On trees and tree-like structures in infinite graphs

Dissertation

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

vorgelegt
am Fachbereich Mathematik
von

Thilo Krill

Hamburg
2024

Vorsitzender der Prüfungskommission: Prof. Dr. Ingenuin Gasser
Erstgutachter und Betreuer: PD Dr. Max Pitz
Zweitgutachter: Prof. Dr. Nathan Bowler

Datum der Disputation: 24.01.2025

Contents

Chapter 1

Introduction

One of the deepest results in graph theory is Robertson and Seymour's graph minor theorem [73]. It states that finite graphs are well-quasi-ordered by the minor relation, i.e. there are no infinite decreasing chains and no infinite antichains of finite graphs with respect to minors. Let us look at two cornerstones on the way to this result. First, Kruskal [52] proved that finite trees are well-quasi-ordered, even by the more restrictive topological minor relation. To push Kruskal's result further, Robertson and Seymour defined classes of graphs that roughly look like trees, called graphs of bounded tree-width. As an important intermediate step for their graph minor theorem, they proved that graphs of bounded tree-width are well-quasi-ordered [72].

Similar approaches have been taken for a number of other famous problems in graph theory. For example, the first major advance towards a graph minor theorem for infinite graphs was Nash-Williams's result [60] that infinite trees are well-quasi-ordered by the topological minor relation. By extending the techniques of Nash-Williams, Thomas [77] proved that infinite graphs of bounded tree-width are well-quasi-ordered by the minor relation. Whether the same is true for all countable graphs is still unknown (but disproved for uncountable graphs [76]).

Another example concerning infinite graphs is Andreae's ubiquity conjecture, which suggests that all locally finite connected graphs are minor-ubiquitous. In a first step, Bowler et. al. [10] showed that infinite trees are topological minor-ubiquitous.[1] In two subsequent papers [11, 12], the same authors proved Andreae's conjecture for graphs of bounded tree-width.

However, one would not do justice to these results by regarding them only as intermediate steps. Trees and tree-like graphs are important graph classes in their own right, with applications for example in algorithmic graph theory (see [6]).

In this dissertation we solve four different problems in infinite graph theory that are once more concerned with (certain classes of) trees and tree-like graphs. A brief overview of these projects is given below. We use the terminology of Diestel's book [26] for graph theoretical concepts, of Bang and Jensen's book [5] for digraph theoretical concepts, of

[1] for definitions, see Section 4.1

Jech's book [44] for set theoretical concepts, and of Engelking's book [34] for topological concepts.

1.1 The number of topological types of trees (Chapter 2)

Two graphs are of the same *topological type* if they are topological minors of each other. We show that there are exactly $\aleph_1$ distinct topological types of countable trees. In general, for any infinite cardinal κ there are exactly κ^+ distinct topological types of trees of size κ. This solves a problem of van der Holst from 2005. Partial results on this topic were obtained by Matthiesen [59], who showed that the number of topological types of locally finite trees is at least $\aleph_1$, and subsequently by Bruno and Szeptycki [15], who showed that the number of topological types of locally finite trees with only countable many rays is exactly $\aleph_1$. Our proof uses a new rank function for trees inspired by Nash-Williams's work on the better-quasi-ordering of infinite trees [60]. An open problem for which the techniques developed here may be useful is to determine the number of distinct bounded tree-width graphs up to minor equivalence.

1.2 The Erdős-Pósa property for infinite graphs (Chapter 3)

A class $\mathcal{G}$ of graphs has the *Erdős-Pósa property (EPP)* if there is a function $f : \mathbb{N} \to \mathbb{N}$ such that for every graph Γ and every $k \in \mathbb{N}$ there are either k disjoint graphs from $\mathcal{G}$ in Γ or there is a set $X \subseteq V(\Gamma)$ of size at most $f(k)$ such that $\Gamma - X$ contains no graph from $\mathcal{G}$. In addition, we also consider the following infinite variant of the EPP: $\mathcal{G}$ has the *κ-EPP*, where κ is an infinite cardinal, if for any graph Γ there are either κ disjoint graphs from $\mathcal{G}$ in Γ or there is a set X of vertices of Γ of size less than κ such that $\Gamma - X$ contains no graph from $\mathcal{G}$.

We investigate which classes of infinite graphs have the EPP or the κ-EPP. In particular, we consider classes $\mathcal{G}$ consisting of a single infinite graph G. We show that if G is tree-like (more formally, if G admits a tree-decomposition into finite parts) and the set of induced subgraphs of G with vertices labelled by a finite set is well-quasi-ordered by the subgraph relation, then the class $\{G\}$ has the κ-EPP for every uncountable cardinal κ. If the condition of being tree-like is replaced by the stronger condition of being rayless, then G also has the $\aleph_0$-EPP and the EPP. As a corollary we get that every graph which does not contain a path of length n for some $n \in \mathbb{N}$ has the EPP and the κ-EPP. Furthermore, we show that the class of all subdivisions of any tree T has the κ-EPP for every uncountable cardinal κ, and if T is rayless, also the $\aleph_0$-EPP and the EPP. Since the $\aleph_0$-EPP is closely related to the ubiquity of graphs, our results yield corollaries concerning ubiquity.

