Farey graph is the unique minor-minimal graph that is infinitely edge-connected but such that every two vertices can be finitely separated. A Π -graph is an infinitely edge-connected graph that does not contain infinitely many independent paths between any two of its vertices. A Π -graph is *typical* if

14. Ubiquity and the Farey graph

it occurs as a minor in every Π -graph. Note that any two typical Π -graphs are minors of each other; we call such graphs *minor-equivalent*.

Theorem 13.1. *Up to minor-equivalence, the Farey graph is the unique typical Π -graph.*

This theorem continues to hold if we require all minors to be tight: A *tight* minor is a minor with finite branch sets. A Π -graph is *tightly* typical if it occurs as a tight minor in every Π -graph. Note that any two tightly typical Π -graphs are tight minors of each other; we call such graphs *tightly minor-equivalent*.

Theorem 13.3.2. *Up to tight minor-equivalence, the Farey graph is the unique tightly typical Π -graph.*

The obvious question this theorem raises is whether it is best possible in the sense that one cannot replace ‘tight’ with ‘topological’. The whirl graph and Theorem 14.1 are needed in Chapter 13 to answer this question in the affirmative:

Theorem 13.3.3. *The whirl graph is a Π -graph that contains the Farey graph as a tight minor but not as a topological minor.*

This theorem in turn raises the two questions how exactly the Farey graph is contained in the whirl graph as a minor and how large the branch sets actually are. We shall use the Cantor set to explicitly determine a Farey graph minor in the whirl graph with branch sets of size two; see Section 14.3 for the explicit description of the Farey graph minor.

Theorem 14.2. *The whirl graph contains the Farey graph as a minor with branch sets of size two, but not as a topological minor.*

This chapter is organised as follows. We introduce the whirl graph in Section 14.2 where we also prove Theorem 14.1, and we prove Theorem 14.2 in Section 14.3.

14.2. Proof of Theorem 1

The *whirl graph*, shown in Figure 14.0.1, is the graph $G = (V, E)$ on $V := \bigcup_{n=1}^{\infty} V_n$ where $V_n := \left\{ \frac{0}{3^n}, \frac{1}{3^n}, \dots, \frac{3^n}{3^n} \right\}$ and with edge set $E := \bigcup_{n=1}^{\infty} E_n$ where

$$E_n := \left\{ \left\{ \frac{3k}{3^n}, \frac{3k+2}{3^n} \right\}, \left\{ \frac{3k+1}{3^n}, \frac{3k+2}{3^n} \right\}, \left\{ \frac{3k+1}{3^n}, \frac{3k+3}{3^n} \right\} \mid k \in \{0, 1, \dots, 3^{n-1} - 1\} \right\}.$$

For every integer $n \geq 1$ we define the three subgraphs

$$G_{\leq n} := (V_n, \bigcup_{k=1}^n E_k) \quad \text{and} \quad G_n := (V_n, E_n) \quad \text{and} \quad G_{\geq n} := (V, \bigcup_{k=n}^{\infty} E_k);$$

see Figure 14.2.1 for an illustration. Note that G_n is a Hamilton path of $G_{\leq n}$ for all n .

For the proof of Theorem 14.1 we need another theorem and a lemma. At the end of one of my talks at Hamburg that involved order-compatible paths, Joshua Erde

14. Ubiquity and the Farey graph




Figure 14.2.1.: $G_1 = G_{\leq 1}$ is black, $G_{\leq 2}$ is the union of black and red, G_2 is red, $G_{\geq 2}$ is the union of red and blue, and $G_{\geq 3}$ is blue

asked: *Is there a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that, for every graph H and every two vertices u and v of H , the existence of at least $f(k)$ many edge-disjoint u - v paths in H implies the existence of k many edge-disjoint pairwise order-compatible u - v paths in H ?* The next day, Jakob Kneip answered the question in the affirmative for $f(k) = k$ the identity on \mathbb{N} :

Theorem 14.2.1 (Kneip). *Let H be any graph, let u and v be any two distinct vertices of H , and let n be any natural number. If H contains n edge-disjoint u - v paths, then H also contains n edge-disjoint pairwise order-compatible u - v paths.*

Proof. Given H, u, v, n we suppose that H contains n edge-disjoint u - v paths. Choose a path-system \mathcal{P} of n edge-disjoint u - v paths in H that uses as few edges of H as possible. Then the paths in \mathcal{P} are pairwise order-compatible: For this, assume for a contradiction that P and Q are paths in \mathcal{P} such that P traverses two vertices x and y as $x <_P y$ while Q traverses them as $y <_Q x$. Then $uPx \cup xQv$ and $uQy \cup yPv$ are connected edge-disjoint subgraphs of $P \cup Q$, so we may choose one u - v path in each of the two. Now replacing P and Q with these two new paths yields a system of n edge-disjoint u - v paths using strictly fewer edges of H than \mathcal{P} , since the edges of xPy and yQx are not used by the new paths (contradiction). \square

Lemma 14.2.2. *Let $u, v \in V$ be any two vertices with $u <_{\mathbb{Q}} v$ and let $n > 1$ be any integer with $u, v \in V_{n-1}$. If $P \subseteq G_{\geq n}$ is any u - v path, then*

$$V_{n-1} \cap [u, v] \subseteq V(P) \subseteq V \cap [u, v]$$

and P traverses the vertices in $V_{n-1} \cap [u, v]$ in the natural order induced by \mathbb{Q} .

Proof. Every vertex $x \in V_{n-1} \setminus \{0, 1\}$ is a cutvertex of $G_{\geq n}$ and the components of $G_{\geq n} - x$ are $G_{\geq n}[V \cap [0, x]]$ and $G_{\geq n}[V \cap (x, 1]]$. This clearly implies the statement of the lemma. \square

Now we prove Theorem 14.1:

14. Ubiquity and the Farey graph

Proof of Theorem 14.1. Clearly, G is planar. It is infinitely edge-connected because it can be written as the edge-disjoint union $\bigcup_{n \in \mathbb{N}} G_n = G$. In particular, it follows from Theorem 14.2.1 that G contains k edge-disjoint pairwise order-compatible paths between any two vertices, for every $k \in \mathbb{N}$.

It remains to show that G does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two vertices. For this, let any two vertices u and v of G be given, say with $u <_{\mathbb{Q}} v$. We pick any integer $N > 1$ such that $u, v \in V_{N-1}$. Since all the edge sets E_0, E_1, \dots are finite, it suffices to show the following assertion:

Whenever P is any u - v path in $G_{\geq N}$ and $M \geq N$ is the minimal integer such that $G_{\leq M}$ contains P , no u - v path in $G_{\geq M+1}$ is order-compatible with P .

Let P and M be given. By the minimality of M the path P must have an edge e in E_M . Let x and y be the two consecutive elements of $V_{M-1} \subseteq \mathbb{Q}$ bounding an interval $[x, y]$ that contains the endvertices of e . We claim that P contains the subpath xG_My of the 0-1 Hamilton path G_M of $G_{\leq M}$.

Indeed, on the one hand the edge e lies on P , so P has a vertex in $V_M \cap (x, y)$. On the other hand, the separator of the separation $\{V_M \setminus (x, y), V_M \cap [x, y]\}$ of $G_{\leq M}$ is $\{x, y\}$ while $u, v \in V_{M-1} \subseteq V_M \setminus (x, y)$ and $G_{\leq M}[V_M \cap [x, y]] = xG_My$. Thus, the u - v path P meeting $V_M \cap (x, y)$ implies that P contains both vertices x and y and that either $xPy = xG_My$ or yPx is the reverse of xG_My . In either case, P contains xG_My .

Now let Q be any u - v path in $G_{\geq M+1}$. We show that Q is not order-compatible with P . For this, we consider the path xG_My that is contained in P . We apply Lemma 14.2.2 twice: First, we apply it to u, v, N and P to establish $V(P) \subseteq V \cap [u, v]$ which ensures $u \leq x < y \leq v$ in \mathbb{Q} . And the second time we apply it to $u, v, M+1$ and Q to establish $V_M \cap [u, v] \subseteq V(Q)$ and that Q traverses the vertices in $V_M \cap [u, v]$ in the natural order induced by \mathbb{Q} . Altogether, we deduce that Q traverses the vertices in $V(xG_My) \subseteq V_M \cap [x, y] \subseteq V_M \cap [u, v]$ in the natural order induced by \mathbb{Q} . Since P contains xG_My and xG_My is the path

$$\frac{3k}{3^M} \frac{3k+2}{3^M} \frac{3k+1}{3^M} \frac{3k+3}{3^M}$$

for the appropriate integer k , the paths P and Q certainly are not order-compatible, completing the proof. □

14.3. Finding the Farey graph in the whirl graph

For the definition of the Farey graph, see Chapter 12. We have shown in Chapter 13 that any graph contains the Farey graph as a minor with finite branch sets if it is infinitely edge-connected and does not contain infinitely many independent paths between any two vertices. As independent paths are order-compatible, it follows that the whirl graph contains the Farey graph as a minor with finite branch sets. The result in Chapter 13, however, does not provide an explicit description of the

14. Ubiquity and the Farey graph

Farey graph minor in the whirl graph, nor does it tell us how large the branch sets actually are. In our particular situation, the latter is especially unsatisfactory, as we already know that infinitely many branch sets must be non-trivial because the whirl graph does not contain the Farey graph as a topological minor. That is why in this section we use the Cantor set to explicitly determine a Farey graph minor in the whirl graph with branch sets of size two.

Recall that the Cantor set is $C := \bigcap_{n=0}^{\infty} \bigcup \mathcal{C}_n$ where $\mathcal{C}_0 := \{[0, 1]\}$ and \mathcal{C}_{n+1} is obtained from \mathcal{C}_n by replacing each interval $[a, a + \Delta] \in \mathcal{C}_n$ with the two intervals $[a, a + \frac{1}{3}\Delta]$ and $[a + \frac{2}{3}\Delta, a + \Delta]$.

We define the subgraph $G^* := (C^*, E^*) \subseteq G$ on $C^* := \bigcup_{n=1}^{\infty} C_n^*$ and with edge set $E^* := \bigcup_{n=1}^{\infty} E_n^*$ where

$$\begin{aligned} C_n^* &:= \left\{ a, a + \frac{1}{3}\Delta, a + \frac{2}{3}\Delta, a + \Delta \mid [a, a + \Delta] \in \mathcal{C}_{n-1} \right\} \\ &= \{x, y \mid [x, y] \in \mathcal{C}_n\} = V_n \cap C \text{ and} \\ E_n^* &:= \left\{ \{a, a + \frac{2}{3}\Delta\}, \{a + \frac{1}{3}\Delta, a + \frac{2}{3}\Delta\}, \{a + \frac{1}{3}\Delta, a + \Delta\} \mid [a, a + \Delta] \in \mathcal{C}_{n-1} \right\} \subseteq E_n. \end{aligned}$$

We shall find the halved Farey graph (minus one edge) as a contraction minor of G^* . For this, we write $G_{\leq n}^* := (C_n^*, \bigcup_{k=1}^n E_k^*)$ and $M := \bigcup_{n=1}^{\infty} M_n$ where

$$M_n := \left\{ \{a + \frac{1}{3}\Delta, a + \frac{2}{3}\Delta\} \mid [a, a + \Delta] \in \mathcal{C}_{n-1} \right\} \subseteq E_n^*;$$

see Figure 14.3.1 for an illustration. We write $M_{\leq n} := \bigcup_{k=1}^n M_k$. If D is an independent set of edges and H is any graph, then we denote by H/D the contraction minor of H obtained by contracting the edges in $D \cap E(H)$.




Figure 14.3.1.: On the left: The red edges form M , and together with the blue edges they form G^* . On the right: $\check{F} - E(\check{F}_0)$ with blue edge set and red vertex set.

Lemma 14.3.1. *There exists an isomorphism $G^*/M \cong \check{F} - E(\check{F}_0)$ that associates 0 and 1 with the two vertices of \check{F}_0 .*

Proof. Let x and y be the two endvertices of \check{F}_0 .

On the one hand, for every $n \in \mathbb{N}$ the x - y Hamilton path of \check{F}_n formed by the blue edges of \check{F}_n induces a linear ordering on the vertex set of \check{F}_n in which x is the

14. Ubiquity and the Farey graph

least element and y is the greatest, and these orderings are compatible for distinct numbers n .

On the other hand, for every integer $n \geq 1$ the vertex set $\{0, 1\} \cup M_{\leq n}$ of $G_{\leq n}^*/M$ inherits the linear ordering \leq_n from \mathbb{Q} in which $0 <_n e <_n 1$ for all $e \in M_{\leq n}$ and $\{u, v\} <_n \{s, t\}$ if and only if $\min\{u, v\} <_{\mathbb{Q}} \min\{s, t\}$ for all $uv \neq st \in M_{\leq n}$, and again these linear orderings are compatible for distinct numbers n .

An induction on $n \geq 1$ shows that the unique order-isomorphism φ_n between the linearly ordered finite vertex sets of $G_{\leq n}^*/M$ and $\check{F}_n - E(\check{F}_0)$ is a graph-isomorphism such that $\varphi_1 \subseteq \dots \subseteq \varphi_n$. Then the ascending union $\bigcup_{n=1}^{\infty} \varphi_n$ of these isomorphisms is the desired graph-isomorphism between G^*/M and $\check{F} - E(\check{F}_0)$ that associates 0 and 1 with the two vertices of \check{F}_0 . \square

In order to find a Farey graph minor in G , we must find two halved Farey graph minors in G . For this, we consider copies of G^* in G that arise by linear transformation. Every permutation π of \mathbb{Q} acts on both the set of graphs H with $V(H) \subseteq \mathbb{Q}$ and the set of edge sets D with $D \subseteq [\mathbb{Q}]^2$ by renaming every vertex v to $\pi(v)$. Then we write πH and πD for the resulting graph and edge set. Now let us consider the two permutations $\pi_1(x) := (1/9)x + 3/9$ and $\pi_2(x) := (1/9)x + 5/9$. These send G^* to copies $\pi_1 G^*$ and $\pi_2 G^*$ of G^* that are subgraphs of G . By Lemma 14.3.1 we have $\pi_1 G^*/\pi_1 M \cong \check{F} - E(\check{F}_0)$ and $\pi_2 G^*/\pi_2 M \cong \check{F} - E(\check{F}_0)$ by isomorphisms that associate the vertices $3/9, 4/9$ and $5/9, 6/9$ with the two vertices of \check{F}_0 . Joining these two halved Farey graph minors appropriately yields the desired Farey graph minor, as shown in Figure 14.3.2:




Figure 14.3.2.: This is $G[V \cap [3/9, 6/9]]$. The two subgraphs $\pi_1 G^*$ and $\pi_2 G^*$ are drawn using both red and blue, like in Figure 14.3.1. Theorem 14.2 states that the Farey graph arises from the subgraph consisting of the coloured edges by contracting red and orange while keeping blue and cyan.