1.3 Ubiquity in directed graphs (Chapters 4 and 5)

A (di)graph H is called *ubiquitous* if every (di)graph D containing k disjoint copies of H for every $k \in \mathbb{N}$ also contains infinitely many disjoint copies of H. The investigation of ubiquity in undirected graphs was initiated by Halin who showed that the ray is ubiquitous [38]. Several other results on the ubiquity or non-ubiquity of undirected graphs followed. These include Halin's result that the double ray is ubiquitous [40] and also the results shown in Chapter 3 of this dissertation.

First results on ubiquity in digraphs are presented in Chapters 4 and 5. An *oriented (double) ray* is a digraph with a (double) ray as underlying undirected graph. In Chapter 4 we characterise which oriented rays are ubiquitous, and in Chapter 5 we investigate the same question for oriented double rays. A *turn* of an oriented (double) ray is a vertex of in-degree 2 or out-degree 2, i.e. a vertex where the orientation "changes". We prove that an oriented ray is ubiquitous if and only if it has a finite number of turns, and that an oriented double ray with at least one turn is ubiquitous if and only if it has a (finite) odd number of turns. It remains an open problem to determine whether the consistently oriented double ray is ubiquitous. In addition, we propose to continue this line of work by investigating the ubiquity of other directed trees.

1.4 End spaces and tree-decompositions (Chapter 6)

The last chapter is of a slightly different character than the previous ones, but also deals with tree structures in infinite graphs. We present a systematic study of how tree-decompositions of finite adhesion capture properties of the topological space $|G|$ formed by a graph G together with its ends. The ends of G interact in a natural way with a tree-decomposition $(T, \mathcal{V})$ of G of finite adhesion. Indeed, as every edge e of T induces a finite order separation $\{A_e, B_e\}$ of G, each end of G has to choose one side of the separation and thus one component of $T - e$. By orienting the edges of T accordingly, we obtain an orientation of T for each end of G. Consider any set Ψ of ends of G. If for every end in Ψ the respective orientation of T points towards a (unique) end of T, and if this correspondence between Ψ and the ends of T is bijective, then we say that the tree-decomposition $(T, \mathcal{V})$ *displays* Ψ.

Although this definition is purely combinatorial, we find surprising connections between sets of ends that can be displayed and their topological properties in the space $|G|$. In particular, we show that the following are equivalent:

- There is a tree-decomposition of finite adhesion displaying Ψ.
- The subspace of $|G|$ consisting of Ψ together with all vertices and edges of G is completely metrizable.
- Ψ is G_δ (i.e. a countable intersection of open sets) in $|G|$.

Since the undominated ends of a graph are easily seen to be G_δ, this provides a structural explanation for Carmesin's result [18] that the set of undominated ends can always be displayed.

Chapter 2

The number of topological types of trees

2.1 Introduction

A graph-theoretic tree T is a *topological minor* of another tree S, written $T \leq S$, if some subdivision of T embeds as a subgraph into S. Nash-Williams [60] proved in 1965 the seminal result that the class of graph-theoretic trees is *well-quasi-ordered* under $\leq$, i.e. that it is a reflexive and transitive relation without infinite strictly decreasing sequences or infinite antichains.

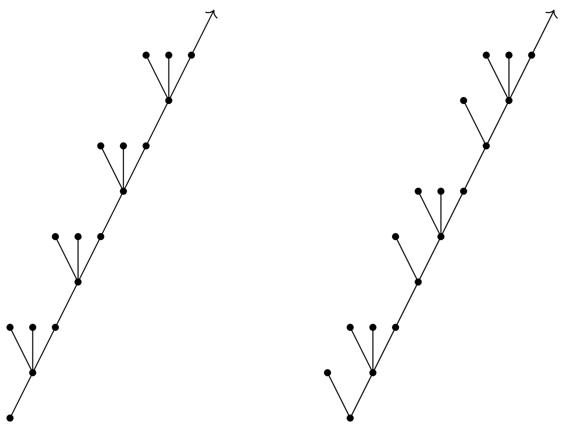

Figure 2.1: Distinct trees of the same topological type.

However, this embedding relation $\leq$ is not anti-symmetric: Two distinct trees T and S may well be topological minors of each other, i.e. $T \leq S$ and $S \leq T$. In this case, we say they are of the same *topological type*, written $T \equiv S$. Describing the hierarchy of graph-theoretic trees under the quasi-ordering $\leq$ means understanding the partial order that $\leq$ induces on the topological types of trees. By Nash-Williams's theorem, this is a well-partial-order. But determining its most fundamental characteristic, namely its cardinality, has been an open problem until now.

Indeed, while up to isomorphism there are exactly $2^{\aleph_0}$ countable trees, determining the exact number of topological types of countable trees is an open problem posed by van der Holst at the Oberwolfach graph theory workshop in 2005 (as reported by Matthiesen [59]). Building on the trees in Figure 2.1, we get examples of $2^{\aleph_0}$ non-isomorphic trees of the same topological type. So a priori, it would have been conceivable that there are only countably many topological types of countable trees. However, Matthiesen [59] showed in 2006, by an indirect proof building on Nash-Williams's theorem, that there are uncountably many topological types of countable trees. Bruno [14] in 2017 gave an explicit construction of uncountably many topological types of subtrees of the binary tree. Recently, Bruno and Szeptycki [15] gave the first indication that this bound could be sharp by establishing that there are exactly $\aleph_1$ many topological types of locally finite trees with only countably many rays.

Our main result confirms this pattern for all trees and all cardinalities. We write κ^+ for the successor cardinal of κ. The size of a tree is the cardinality of its vertex set.