14. Ubiquity and the Farey graph

Theorem 14.2. *The whirl graph contains the Farey graph as a minor with branch sets of size two:*

$$F \cong \left(\pi_1 G^* \cup \pi_2 G^* + \left\{ \frac{3}{9}, \frac{5}{9} \right\} + \left\{ \frac{4}{9}, \frac{6}{9} \right\} + \left\{ \frac{3}{9}, \frac{6}{9} \right\} \right) / \left(\pi_1 M \cup \pi_2 M + \left\{ \frac{3}{9}, \frac{5}{9} \right\} + \left\{ \frac{4}{9}, \frac{6}{9} \right\} \right)$$

where $\pi_1 G^* \cup \pi_2 G^* + \left\{ \frac{3}{9}, \frac{5}{9} \right\} + \left\{ \frac{4}{9}, \frac{6}{9} \right\} + \left\{ \frac{3}{9}, \frac{6}{9} \right\} \subseteq G$. But the whirl graph does not contain the Farey graph as a topological minor. □

Appendix

15. English summary

Part I:

Chapter 3. We show that the tangle space of a graph, which compactifies it, is a quotient of its Stone-Ćech remainder obtained by contracting the connected components.

Chapter 4. Carmesin has extended Robertson and Seymour’s tree-of-tangles theorem to the infinite-order tangles of locally finite infinite graphs. We extend it further to the infinite-order tangles of all infinite graphs.

Our result has a number of applications for the topology of infinite graphs, such as their end spaces and their compactifications.

Part II:

Chapter 5. Extending the well-known star-comb lemma for infinite graphs, we characterise the graphs that do not contain an infinite comb or an infinite star, respectively, attached to a given set of vertices.

We offer several characterisations: in terms of normal trees, tree-decompositions, ranks of rayless graphs and tangle-distinguishing separators.

Chapter 6. In a series of four chapters we determine structures whose existence is dual, in the sense of complementary, to the existence of stars or combs. Here, in the second chapter of the series, we present duality theorems for combinations of stars and combs: 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 or tree-decompositions.

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.

Chapter 7. In a series of four chapters we determine structures whose existence is dual, in the sense of complementary, to the existence of stars or combs. Here, in the third chapter of the series, we present duality theorems for a combination of stars and combs: undominated combs. We describe their complementary structures in terms of rayless trees and of tree-decompositions.

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. Only two such graphs had so far been constructed, by Seymour and Thomas [76] and by Thomassen [83]. 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.

Another application settles a problem left unsolved by Carmesin [19]: The

graphs whose undominated ends are reflected by a suitable spanning tree can be characterised in terms of normal spanning trees. In particular, we show that every graph that has a normal spanning tree does have a spanning tree reflecting its undominated ends.

Chapter 8. In a series of four chapters we determine structures whose existence is dual, in the sense of complementary, to the existence of stars or combs. In the first chapter of our series we determined structures that are complementary to arbitrary stars or combs. Stars and combs can be combined, positively as well as negatively. In the second and third chapter of our series we provided duality theorems for all but one of the possible combinations.

In this fourth and final chapter of our series, we complete our solution to the problem of finding complementary structures for stars, combs, and their combinations, by presenting duality theorems for the missing piece: for undominating stars. Our duality theorems are phrased in terms of end-compactified subgraphs, tree-decompositions and tangle-distinguishing separators.

Chapter 9. Schmidt 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 asked whether Schmidt's rank function can be generalised to characterise other important classes of graphs. In this chapter 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 .

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. Investigating the class of normally traceable graphs further, we prove that all its graphs have spanning trees reflecting their undominated ends. Our proofs rely on a characterisation of the class of normally traceable graphs by an ordinal rank function that we provide.

Part III:

Chapter 10. We show that every connected graph can be approximated by a normal tree, up to some arbitrarily small error phrased in terms of neighbourhoods around its ends. The existence of such approximate normal trees has consequences of both combinatorial and topological nature.

On the combinatorial side, we show that a graph has a normal spanning tree as soon as it has normal spanning trees locally at each end; i.e., the only obstruction for a graph to having a normal spanning tree is an end for which none of its neighbourhoods has a normal spanning tree.

On the topological side, we show that the end space $\Omega(G)$, as well as the spaces

$|G| = G \cup \Omega(G)$ naturally associated with a graph G , are always paracompact. This gives unified and short proofs for a number of results by Diestel, Sprüssel and Polat, and answers an open question about metrizability of end spaces by Polat.

Chapter 11. The *directions* of an infinite graph G are a tangle-like description of its ends: they are choice functions that choose compatibly for all finite vertex sets $X \subseteq V(G)$ a component of $G - X$.

Although every direction is induced by a ray, there exist directions of graphs that are not uniquely determined by any countable subset of their choices. We characterise these directions and their countably determined counterparts in terms of star-like substructures or rays of the graph.

Curiously, there exist graphs whose directions are all countably determined but which cannot be distinguished all at once by countably many choices. We structurally characterise the graphs whose directions can be distinguished all at once by countably many choices, and we structurally characterise the graphs which admit no such countably many choices. Our characterisations are phrased in terms of normal trees and tree-decompositions.

Our four (sub)structural characterisations imply combinatorial characterisations of the four classes of infinite graphs that are defined by the first and second axiom of countability applied to their end spaces: the two classes of graphs whose end spaces are first countable or second countable, respectively, and the complements of these two classes.

Part IV:

Chapter 12. We show that every infinitely edge-connected graph contains the Farey graph or $T_{\aleph_0} * t$ as a minor. These two graphs are unique with this property up to minor-equivalence.

Chapter 13. We show that, up to minor-equivalence, the Farey graph is the unique minor-minimal graph that is infinitely edge-connected but such that every two vertices can be finitely separated.

Chapter 14. We construct a countable planar graph which, for any two vertices u, v and any integer $k \geq 1$, contains k edge-disjoint order-compatible $u-v$ paths but not infinitely many. This graph has applications in Ramsey theory, in the study of connectivity and in the characterisation of the Farey graph.

16. Deutsche Zusammenfassung

Part I:

Chapter 3. Wir zeigen, dass der Tangle-Raum eines Graphen G , welcher ihn kompaktifiziert, ein Quotient seines Stone-Čech Restes $(\beta G) \setminus G$ ist der durch Kontraktion der Zusammenhangskomponenten entsteht.

Chapter 4. Carmesin hat Robertson und Seymours Baum-von-Tangles Satz für Tangles unendlicher Ordnung von lokal endlichen unendlichen Graphen verallgemeinert.

Unser Resultat hat einige Anwendungen für die Topologie unendlicher Graphen, insbesondere für die Endenräume und Kompaktifizierungen unendlicher Graphen.

Part II:

Chapter 5. Das Stern-Kamm Lemma für unendliche Graphen erweiternd, charakterisieren wir die Graphen die keinen unendlichen Kamm oder unendlichen Stern an einer gegebenen Eckenmenge enthalten.

Wir präsentieren mehrere Charakterisierungen: durch normale Bäume, Baumzerlegungen, Ränge von strahlenlosen Graphen und Tangle-unterscheidende Trenner.

Chapter 6. In einer Serie von vier Kapiteln bestimmen wir Strukturen, deren Existenz dual ist, im Sinne von komplementär, zur Existenz von Sternen oder Kämmen. Hier, im zweiten Kapitel der Serie, präsentieren wir Dualitätssätze für Kombinationen von Sternen und Kämmen: dominierende Sterne und dominierte Kämmen. Da dominierende Sterne genau dann existieren, wenn dominierende Kämmen existieren, fallen ihre komplementären Strukturen zusammen. Wie für beliebige Sterne und Kämmen sind unsere Dualitätssätze für dominierte Kämmen (und dominierende Sterne) formuliert hinsichtlich normaler Bäume und Baumzerlegungen.

Die komplementären Strukturen für dominierte Kämmen, die wir bereitstellen, vereinigen jene für Sterne und Kämmen und erlauben uns, unsere Dualitätssätze für Sterne und Kämmen aus denen für dominierte Kämmen herzuleiten. Das ist überraschend vor dem Hintergrund, dass unsere komplementären Strukturen für Sterne und Kämmen recht verschieden sind: jene für Sterne sind lokal endlich wohingegen jene für Kämmen strahlenlos sind.

Chapter 7. In einer Serie von vier Kapiteln bestimmen wir Strukturen, deren Existenz dual ist, im Sinne von komplementär, zur Existenz von Sternen oder Kämmen. Hier, im dritten Kapitel der Serie, präsentieren wir Dualitätssätze für eine Kombination von Sternen und Kämmen: undominierte Kämmen. Wir beschreiben ihre komplementären Strukturen durch strahlenlose Bäume und durch Baumzerlegungen.

Anwendungen beinhalten eine vollständige Charakterisierung durch normale Bäume, von den Graphen, deren Strahlen dominiert sind, die aber keine strahlenlosen Bäume haben. Nur zwei solche Graphen wurden zuvor konstruiert, von Seymour und Thomas [76] und von Thomassen [83]. Als Korollar zeigen wir, dass

Graphen mit normalen Spann­bäumen genau dann einen strahlenlosen Spannbaum haben, wenn alle ihre Strahlen dominiert sind.

Eine weitere Anwendung löst ein Problem, das Carmesin [19] offen gelassen hat: Die Graphen, deren undominierte Enden durch einen geeigneten Spannbaum reflektiert werden, können durch normale Spann­bäume charakterisiert werden. Insbesondere zeigen wir, dass jeder Graph, der einen normalen Spannbaum hat, auch einen Spannbaum hat der seine undominierten Enden reflektiert.

Chapter 8. In einer Serie von vier Kapiteln bestimmen wir Strukturen, deren Existenz dual ist, im Sinne von komplementär, zur Existenz von Sternen oder Kämmen. Im ersten Kapitel unserer Serie haben wir Strukturen bestimmt, die komplementär sind zu beliebigen Sternen und Kämmen. Sterne und Kämmen können kombiniert werden, positiv wie negativ. Im zweiten und dritten Kapitel unserer Serie haben wir Dualitätssätze für alle bis auf eine der möglichen Kombinationen bereitgestellt.

In diesem vierten und finalen Kapitel unserer Serie vervollständigen wir unsere Lösung des Problems, komplementäre Strukturen für Sterne, Kämmen und ihre Kombinationen zu finden, indem wir Dualitätssätze für das letzte Puzzleteil finden: für undominierende Sterne. Unsere Dualitätssätze sind formuliert hinsichtlich Enden-kompaktifizierter Teilgraphen, Baumzerlegungen und Tangle-unterscheidender Trenner.

Chapter 9. Schmidt hat die Klasse der strahlenlosen Graphen durch eine ordinale Rangfunktion charakterisiert, was es ermöglicht, Aussagen über strahlenlose Graphen durch transfiniten Induktion zu beweisen. Halin hat gefragt, ob Schmidts Rangfunktion verallgemeinert werden kann, um andere wichtige Graphenklassen zu charakterisieren. In diesem Kapitel beantworten wir Halins Frage positiv: wir charakterisieren zwei wichtige Graphenklassen durch eine ordinale Rangfunktion.

Seymour und Thomas haben für jede überabzählbare Kardinalzahl κ die Klasse der Graphen ohne einen T_κ Minor charakterisiert. Wir erweitern ihre Charakterisierungen um eine ordinale Rangfunktion, eine für jede überabzählbare Kardinalzahl κ .

Ein weiteres weitgehend offenes Problem von Halin fragt nach einer Charakterisierung der Klasse der Graphen, die einen endentreuen Spannbaum haben. Eine gründlich studierte Teilklasse besteht aus den Graphen mit normalen Spann­bäumen. Wir bestimmen eine größere Teilklasse, die Klasse der normal verfolg­baren Graphen, welche aus den zusammenhängenden Graphen besteht, die eine strahlenlose Baumzerlegung in normal aufgespannte Teile besitzen. Die Klasse der normal verfolg­baren Graphen weiter studierend zeigen wir, dass alle ihre Graphen Spann­bäume besitzen, welche ihre undominierte Enden reflektieren. Unsere Beweise bauen auf eine Charakterisierung der Klasse der normal verfolg­baren Graphen durch eine ordinale Rangfunktion, die wir bereitstellen.

Part III:

Chapter 10. Wir zeigen, dass jeder zusammenhängende Graph durch einen

normalen Baum approximiert werden kann, bis auf einen beliebig kleinen Fehler formuliert hinsichtlich Nachbarschaften um seine Enden. Die Existenz von solchen approximierenden normalen Bäumen hat Konsequenzen von sowohl kombinatorischer als auch topologischer Natur.

Auf der kombinatorischen Seite zeigen wir, dass jeder Graph einen normalen Spannbaum hat, sobald er normale Spannbäume lokal an jedem seiner Enden hat; das heißt, das einzige Hindernis dafür, dass ein Graph einen normalen Spannbaum hat, ist ein Ende ohne normal aufgespannte Nachbarschaft.

Auf der topologischen Seite zeigen wir, dass jeder Endenraum $\Omega(G)$, als auch die Räume $|G| = G \cup \Omega(G)$ die natürlicherweise mit einem Graphen G assoziiert werden, immer parakompakt sind. Dies ermöglicht einheitliche und kurze Beweise von einigen Resultaten von Diestel, Sprüssel und Polat, und beantwortet eine offene Frage von Polat betreffend die Metrisierbarkeit von Endenräumen.

Chapter 11. Die *Richtungen* eines unendlichen Graphen G sind eine Tangle-artige Beschreibung seiner Enden: sie sind Auswahlfunktionen, die kompatibel für jede endliche Eckenmenge $X \subseteq V(G)$ eine Komponente von $G - X$ auswählen.

Obwohl jede Richtung von einem Strahl induziert ist, gibt es Richtungen von Graphen, die nicht eindeutig bestimmt sind durch eine abzählbare Teilmenge ihrer Entscheidungen. Wir charakterisieren diese Richtungen und ihre abzählbar bestimmten Gegenstücke durch Stern-artige Teilstrukturen oder Strahlen des Graphen.

Kurioserweise gibt es Graphen, deren Richtungen alle abzählbar bestimmt aber nicht alle gleichzeitig durch abzählbar viele Entscheidungen unterscheidbar sind. Wir charakterisieren strukturell die Graphen, deren Richtungen sich alle gleichzeitig durch abzählbar viele Entscheidungen unterscheiden lassen, und wir charakterisieren strukturell die Graphen, die keine solche abzählbar vielen Entscheidungen aufweisen. Unsere Charakterisierungen sind formuliert hinsichtlich normaler Bäume und Baumzerlegungen.