Theorem. *For any infinite cardinal κ there are exactly κ^+ distinct topological types of trees of size κ.*

Our proof uses a new rank function for trees inspired by Nash-Williams's work on the better-quasi-ordering of infinite trees.

We end this section with two open problems. We say that two graphs are of the same *minor type* if they are minors of each other.

Problem 2.1.1. *What is the number of minor types of graphs with finitely bounded tree-width of given cardinality?*

Problem 2.1.2. *What is the number of minor types of countable graphs?*

For Problem 2.1.1, Thomas's result [77] that graphs of finitely bounded tree-width are well-quasi-ordered by minors may be helpful.

2.2 The lower bound

We recall Schmidt's rank function for rayless graphs [74], see also [43] for an English account: We say that a graph G has *rank* 0 if G is finite. Given an ordinal $\alpha > 0$, we assign *rank* α to G if G does not already have a rank $< \alpha$ and there exists a finite set of vertices X in G such that all components of $G - X$ have a rank $< \alpha$. Schmidt proved that a graph has a rank if and only if it is rayless. Moreover, a routine induction on the rank shows that Schmidt's rank function is non-decreasing with respect to the (topological) minor relation, see e.g. [43, Proposition 4.4]:

Lemma 2.2.1. *If a rayless graph H is a (topological) minor of a rayless graph G, then the rank of H is at most the rank of G.* □

We can now give the argument for the lower bound in our main theorem:

Lemma 2.2.2. *For any infinite cardinal κ there are at least κ^+ distinct topological types of (rayless) trees of size κ.*

Proof. We show that for all ordinals $0 < \alpha < \kappa^+$, there exists a rayless tree T_α of size κ and rank α. Then it follows from Lemma 2.2.1 that all T_α belong to different topological types, establishing the assertion of the lemma.

We construct the T_α by recursion on α, beginning with T_1 as the κ-star. For successor steps, take countably many disjoint copies T_n ($n \in \mathbb{N}$) of T_α and obtain $T_{\alpha+1}$ by adding a new vertex v to $\bigsqcup_{n\in\mathbb{N}} T_n$ and connecting it to the root of every T_n. Then deleting $\{v\}$ witnesses that $T_{\alpha+1}$ has rank at most $\alpha+1$. On the other hand, every finite set of vertices X of $T_{\alpha+1}$ avoids infinitely many copies of T_α, so there are components of $T_{\alpha+1} - X$ containing copies of T_α. Every such component has rank at least α by Lemma 2.2.1, showing that $T_{\alpha+1}$ has rank at least $\alpha + 1$.

For limit steps, obtain T_ℓ by adding a new vertex v to $\bigsqcup_{\alpha<\ell} T_\alpha$ and connecting it to the root of every T_α for $\alpha < \ell$. Then deleting $\{v\}$ witnesses that T_ℓ has rank at most ℓ; on the other hand, every finite set of vertices X of T_ℓ avoids almost all T_α copies for $\alpha < \ell$, so $T_\ell - X$ contains components of arbitrarily large rank below ℓ by Lemma 2.2.1, showing that T_ℓ has rank at least ℓ. □

2.3 Better-quasi-orderings

A *quasi-ordering* is a binary relation that is reflexive and transitive. A quasi-ordering $\leq$ on set Q is a *well-quasi-ordering* if for every sequence $q_1, q_2, q_3, \dots$ of elements in Q there are indices $n < m \in \mathbb{N}$ such that $q_n \leq q_m$. We define an equivalence relation $\equiv$ on Q by $q \equiv q'$ if both $q \leq q'$ and $q' \leq q$. We abbreviate $|Q|_\equiv := |Q/{\equiv}|$.

Let Q be quasi-ordered and κ an infinite cardinal. We say that $q \in Q$ is *κ-embeddable in Q* if there exist at least κ many elements $q' \in Q$ with $q \leq q'$. We need the following routine result, a proof of which can be found e.g. in [10, Lemma 3.3]:

Lemma 2.3.1. *For any well-quasi-order Q and infinite cardinal κ, the number of elements of Q which are not κ-embeddable in Q is less than κ.* □

Let $(Q, \leq_Q)$ be a quasi-order. Following Nash-Williams [60], we consider the quasi-ordering on the power set $\mathcal{P}(Q)$ where for $A, B \subseteq Q$ we have $A \leq B$ if there is an injective function $f\colon A \to B$ such that $a \leq_Q f(a)$ for all $a \in A$. Recall that $\mathcal{P}(Q)$ is not necessarily well-quasi-ordered if Q is well-quasi-ordered (see [70]). This is remedied by the introduction of the concept of a *better-quasi-ordering*. We shall not define this concept precisely; we only use as a blackbox that every better-quasi-ordered set is also well-quasi-ordered, that $\mathcal{P}(Q)$ is better-quasi-ordered if Q is better-quasi-ordered, and that the class of all trees is better-quasi-ordered under the topological minor relation [60] (also see [57, 58]).

We write $\mathcal{P}_\kappa(Q)$ for the set of subsets of Q of size exactly κ and $\mathcal{P}_{\leq\kappa}(Q)$ for the set of subsets of Q of size at most κ. Extending [15, Theorem 3], we prove the following result on the number of equivalence classes in $\mathcal{P}_\kappa(Q)$:

Lemma 2.3.2. *Let μ be an infinite cardinal and Q a better-quasi-ordered set with $|Q|_\equiv = \mu$. Then $|\mathcal{P}_\kappa(Q)|_\equiv = \mu$ for all cardinals $\kappa < \aleph_{\mu^+}$.*

Proof. By induction on κ. Suppose for a contradiction that $|\mathcal{P}_\kappa(Q)|_\equiv \geq \mu^+$. By the Erdős-Dushnik-Miller theorem, every partial order $(P, \leq)$ contains an infinite antichain or a chain of size $|P|$, see [32, Theorem 5.25]. As $(Q, \leq_Q)$ is better-quasi-ordered, $(\mathcal{P}_\kappa(Q), \leq)$ is well-quasi-ordered [60, Corollary 28A]. So the partial order $\mathcal{P}_\kappa(Q)/\equiv$ contains no infinite antichains and thus contains a chain of size μ^+. Since $\mathcal{P}_\kappa(Q)/\equiv$ is well-founded, this chain is well-ordered. Hence, there is a strictly increasing chain $\mathcal{A} = (A_\alpha \colon \alpha < \mu^+)$ in $\mathcal{P}_\kappa(Q)$.

By applying Lemma 2.3.1 to each induced suborder $(A_\alpha, \leq)$ of $(Q, \leq_Q)$, we obtain for every $A_\alpha \in \mathcal{A}$ a subset $X_\alpha \subseteq A_\alpha$ with $|X_\alpha| < \kappa$ such that all elements of $A_\alpha \setminus X_\alpha$ are κ-embeddable in $(A_\alpha, \leq)$. Since $|X_\alpha| < \kappa < \aleph_{\mu^+}$ for all $\alpha < \mu^+$, there are at most μ different possible cardinalities for the sets X_α (since $\kappa = \aleph_i < \aleph_{\mu^+}$ has at most $|i| \leq \mu$ many $\aleph$'s preceding it).

Since μ^+ is a successor cardinal and thus regular (i.e. any union of fewer than μ^+ sets each containing fewer than μ^+ elements has size less than μ^+), there is a cardinal $\nu < \kappa$ and a μ^+-sized subchain of $\mathcal{A}$ for which the respective $|X_\alpha|$ all take the same value ν. Hence, by restricting to that subchain and relabelling, we may assume that $|X_\alpha| = \nu$ for all $\alpha < \mu^+$ and some cardinal $\nu < \kappa$. Furthermore, by similar considerations, we may assume that all sets X_α for $\alpha < \mu^+$ are pairwise equivalent with respect to $\equiv$, since $|\mathcal{P}_\nu(Q)|_\equiv = \mu$ by induction.

Next, let $\{q_\beta \colon \beta < \mu\}$ be a representation system for the equivalence classes of $Q/\equiv$. For every q_β that is κ-embeddable in some $A \in \mathcal{A}$, we pick a suitable $A_{\alpha(\beta)} \in \mathcal{A}$ witnessing this. Let $\alpha^* := \sup\{\alpha(\beta) \colon \beta < \mu\} < \mu^+$. We arrive at the desired contradiction once we have proved that $A_\alpha \equiv A_{\alpha^*}$ for all $\alpha > \alpha^*$. Since $X_\alpha \equiv X_{\alpha^*}$ already, it suffices to show that

$$A_\alpha \setminus X_\alpha \leq A_{\alpha^*} \setminus X_{\alpha^*}$$

for all $\alpha > \alpha^*$. For this, we need an injective function $f : A_\alpha \setminus X_\alpha \to A_{\alpha^*} \setminus X_{\alpha^*}$ that satisfies $a \leq_Q f(a)$ for all $a \in A_\alpha \setminus X_\alpha$. Enumerate $A_\alpha \setminus X_\alpha = \{a_i \colon i < \kappa\}$, let $i < \kappa$, and suppose that f has been defined on a_j for all $j < i$. Since a_i is κ-embeddable in A_α, it is also κ-embeddable in $A_{\alpha'}$ for some $\alpha' \leq \alpha^*$ by the definition of α^*. Since $A_{\alpha'} \leq A_{\alpha^*}$, the element a_i is also κ-embeddable in A_{α^*}. Hence we can find an element $b \in A_{\alpha^*} \setminus X_{\alpha^*}$ such that $a_i \leq_Q b$ and b is distinct from all values of f that have already been defined. We set $f(a_i) := b$, which completes the construction of f. □

Corollary 2.3.3. *Let μ be an infinite cardinal and Q a better-quasi-ordered set with $|Q|_\equiv = \mu$. Then $|\mathcal{P}_{\leq\kappa}(Q)|_\equiv = \mu$ for all cardinals $\kappa < \aleph_{\mu^+}$.*

Proof. Since $\kappa < \aleph_{\mu^+}$, there exist at most μ cardinals $\leq \kappa$. Hence $|\mathcal{P}_{\leq\kappa}(Q)|_\equiv \leq \mu \cdot \mu = \mu$ by Theorem 2.3.2 applied to $\mathcal{P}_\nu(Q)$ for all cardinals $\nu \leq \kappa$. □

2.4 The upper bound

We consider rooted, graph theoretic trees and tree-order preserving topological minors. For this, we introduce a minimal amount of notation, cf. [26, §12.2]. Recall that fixing a root r of a graph-theoretic tree T yields a natural tree-order $\leq$ on T where $t \leq s$ if t lies on the unique path from r to s in T. Given a rooted tree, write $\mathrm{br}_T(t)$ (or simply $\mathrm{br}(t)$) for the subtree of T induced by the set $\{t' \in T\colon t \leq_r t'\}$ with root t. The neighbours of t in $\mathrm{br}(t)$ are the *successors* of t, denoted by the set $\mathrm{succ}_T(t)$ (or simply $\mathrm{succ}(t)$). Given rooted trees T and S, we write $T \leq S$ if there exists a topological minor embedding $\varphi\colon T \to S$ that preserves the tree-order: If $x \leq y$ in T then $\varphi(x) \leq \varphi(y)$ in S.