Unsere vier (teil)strukturellen Charakterisierungen implizieren kombinatorische Charakterisierungen der vier Klassen unendlicher Graphen, die definiert sind durch das erste und zweite Axiom der Abzählbarkeit, angewandt auf ihre Endenräume: die zwei Klassen von Graphen, deren Endenräume das erste oder zweite Abzählbarkeitsaxiom erfüllen, und die Komplemente dieser zwei Klassen.

Part IV:

Chapter 12. Wir zeigen, dass jeder unendlich kantenzusammenhängende Graph den Farey-Graphen oder $T_{\mathbb{N}_0} * t$ als Minor enthält. Diese zwei Graphen sind eindeutig mit dieser Eigenschaft bis auf Minoren-Äquivalenz.

Chapter 13. Wir zeigen, dass der Farey-Graph, bis auf Minoren-Äquivalenz, der eindeutige Minoren-minimale Graph ist, der unendlich kantenzusammenhängend ist aber sodass alle zwei Ecken endlich trennbar sind.

Chapter 14. Wir konstruieren einen abzählbaren planaren Graphen, welcher für alle zwei Ecken u, v und jede ganze Zahl $k \geq 1$ eine Anzahl von k kantendis-

16. Deutsche Zusammenfassung

junkten paarweise ordnungskompatiblen u - v Wegen enthält, aber nicht unendlich viele. Dieser Graph hat Anwendungen in der Ramseytheorie, im Studium des Zusammenhangs und in der Charakterisierung des Farey-Graphen.

17. Publications related to this dissertation

The following articles are related to this dissertation:

Part I:

- (i) Chapter 3 is based on [63].
- (ii) Chapter 4 is based on [35].

Part II:

- (iii) Chapter 5 is based on [14].
- (iv) Chapter 6 is based on [15].
- (v) Chapter 7 is based on [16].
- (vi) Chapter 8 is based on [17].
- (vii) Chapter 9 is based on [18].

Part III:

- (viii) Chapter 10 is based on [61].
- (ix) Chapter 11 is based on [60].

Part IV:

- (x) Chapter 12 is based on [56].
- (xi) Chapter 13 is based on [58].
- (xii) Chapter 14 is based on [59].

18. Declaration on my contributions

Part I:

Chapter 3. This chapter is based on the paper [63] that I wrote together with Max Pitz. We conducted the research in this chapter together. I drafted Sections 3.3 and 3.7, and I drafted half of each of the Sections 3.1, 3.2, 3.5 and 3.6.

Chapter 4. This chapter is based on the paper [35] that I wrote together with Ann-Kathrin Elm. We conducted the research in this chapter together. We drafted Sections 4.1 and 4.2 together. I drafted Section 4.3 except Section 4.3.1, and I drafted Sections 4.5 and 4.6. Of Section 4.8 I drafted the second half starting at Section 4.8.2. Of the figures, I drafted Figures 4.3.1, 4.3.2 and Figure 4.5.3.

Part II:

Chapters 5–8. This series of chapters is based on the four papers [14–17] that I wrote together with Carl Bürger. The series started as a single draft paper that I created, and on which [14] is based on. When I presented my draft in a talk, Carl had some exciting suggestions to extend it, so I invited him to join the project, and together we extended the first paper. Then Carl some day wondered about combinations of stars and combs, so we looked into those. Quickly we realised that these are exciting too, and in close collaboration we conducted a huge research effort of one and a half years that, ultimately, resulted in papers 2–4 of the series, and more additions to the first paper. One can say that I drafted most of the first and fourth paper, while Carl drafted most of the second and third paper.

Chapter 9. This chapter is based on the paper [18] that I wrote together with Carl after the star-comb project. We conducted the research in this chapter together. We drafted Sections 9.1 and 9.2 together. Of the remaining four sections, I drafted the first half, Sections 9.3 and 9.4.

Part III:

Chapter 10. This chapter is based on the paper [61] that I wrote together with Max Pitz and Ruben Melcher. We conducted the research in this chapter together. I drafted a third of the introduction and I drafted the final lemma. Furthermore, when an early draft of the proof of the main result had been created, I completed it.

Chapter 11. This chapter is based on the draft paper [60] that I am writing together with Ruben Melcher. We are conducting the research in this paper together. We drafted Sections 11.1 and 11.2 together. Ruben wrote an early draft for Section 11.3 and I wrote an early draft for Section 11.4. I polished both early drafts in order to include them in this dissertation before the publication of the paper.

18. *Declaration on my contributions*

Part IV:

Chapters 12–14. I created this entire series of three chapters on my own.

Acknowledgement

I am grateful to Reinhard Diestel for intensively supporting my research over the recent years to this day.

I am grateful to Max Pitz for introducing me to mathematical paper writing early and for pleasant collaborations.

Next, I would like to thank my co-author-officemates: Carl Bürger and Ruben Melcher. We had a great time together, and I dearly miss seeing you in person every day since that guy ate a bat.

I would like to thank Ann-Kathrin Elm for a nice collaboration that influenced some of my later work.

Also, I am grateful to Konstantinos Stavropoulos for stimulating conversations ;) And thank you for showing me a nice burger place!

Thank you to everyone in the DM group at Hamburg for creating a friendly and helpful working environment that makes all math problems solve themselves.

Lovely thanks to all my family and friends for all the fun times!

Bibliography

- [1] J.M. Aarts and T. Nishiura, *Dimension and extensions*, Vol. 48, Elsevier, 1993. MR1206002 ↑1.1.1, 3.1
- [2] T. Andreae, *Bemerkung zu einem Problem aus der Theorie der unendlichen Graphen*, Abh. Math. Sem. Univ. Hamburg **46** (1977), no. 1, 91, DOI 10.1007/BF02993015. MR0505832 ↑14.1
- [3] ———, *Classes of locally finite ubiquitous graphs*, J. Combin. Theory (Series B) **103** (2013), no. 2, 274–290, DOI 10.1016/j.jctb.2012.11.003. ↑14.1
- [4] ———, *On disjoint configurations in infinite graphs*, J. Graph Theory **39** (2002), no. 4, 222–229, DOI 10.1002/jgt.10016. ↑14.1
- [5] ———, *Über eine Eigenschaft lokalfiniter, unendlicher Bäume*, J. Combin. Theory (Series B) **27** (1979), no. 2, 202–215, DOI 10.1016/0095-8956(79)90082-0. MR0546863 ↑14.1
- [6] Nathan Bowler. Hamburg workshop on abstract separation systems, Spiekeroog 2014. ↑4.3.1
- [7] N. Bowler and J. Carmesin and J. Pott, *Edge-disjoint double rays in infinite graphs: A Halin type result*, J. Combin. Theory (Series B) **111** (2015), 1–16, DOI 10.1016/j.jctb.2014.08.005, available at arXiv:1307.0992. MR3315597 ↑14.1
- [8] N. Bowler and C. Elbracht and J. Erde and P. Gollin and K. Heuer and P. Pitz and M. Teegen, *Ubiquity in graphs I: Topological ubiquity of trees* (2018), available at arXiv:1806.04008. Submitted. ↑14.1
- [9] ———, *Ubiquity in graphs II: Ubiquity of graphs with non-linear end structure* (2018), available at arXiv:1809.00602. Submitted. ↑14.1
- [10] ———, *Ubiquity in graphs III: Ubiquity of graphs with extensive tree decompositions* (In preparation). ↑14.1
- [11] N. Bowler and G. Geschke and M. Pitz, *Minimal obstructions for normal spanning trees*, Fund. Math. **241** (2018), 245–263, DOI 10.4064/fm337-10-2017. MR3778904 ↑1.3.1, 10.1
- [12] H. Bruhn and R. Diestel, *Duality in infinite graphs*, Comb., Probab. & Comput. **15** (2006), 75–90, DOI 10.1017/S0963548305007261. MR2195576 ↑7.1, 7.1, 7.3, 7.3, 12.5.1
- [13] H. Bruhn and R. Diestel and A. Georgakopoulos and P. Sprüssel, *Every rayless graph has an unfriendly partition*, Combinatorica **30** (2010), no. 5, 521–532, DOI 10.1007/s00493-010-2590-3. MR2776717 ↑5.1, 5.3.4, 9.1
- [14] C. Bürger and J. Kurkofka, *Duality theorems for stars and combs I: Arbitrary stars and combs* (2020), available at arXiv:2004.00594. Submitted. ↑(iii), 18
- [15] ———, *Duality theorems for stars and combs II: Dominating stars and dominated combs* (2020), available at arXiv:2004.00593. Submitted. ↑(iv), 18
- [16] ———, *Duality theorems for stars and combs III: Undominated combs* (2020), available at arXiv:2004.00592. Submitted. ↑(v), 18
- [17] ———, *Duality theorems for stars and combs IV: Undominating stars* (2020), available at arXiv:2004.00591. Submitted. ↑(vi), 18
- [18] ———, *End-faithful spanning trees in graphs without normal spanning trees* (2020), available at arXiv:2006.01071. Submitted. ↑(vii), 18
- [19] J. Carmesin, *All graphs have tree-decompositions displaying their topological ends*, Combinatorica **39** (2019), no. 3, 545–596, DOI 10.1007/s00493-018-3572-0. MR3989261 ↑1.1.2, 4.1, 4.2.3, 4.2.3, 4.2.11, 4.7, 5.2.2, 5.2.7, 5.3.5, 6.1, 6.3, 6.3.2, 7.1, 7.3, 8.3, 9.1, 15, 16
- [20] J. Carmesin and M. Hamann and B. Miraftab, *Canonical trees of tree-decompositions* (2020), available at arXiv:2002.12030. Submitted. ↑4.2.3, 11.2.2, 11.2.4, 11.2.2

BIBLIOGRAPHY

- [21] D.I. Cartwright and P.M. Soardi and W. Woess, *Martin and end compactifications for non-locally-finite graphs*, Trans. Am. Math. Soc. **338** (1993), 679–693, DOI [10.2307/2154423](https://doi.org/10.2307/2154423). MR1102885 ↑[1.1.1](#), [3.1](#)
- [22] M. Clay and D. Margalit, *Office Hours with a Geometric Group Theorist*, Princeton University Press, 2017. MR3645425 ↑[1.4](#), [12.1](#), [12.2](#), [13.1](#), [14.1](#)
- [23] R. Diestel, *Abstract Separation Systems*, Order **35** (2018), no. 1, 157–170, DOI [10.1007/s11083-017-9424-5](https://doi.org/10.1007/s11083-017-9424-5), available at [arXiv:1406.3797v6](https://arxiv.org/abs/1406.3797v6). MR3774512 ↑[2.3](#), [1](#), [4.2.1](#), [4.3.1](#)
- [24] ———, *End spaces and spanning trees*, J. Combin. Theory (Series B) **96** (2006), no. 6, 846–854, DOI [10.1016/j.jctb.2006.02.010](https://doi.org/10.1016/j.jctb.2006.02.010). MR2274079 ↑[1.3.1](#), [4.1](#), [5.2.2](#), [6.1](#), [6.2](#), [6.2](#), [6.2.9](#), [6.2](#), [8.1](#), [8.2](#), [10.1](#), [10.4](#), [10.4](#), [11.1](#), [11.1](#), [11.4](#)
- [25] ———, *Ends and Tangles*, Abh. Math. Sem. Univ. Hamburg **87** (2017), no. 2, 223–244, DOI [10.1007/s12188-016-0163-0](https://doi.org/10.1007/s12188-016-0163-0), available at [arXiv:1510.04050v3](https://arxiv.org/abs/1510.04050v3). Special issue in memory of Rudolf Halin. MR3696148 ↑[1.1.1](#), [3.1](#), [3.2.2](#), [3.2.2](#), [3.3](#), [4.1](#), [4.1](#), [4.2.1](#), [4.2.2](#), [4.2.4](#), [4.3.1](#), [4.5](#), [4.6.2](#), [5.1](#), [5.2.3](#), [8.1](#), [11.1](#)
- [26] ———, *Graph Theory*, 5th, Springer, 2016. ↑[1](#), [1.1.1](#), [1.2.1](#), [1.2.2](#), [2](#), [2.1](#), [2.3](#), [2.3.1](#), [2.3.3](#), [2.4.1](#), [3.1](#), [3.7](#), [4.1](#), [4.3.1](#), [5.1](#), [5.1](#), [5.1](#), [5.1](#), [5.2.1](#), [5.2.2](#), [5.2.1](#), [4](#), [5.2.4](#), [5.2.4](#), [5.2.6](#), [5.2.14](#), [5.2.6](#), [5.3](#), [5.3.4](#), [5.3.8](#), [5.3.4](#), [5.3.5](#), [6.1](#), [6.1](#), [6.2](#), [6.3.3](#), [6.3.3](#), [7.1](#), [7.1](#), [7.1](#), [7.1](#), [8.1](#), [9.1](#), [9.1](#), [9.1](#), [9.3](#), [10.1](#), [10.3](#), [11.1](#), [11.1](#), [11.2.1](#), [11.2.2](#), [12.1](#), [12.5.2](#), [12.5.6](#), [14.1](#), [14.1](#)
- [27] ———, *Locally finite graphs with ends: a topological approach*, Discrete Math. **310**–**312** (2010), 2750–2765 (310); 1423–1447 (311); 21–29 (312), available at [arXiv:0912.4213v3](https://arxiv.org/abs/0912.4213v3). ↑[4](#)
- [28] ———, *The end structure of a graph: recent results and open problems*, Disc. Math. **100** (1992), no. 1, 313–327, DOI [10.1016/0012-365X\(92\)90650-5](https://doi.org/10.1016/0012-365X(92)90650-5). MR1172358 ↑[1.1.2](#), [5.2.7](#), [10.1](#), [10.4](#), [10.5](#)
- [29] ———, *Tree Sets*, Order **35** (2018), no. 1, 171–192, DOI [10.1007/s11083-017-9425-4](https://doi.org/10.1007/s11083-017-9425-4), available at [arXiv:arXiv:1512.03781v3](https://arxiv.org/abs/1512.03781v3). ↑[2.3](#), [4.2.3](#), [8.2](#)
- [30] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, J. Combin. Theory (Series B) **87** (2003), 197–206, DOI [10.1016/S0095-8956\(02\)00034-5](https://doi.org/10.1016/S0095-8956(02)00034-5). MR1967888 ↑[2.4.3](#), [2.4.1](#), [4.1](#), [5.2.1](#), [5.2.2](#), [5.2.7](#), [5.4.2](#), [9.3](#), [10.1](#), [11.1](#), [11.1](#)
- [31] ———, *Topological Paths, Cycles and Spanning Trees in Infinite Graphs*, Europ. J. Comb. **25** (2004), 835–862, DOI [10.1016/j.ejc.2003.01.002](https://doi.org/10.1016/j.ejc.2003.01.002). MR2079902 ↑[7.1](#)
- [32] R. Diestel and I. Leader, *Normal spanning trees, Aronszajn trees and excluded minors*, J. London Math. Soc. **63** (2001), 16–32, DOI [10.1112/S0024610700001708](https://doi.org/10.1112/S0024610700001708). MR1801714 ↑[1.3.1](#), [5.3.1](#), [7.1](#), [7.2.10](#), [9.1](#), [9.4](#), [10.1](#)
- [33] J.A. Dieudonné, *Une généralisation des espaces compacts*, J. Math. Pures. Appl. **23** (1944), 65–76. MR0013297 ↑[10.5](#)
- [34] C. Elbracht and J. Kneip and M. Teegen, *Trees of tangles in infinite separation systems*, available at [arXiv:2005.12122](https://arxiv.org/abs/2005.12122). Submitted. ↑[1.1.2](#), [4.1](#), [4.6](#)
- [35] A.K. Elm and J. Kurkofka, *A tree-of-tangles theorem for infinite-order tangles* (2020), available at [arXiv:2003.02535](https://arxiv.org/abs/2003.02535). Submitted. ↑(ii), [18](#)
- [36] R. Engelking, *General Topology*, 2nd ed., Sigma Series in Pure Mathematics, vol. 6, Heldermann Verlag, Berlin, 1989. MR1039321 ↑[2.2](#), [3.3](#), [3.3](#), [3.4.1](#), [3.4.1](#), [3.4.4](#), [3.4.1](#), [3.4.1](#), [3.4.1](#), [3.5.2](#), [3.5.11](#), [3.5.4](#), [3.7](#), [4.6.2](#), [4.6.2](#), [2](#), [5.2.2](#), [10.2](#), [10.4](#), [10.4](#), [10.5](#), [11.2](#)
- [37] H. Freudenthal, *Neuaufbau der Endentheorie*, Annals of Mathematics, posted on 1942, 261–279, DOI [10.2307/1968869](https://doi.org/10.2307/1968869). ↑[4.1](#)
- [38] J. Geelen and B. Joeris, *A generalization of the Grid Theorem* (2016), available at [arXiv:1609.09098](https://arxiv.org/abs/1609.09098). Submitted. ↑[5.2.1](#), [12.1](#), [14.1](#)