We now introduce a new rank function inspired by the proof methods of the better-quasi-ordering of trees due to Nash-Williams:

Definition 2.4.1. We say that a tree T has rank 0 if $\mathrm{br}(t) \equiv T$ holds for all $t \in T$. Given an ordinal $\alpha > 0$, we assign rank α to T if T does not already have a rank $< \alpha$ and for all $t \in T$, we have either $\mathrm{br}(t) \equiv T$ or $\mathrm{br}(t)$ has rank $< \alpha$. We also write $\mathrm{rank}(T)$ for the rank of T.

We remark that this rank function is not monotone with respect to subtrees. For example, a path on $n+1$ vertices rooted in one of its endvertices has rank n for all $n \in \mathbb{N}$. However, a ray rooted in its unique vertex of degree 1 has rank 0. Similarly, also the infinite rooted binary tree has rank 0. By taking disjoint paths of all finite lengths and adding a root which we connect to an endvertex of each path, we obtain a graph of infinite rank ω.

Lemma 2.4.2. *Every tree of size at most κ has a rank $< \kappa^+$.*

Proof. Suppose for a contradiction that there is a tree of size at most κ which does not have a rank $< \kappa^+$. Since rooted trees are well-quasi-ordered under $\leq$ by Nash-Williams's theorem [60], there exists a $\leq$-minimal such tree T. Then for every $t \in T$ with $\mathrm{br}(t) \not\equiv T$ we have $\mathrm{rank}(\mathrm{br}(t)) < \kappa^+$ by minimality of T. However, the rank of T is at most

$$\sup\{\mathrm{rank}(\mathrm{br}(t))\colon t \in T, \mathrm{br}(t) \not\equiv T\} + 1,$$

which is an ordinal $< \kappa^+$ since $|T| \leq \kappa$. This contradicts the choice of T. □

For the remainder of this section, let κ be a fixed infinite cardinal. We write $\mathcal{T}$ for the class of rooted trees of size at most κ, and $\mathcal{C}$ for the set of cardinals of size at most κ.

Let T be a tree in $\mathcal{T}$. For all $t \in T$, let

$$\Gamma(t) := (|\{s \in \mathrm{succ}(t)\colon \ \mathrm{br}(s) \equiv T\}|, \{\mathrm{br}(s)\colon s \in \mathrm{succ}(t), \mathrm{br}(s) \not\equiv T\}) \in \mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}).$$

Furthermore, we define

$$\Theta(T) := \{\Gamma(t)\colon t \in T\} \in \mathcal{P}_{\leq\kappa}(\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T})).$$

Given two quasi-orderings $(Q, \leq)$ and $(R, \leq)$, we define a quasi-ordering on $Q \times R$ by letting $(q, r) \leq (q', r')$ if $q \leq q'$ and $r \leq r'$. Together with the quasi-ordering on $\mathcal{P}(Q)$ defined in Section 2.3, this yields a quasi-ordering on the set $\mathcal{P}_{\leq\kappa}(\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}))$ considered in the definition of $\Theta(T)$ above. Nash-Williams showed in [60, Lemma 29]:

Lemma 2.4.3. *For all rooted trees with $\Theta(T) \leq \Theta(S)$, we have $T \leq S$.* □

Finally, we give the argument for the upper bound in our main theorem, in a stronger version for rooted trees:

Theorem 2.4.4. *For any infinite cardinal κ there are at most κ^+ distinct topological types of rooted trees of size κ.*

To see that Theorem 2.4.4 yields the same bound for unrooted trees, consider the map f mapping any rooted tree of size κ to its corresponding unrooted tree. Since the images under f of two equivalent rooted trees are also equivalent as unrooted trees, the map f induces a surjection from the topological types of κ-sized rooted trees to the topological types of κ-sized unrooted trees.

Proof. For all ordinals $\alpha < \kappa^+$, we write $\mathcal{T}_\alpha$ for the class of all trees of size at most κ and rank α and $\mathcal{T}_{<\alpha}$ for the class of all trees of size at most κ and rank $< \alpha$. We show by induction on α that $|\mathcal{T}_\alpha|_\equiv \leq \kappa$ holds for all $\alpha < \kappa^+$. Then it follows from Lemma 2.4.2 that

$$|\mathcal{T}|_\equiv = \left| \bigcup_{\alpha<\kappa^+} \mathcal{T}_\alpha \right|_\equiv \leq \kappa^+,$$

completing the proof.