BIBLIOGRAPHY

- [39] L. Gillman and M. Jerison, *Rings of Continuous Functions*, Springer-Verlag, New York, 1976. MR0407579 ↑3.4.1
- [40] J.P. Gollin and K. Heuer, *Characterising k -connected sets in infinite graphs* (2018), available at [arXiv:1811.06411](https://arxiv.org/abs/1811.06411). Submitted. ↑5.2.1, 5.2.1, 12.1, 14.1
- [41] R. Halin, *Charakterisierung der Graphen ohne unendliche Wege*, Arch. Math. **16** (1965), no. 1, 227–231, DOI [10.1007/BF01220026](https://doi.org/10.1007/BF01220026). MR0179104 ↑5.1, 5.3.3
- [42] ———, *Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen*, Math. Nachr. **44** (1970), no. 1–6, 119–127, DOI [10.1002/mana.19700440109](https://doi.org/10.1002/mana.19700440109). MR0270953 ↑14.1
- [43] ———, *Graphen ohne unendliche Wege*, Math. Nachr. **31** (1966), no. 1–2, 111–123, DOI [10.1002/mana.19660310110](https://doi.org/10.1002/mana.19660310110). MR0191838 ↑5.1
- [44] ———, *Miscellaneous Problems on Infinite Graphs*, J. Graph Theory **35** (2000), 128–151, DOI [10.1002/1097-0118\(200010\)35:2<128::AID-JGT6>3.0.CO;2-6](https://doi.org/10.1002/1097-0118(200010)35:2<128::AID-JGT6>3.0.CO;2-6). MR1781193 ↑1.2.2, 7.3, 9.1
- [45] ———, *Simplicial decompositions of infinite graphs*, Advances in Graph Theory, Annals of Discrete Mathematics, 1978, DOI [10.1016/S0167-5060\(08\)70500-4](https://doi.org/10.1016/S0167-5060(08)70500-4). MR0499113 ↑5.2.1, 6.2, 7.3, 12.1, 14.1
- [46] ———, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Math. Nachr. **30** (1965), no. 1–2, 63–85, DOI [10.1002/mana.19650300106](https://doi.org/10.1002/mana.19650300106). MR0190031 ↑14.1
- [47] ———, *Über unendliche Wege in Graphen*, Math. Annalen **157** (1964), 125–137, DOI [10.1007/BF01362670](https://doi.org/10.1007/BF01362670). MR0170340 ↑1.1.2, 1.2.2, 2.4.1, 5.2.7, 7.3, 9.1, 11.1
- [48] K.P. Hart, *The Čech-Stone compactification of the real line*, Recent Progress in General Topology, North-Holland, Amsterdam (1992), 317–352. MR1229130 ↑3.4.1, 3.4.5, 3.4.6, 3.4.1, 3.4.2, 3.7
- [49] A. Hatcher, *Topology of numbers*, Book in preparation (2017). Available [online](#). ↑1.4, 12.1, 12.2, 13.1, 14.1
- [50] B. Hughes, *Trees and ultrametric spaces: a categorical equivalence*, Advances in Mathematics **189** (2004), no. 1, 148–191, DOI [10.1016/j.aim.2003.11.008](https://doi.org/10.1016/j.aim.2003.11.008). MR2093482 ↑10.4
- [51] B. Joeris, *Connectivity, tree-decompositions and unavoidable-minors*, University of Waterloo, 2015. ↑5.2.1, 12.1, 14.1
- [52] H.A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, Math. Nachr. **41** (1969), 1–22, DOI [10.1002/mana.19690410102](https://doi.org/10.1002/mana.19690410102). ↑5.1, 5.3, 5.3.5, 6.2, 7.1, 9.2, 9.2.3, 10.2.1, 11.2.3
- [53] J. Kneip and P. Gollin, *Representations of infinite tree sets*, Order, posted on 2020, DOI [10.1007/s11083-020-09529-0](https://doi.org/10.1007/s11083-020-09529-0), available at [arXiv:1908.10327](https://arxiv.org/abs/1908.10327). ↑2.3.4, 2.3.1, 4.2.3
- [54] P. Komjáth, *Martin’s axiom and spanning trees of infinite graphs*, J. Combin. Theory (Series B) **56** (1992), no. 1, 141–144, DOI [10.1016/0095-8956\(92\)90013-N](https://doi.org/10.1016/0095-8956(92)90013-N). MR1182464 ↑7.2
- [55] K. Kunen, *Set Theory*, Studies in logic, College Publications, 2011. MR2905394 ↑5.2.1
- [56] J. Kurkofka, *Every infinitely edge-connected graph contains the Farey graph or $T_{\aleph_0} * t$ as a minor* (2020), available at [arXiv:2004.06710](https://arxiv.org/abs/2004.06710). Submitted. ↑(x)
- [57] ———, *On the tangle compactification of infinite graphs* (2017), available at [arXiv:1908.10212](https://arxiv.org/abs/1908.10212). ↑13.5
- [58] ———, *The Farey graph is uniquely determined by its connectivity* (2020), available at [arXiv:2006.12472](https://arxiv.org/abs/2006.12472). Submitted. ↑(xi)
- [59] ———, *Ubiquity and the Farey graph* (2019), available at [arXiv:1912.02147](https://arxiv.org/abs/1912.02147). Submitted. ↑(xii)