Let $\alpha < \kappa^+$ and suppose $|\mathcal{T}_\beta|_\equiv \leq \kappa$ for all $\beta < \alpha$. Consider the function

$$\mathcal{T}_\alpha \to \mathcal{P}_{\leq\kappa}(\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha})),\, T \mapsto \Theta(T).$$

If $T, S \in \mathcal{T}_\alpha$ belong to different topological types, then also $\Theta(T)$ and $\Theta(S)$ belong to different equivalence classes of $\mathcal{P}_{\leq\kappa}(\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha}))$ by Lemma 2.4.3. We conclude that

$$|\mathcal{T}_\alpha|_\equiv \leq |\mathcal{P}_{\leq\kappa}(\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha}))|_\equiv.$$

Thus it suffices to show

$$|\mathcal{P}_{\leq\kappa}(\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha}))|_\equiv \leq \kappa.$$

First, we argue that $|\mathcal{T}_{<\alpha}|_\equiv \leq \kappa$: This is clear if $\alpha = 0$. If $\alpha > 0$, we have $|\mathcal{T}_\beta|_\equiv \leq \kappa$ for all $\beta < \alpha$ and hence $|\mathcal{T}_{<\alpha}|_\equiv = |\bigcup_{\beta<\alpha} \mathcal{T}_\beta|_\equiv \leq \kappa$ since $\alpha < \kappa^+$. Next, $\mathcal{T}$ and therefore $\mathcal{T}_{<\alpha}$ is better-quasi-ordered by [60] and thus $|\mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha})|_\equiv \leq \kappa$ by Corollary 2.3.3. Then it follows from cardinal arithmetic that also $|\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha})|_\equiv \leq \kappa$. Finally, since $\mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha})$ is better-quasi-ordered by [60, Corollary 28A] and hence $\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha})$ is better-quasi-ordered by [60, Corollary 22A], applying Corollary 2.3.3 again yields $|\mathcal{P}_{\leq\kappa}(\mathcal{C} \times \mathcal{P}_{\leq\kappa}(\mathcal{T}_{<\alpha}))|_\equiv \leq \kappa$. □

Chapter 3

The Erdős-Pósa property for infinite graphs

3.1 Introduction

3.1.1 The Erdős-Pósa property

Erdős and Pósa [35] established the following landmark result:

Theorem 3.1.1. *There is a function* $f : \mathbb{N} \to \mathbb{N}$ *such that for every graph* Γ *and every* $k \in \mathbb{N}$ *one of the following holds:*

- Γ *contains* k *disjoint cycles, or*
- *there is set* $X \subseteq V(\Gamma)$ *of size at most* $f(k)$ *such that* $\Gamma - X$ *is a forest.*

The following definition makes it possible to investigate whether theorems of this type hold for graphs other than cycles:

Definition 3.1.2. A class $\mathcal{G}$ of graphs has the *Erdős-Pósa property (EPP)* if there is a function $f : \mathbb{N} \to \mathbb{N}$ such that for every graph Γ (referred to as host graph) and every $k \in \mathbb{N}$ one of the following holds:

- Γ contains k disjoint copies of graphs from $\mathcal{G}$ as subgraphs, or
- there is set $X \subseteq V(\Gamma)$ of size at most $f(k)$ such that $\Gamma - X$ does not contain any copy of a graph from $\mathcal{G}$ as a subgraph.

For an overview of known results on classes of graphs $\mathcal{G}$ which have or do not have the EPP, see [71]. The graphs in $\mathcal{G}$ and the graph Γ in Definition 3.1.2 are usually required to be finite, but here we allow them to be infinite. A first natural question would be whether known results for finite graphs (such as Theorem 3.1.1) still hold if the host graph Γ can be infinite. In fact, this follows from a simple compactness argument (Proposition 3.3.1).

However, a completely new type of problem arises when not only Γ but also the graphs in $\mathcal{G}$ are infinite. The author is not aware of any previous work on this topic. We focus mainly on the special case where $\mathcal{G}$ consists of a single infinite graph. If $\mathcal{G} = \{G\}$ and $\mathcal{G}$ has the EPP, then we also say that G has the *EPP*.

Problem 3.1.3. *Which infinite graphs have the EPP?*

This problem may seem surprising, considering that every finite graph trivially has the EPP (Proposition 3.3.2). However, not every infinite graph has the EPP; the simplest counterexample is the ray (Proposition 3.3.4). More counterexamples are given in Proposition 3.3.6. On the other hand, we prove that rayless graphs with a certain additional property have the EPP: we say that a class $\mathcal{C}$ of graphs is *labelled well-quasi-ordered (lwqo)* if, for any finite set L, the class of all graphs from $\mathcal{C}$ with vertices labelled by elements of L is well-quasi-ordered (wqo) by label-preserving subgraph embeddings.

Theorem 3.1.4. *Let G be any rayless graph such that the set of induced subgraphs of G is lwqo. Then G has the EPP.*

An example of a class of graphs that is lwqo is the class of all graphs that exclude a path of fixed length as a subgraph. This was shown for finite graphs by Ding [31] and generalised to infinite graphs by Jia [45]. Hence we obtain the following corollary of Theorem 3.1.4:

Corollary 3.1.5. *Let G be a graph that does not contain a path of length n as a subgraph for some $n \in \mathbb{N}$. Then G has the EPP.*

For more results on labelled well-quasi-ordering see [4, 13, 21, 49, 69]. We believe that extending lwqo results to infinite graphs is an interesting problem and will lead to further applications of Theorem 3.1.4.

3.1.2 The κ-Erdős-Pósa property

We also consider the following infinite version of the Erdős-Pósa property, which will turn out to behave similarly to the "normal" Erdős-Pósa property:

Definition 3.1.6. Let κ be any infinite cardinal. We say that a class $\mathcal{G}$ of graphs has the *κ-Erdős-Pósa property (κ-EPP)* if for every graph Γ one of the following holds:

- Γ contains κ disjoint copies of graphs from $\mathcal{G}$ as subgraphs, or
- there is set $X \subseteq V(\Gamma)$ of size less than κ such that $\Gamma - X$ does not contain any copy of a graph from $\mathcal{G}$ as a subgraph.

Again, the author is not aware of any previous work on this topic. We say that a graph G has the EPP if the class $\{G\}$ does.

Problem 3.1.7. *Which infinite graphs have the κ-EPP?*

Addressing this problem, we offer the following results:

Theorem 3.1.8. *Let G be any rayless graph such that the set of induced subgraphs of G is lwqo. Then G has the $\aleph_0$-EPP.*

Theorem 3.1.9. *Let G be any graph admitting a tree-decomposition into finite parts such that the set of induced subgraphs of G is lwqo. Then G has the κ-EPP for every uncountable cardinal κ.*

The conditions on G in Theorems 3.1.4 and 3.1.8 are the same. In Theorem 3.1.9, the condition that G is rayless is replaced by the much weaker condition of admitting a tree-decomposition into finite parts. Indeed, every graph with a normal spanning tree has such a tree-decomposition (see [1, Theorem 2.2]) and thus also every graph that does not contain a subdivision of an infinite clique by [42, Theorem 10.1]. For more graphs with normal spanning trees see [27, 46, 63–65].

Theorems 3.1.4, 3.1.8 and 3.1.9 are the main results of this chapter. Their proofs rely on the same core ideas and we will prove all three theorems in Section 3.7.

Similar to Corollary 3.1.5, we deduce from Theorems 3.1.8 and 3.1.9:

Corollary 3.1.10. *Let G be a graph that does not contain a path of length n as a subgraph for some $n \in \mathbb{N}$. Then G has the κ-EPP for every infinite cardinal κ.*

3.1.3 Graphs without the κ-EPP for uncountable κ

While it is easy to find graphs that do not have the EPP or the $\aleph_0$-EPP (Proposition 3.3.4), a bit more work is needed to prove the following theorem:

Theorem 3.1.11. *For every uncountable cardinal κ there is a graph that does not have the κ-EPP.*

An important step in our proof of this theorem is to construct for each uncountable cardinal κ a graph of size κ which is not a proper subgraph of itself. Since it does not make our proof more difficult, we prove the following stronger theorem:

Theorem 3.1.12. *For every uncountable cardinal κ there is a graph of size κ which is not a proper minor of itself.*

This generalises Oporowski's result [61] that there exists a continuum-sized graph which is not a proper minor of itself.

3.1.4 Trees without the κ-EPP for uncountable κ

In an attempt to strengthen Theorem 3.1.11, we ask:

Problem 3.1.13. *Is there, for every uncountable cardinal κ, a tree that does not have the κ-EPP?*

Regarding Problem 3.1.13, we are only able to prove a consistency result, which is based on the assumption that there are no weak limit cardinals:

Theorem 3.1.14. *It is consistent with ZFC that for every uncountable cardinal κ there is a tree that does not have the κ-EPP.*

For the proof we need the following theorem, which is based on the same set-theoretic assumption as above:

Theorem 3.1.15. *It is consistent with ZFC that for every uncountable cardinal κ there is a tree of size κ which is not a proper subgraph of itself.*

It remains open whether this also holds in ZFC:

Problem 3.1.16. *Is there, for every uncountable cardinal κ, a tree of size κ which is not a proper subgraph of itself?*

In Section 3.6 we will see that a positive answer to the following problem would imply a positive answer to the latter:

Problem 3.1.17. *Is there, for every uncountable cardinal κ, a κ-sized subgraph-antichain of trees of size at most κ?*

Problem 3.1.18. *Do the answers to Problems 3.1.13, 3.1.16 or 3.1.17 change if we additionally require the trees to be rayless?*

3.1.5 The (κ-)EPP for classes defined by topological minors and minors

For a graph G, write $\mathcal{T}(G)$ for the class of graphs containing G as a topological minor and $\mathcal{M}(G)$ for the class of graphs containing G as a minor. Classes of this form are typically considered in results on the EPP for finite graphs. In Section 3.8 (Theorem 3.8.1), we formulate and prove a version of Theorems 3.1.4, 3.1.8, and 3.1.9 for classes of the form $\mathcal{T}(G)$. Note that similar results on the EPP or $\aleph_0$-EPP are not possible for classes of the form $\mathcal{M}(G)$, since not even $\mathcal{M}(K_{1,\aleph_0})$ has the EPP or the $\aleph_0$-EPP (Proposition 3.3.7).

Nash-Williams [60] showed that the class of all trees is well-quasi-ordered by the topological minor relation and Laver [58] proved a more general labelled version of this result. With help of Laver's theorem, we deduce the following result:

Corollary 3.1.19. *$\mathcal{T}(T)$ has the κ-EPP for every uncountable cardinal κ and every tree T. If T is rayless, then $\mathcal{T}(T)$ also has the $\aleph_0$-EPP.*

A corresponding statement for the EPP does not hold, since Thomassen [78] showed that there are finite trees T for which $\mathcal{T}(T)$ does not have the EPP.

An argument similar to Proposition 3.3.4 shows that $\mathcal{T}(R)$ and $\mathcal{M}(R)$, where R is a ray, do not have the EPP or the $\aleph_0$-EPP. However, we know of no such counterexamples for the κ-EPP for uncountable κ.