BIBLIOGRAPHY

- [60] J. Kurkofka and R. Melcher, *Countably determined ends and graphs* (2020). In preparation. [↑11.2.2, 11.3, \(ix\), 18](#)
- [61] J. Kurkofka and R. Melcher and M. Pitz, *Approximating infinite graphs by normal trees* (2020), available at [arXiv:2002.08340](#). Submitted. [↑\(viii\), 18](#)
- [62] J. Kurkofka and M. Pitz, *Ends, tangles and critical vertex sets*, Math. Nachr. **292** (2019), no. 9, 2072–2091, DOI [10.1002/mana.201800174](#), available at [arXiv:1804.00588](#). [↑3.2.3, 3.4.2, 3.7, 4.1, 4.1, 4.2.1, 4.2.2, 4.2.2, 4.2.3, 4.2.2, 4.6.2, 4.6.2, 5.1, 5.2.3, 8.1, 8.2, 11.1](#)
- [63] ———, *Tangles and the Stone-Čech compactification of infinite graphs* (2018), available at [arXiv:1806.00220](#). Submitted. [↑\(i\), 18](#)
- [64] J. Lake, *A problem concerning infinite graphs*, Disc. Math. **14** (1976), no. 4, 343–345, DOI [10.1016/0012-365X\(76\)90066-2](#). MR0419297 [↑14.1](#)
- [65] B. Oporowski and J. Oxley and R. Thomas, *Typical Subgraphs of 3- and 4-connected Graphs*, J. Combin. Theory (Series B) **57** (1993), no. 2, 239–257, DOI [10.1006/jctb.1993.1019](#). MR1207490 [↑5.2.1, 12.1, 14.1](#)
- [66] M. Pitz, *A new obstruction for normal spanning trees* (2020), available at [arXiv:2005.04150](#). Submitted. [↑9.1](#)
- [67] N. Polat, *End-Faithful Spanning Trees in T_{\aleph_1} -Free Graphs*, J. Graph Theory **26** (1997), no. 4, 175–181, DOI [10.1002/\(SICI\)1097-0118\(199712\)26:4<175::AID-JGT1>3.0.CO;2-N](#). MR1487495 [↑7.2, 9.1, 9.1](#)
- [68] ———, *Ends and multi-endings I*, J. Combin. Theory (Series B) **67** (1996), 86–110, DOI [10.1006/jctb.1996.0035](#). MR1385385 [↑1.3.1, 4.1, 4.6.2, 10.1, 10.2, 10.3, 10.4, 11.1](#)
- [69] ———, *Ends and multi-endings II*, J. Combin. Theory (Series B) **1** (1996), 56–86, DOI [10.1006/jctb.1996.0057](#). MR1405706 [↑10.1, 10.4, 11.1](#)
- [70] ———, *Développements terminaux des graphes infinis III. Arbres maximaux sans rayon, cardinalité maximum des ensembles disjoints de rayons*, Math. Nachr. **115** (1984), no. 1, 337–352, DOI [10.1002/mana.19841150126](#). [↑7.1, 7.2](#)
- [71] ———, *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., posted on 1990, 197–220, DOI [10.1007/978-94-009-0517-7_16](#). MR1096994 [↑1.1.1, 3.1, 7.1, 7.2](#)
- [72] L. Ribes and P. Zalesskii, *Profinite Groups*, Springer, 2010. MR2599132 [↑2.2, 2](#)
- [73] N. Robertson and P. D. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Combin. Theory Ser. B **52** (1991), no. 2, 153–190, DOI [10.1016/0095-8956\(91\)90061-N](#). MR1110468 [↑3.1, 3.2.2, 4.1, 4.2.1, 5.1, 8.1](#)
- [74] N. Robertson and P.D. Seymour and R. Thomas, *Excluding infinite minors*, Disc. Math. **95** (1991), no. 1, 303–319, DOI [10.1016/0012-365X\(91\)90343-Z](#). MR1141945 [↑12.7](#)
- [75] ———, *Excluding subdivisions of infinite cliques*, Trans. Amer. Math. Soc. **332** (1992), 211–223, DOI [10.1090/S0002-9947-1992-1079057-3](#). MR1079057 [↑12.7](#)
- [76] P. Seymour and R. Thomas, *An end-faithful spanning tree counterexample*, Disc. Math. **95** (1991), no. 1, 321–330, DOI [10.1016/0012-365X\(91\)90344-2](#). MR1045600 [↑1.2.2, 7.1, 7.2, 7.2.2, 7.2, 9.1, 15, 16](#)
- [77] P.D. Seymour and R. Thomas, *Excluding infinite trees*, Trans. Amer. Math. Soc. **335** (1993), 597–630, DOI [10.1090/S0002-9947-1993-1079058-6](#). MR1079058 [↑1.2.2, 9.1, 9.1, 9.3, 9.3.1, 12.7](#)
- [78] 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, DOI [10.1007/BF01192782](#). MR0701276 [↑1.2.2, 5.1, 5.3.4, 9.1, 9.3](#)

BIBLIOGRAPHY

- [79] J. Širáň, *End-faithful forests and spanning trees in infinite graphs*, Disc. Math. **95** (1991), no. 1, 331–340, DOI [10.1016/0012-365X\(91\)90345-3](https://doi.org/10.1016/0012-365X(91)90345-3). MR1141946 ↑[7.1](#), [7.2](#), [7.2](#)
- [80] P. Sprüssel, *End spaces of graphs are normal*, J. Combin. Theory (Series B) **98** (2008), 798–804, DOI [10.1016/j.jctb.2007.10.006](https://doi.org/10.1016/j.jctb.2007.10.006). MR2418772 ↑[1.3.1](#), [4.1](#), [4.6.2](#), [10.1](#), [10.3](#), [10.5](#), [11.1](#)
- [81] J.M. Teegen, *Abstract Tangles as an Inverse Limit, and a Tangle Compactification for Topological Spaces*, Universität Hamburg, 2017. ↑[3.2.3](#)
- [82] C. Thomassen, *Duality of Infinite Graphs*, J. Combin. Theory (Series B) **33** (1982), 137–160, DOI [10.1016/0095-8956\(82\)90064-8](https://doi.org/10.1016/0095-8956(82)90064-8). MR0685062 ↑[7.3](#)
- [83] ———, *Infinite connected graphs with no end-preserving spanning trees*, J. Combin. Theory (Series B) **54** (1992), no. 2, 322–324, DOI [10.1016/0095-8956\(92\)90059-7](https://doi.org/10.1016/0095-8956(92)90059-7). ↑[1.2.2](#), [7.1](#), [7.2](#), [9.1](#), [15](#), [16](#)
- [84] H. Whitney, *Non-separable and planar graphs*, Trans. Am. Math. Soc. **34** (1932), 339–362, DOI [10.1090/S0002-9947-1932-1501641-2](https://doi.org/10.1090/S0002-9947-1932-1501641-2). ↑[7.1](#)
- [85] D.R. Woodall, *A note on a problem of Halin's*, J. Combin. Theory (Series B) **21** (1976), no. 2, 132–134, DOI [10.1016/0012-365X\(78\)90181-4](https://doi.org/10.1016/0012-365X(78)90181-4). MR0427154 ↑[14.1](#)

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.