Problem 3.1.20. *Is there an uncountable cardinal κ and a graph G such that $\mathcal{T}(G)$ or $\mathcal{M}(G)$ does not have the κ-EPP? Does such a graph exist for every uncountable cardinal κ?*

3.1.6 Comparing and evaluating the results

Considering Theorems 3.1.4 and 3.1.8, Proposition 3.3.6, and Corollary 3.1.19, the $\aleph_0$-EPP and the EPP seem to behave very similarly. The κ-EPP for uncountable κ behaves differently. In fact, it is easy to see that any graph of size less than κ has the κ-EPP (Proposition 3.3.3). So in particular, the ray has the κ-EPP for all uncountable κ, but not the EPP or the $\aleph_0$-EPP (Proposition 3.3.4). However, if we restrict ourselves to graphs of size at least κ, we do not know whether such examples exist:

Problem 3.1.21. *Let $\mu \leq \kappa$ be infinite cardinals and let G be a graph of size at least κ. Does G have the κ-EPP if and only if G has the μ-EPP if and only if G has the EPP?*

Viewed from a distance, the question whether a graph G has the (κ-)EPP seems to be related to the self-similarity of G: in the proofs of Theorems 3.1.4, 3.1.8, and 3.1.9 the lwqo property is used to find various proper subgraph embeddings of G into itself. On the other hand, the graphs without the κ-EPP, which we construct in the proof of Theorem 3.1.11, do not admit any proper subgraph embeddings into themselves.

3.1.7 Ubiquity

The results of this chapter have applications in the field of ubiquity. Call a class $\mathcal{G}$ of graphs ubiquitous if every graph Γ that contains n disjoint copies of graphs from $\mathcal{G}$ as subgraphs for all $n \in \mathbb{N}$ also contains infinitely many such copies. It is easy to see that if $\mathcal{G}$ has the $\aleph_0$-EPP, then $\mathcal{G}$ is ubiquitous (Proposition 3.3.8). However, the converse is not true, since the ray is ubiquitous by a result of Halin [38].

Thus, all graphs as in Theorem 3.1.8 and Corollary 3.1.10 are ubiquitous, which was not known before. Furthermore, by Corollary 3.1.19 the class $\mathcal{T}(T)$ is ubiquitous for any rayless tree T. With that, we reobtain a special case of the result of Bowler et. al. [10] that $\mathcal{T}(T)$ is ubiquitous for every tree T.

3.2 Preliminaries

3.2.1 Minors

Recall that a graph H is a *minor* of a graph G if there is a family $\{V_h : h \in V(H)\}$ of connected, pairwise disjoint, non-empty subsets of $V(G)$ such that there is a V_h–$V_{h'}$ edge in G whenever there is a h–h' edge in H. The sets V_h for $h \in V(H)$ are called *branch sets*.

3.2.2 Tree order

All trees in this chapter have a fixed root vertex, which we will usually not explicitly specify. (In some parts of the chapter, root vertices of trees are not needed and can be ignored.) Recall that the tree-order of a tree T with root r is defined by $x \leq y$ if x lies on the unique r–y path in T. As in the previous chapter, we denote by $\mathrm{br}_T(x)$ the subtree of T with vertex set $\{y \in V(T) : x \leq y\}$ and we write $\mathrm{succ}_T(x)$ for the set of immediate successors of x in $V(T)$. Furthermore, we call a subtree of T *rooted* if it contains r.

3.2.3 Tree-decompositions

A *tree-decomposition* of a graph G is a pair $(T, \mathcal{V})$, where T is a tree and $\mathcal{V} = (V_t : t \in V(T))$ is a family of subsets of $V(G)$ such that:

- $V(G) = \bigcup\{V_t : t \in V(T)\}$,
- for every edge vw of G there is a $t \in V(T)$ such that $v \in V_t$ and $w \in V_t$, and
- if $t_1, t_2, t_3 \in V(T)$ such that t_2 lies on the unique t_1–t_3 path in T, then $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$.

The sets V_t for $t \in V(T)$ are called *parts* of the tree-decomposition.

Halin [39] showed:

Lemma 3.2.1. *Every rayless graph admits a tree-decomposition $(T, \mathcal{V})$ into finite parts such that T is rayless.*

3.2.4 Compactness

Following Diestel's book [26, Appendix A], we describe the Compactness Principle, a combinatorial framework for compactness proofs.

Let V be an arbitrary set, let $\mathcal{F}$ be a set of finite subsets of V, and let S be any finite set. Suppose for each $U \in \mathcal{F}$ we have fixed a set of *admissible* functions $U \to S$. We call a subset $\mathcal{U}$ of $\mathcal{F}$ *compatible* if there is a function $g : V \to \mathcal{F}$ such that $g \restriction U$ is admissible for all $U \in \mathcal{U}$.

Theorem 3.2.2 (Compactness Principle)**.** *If all finite subsets of $\mathcal{F}$ are compatible, then $\mathcal{F}$ is compatible.*

3.3 Propositions

We say that a class $\mathcal{G}$ of finite graphs has the *EPP for finite host graphs* if it has the EPP as defined in the beginning of Section 3.1 with the additional requirement that the host graph Γ is finite.

Proposition 3.3.1. *Any class $\mathcal{G}$ of finite graphs has the EPP for finite host graphs if and only if it has the EPP.*

Proof. If $\mathcal{G}$ has the EPP, then it clearly has the EPP for finite host graphs. Conversely, let $f : \mathbb{N} \to \mathbb{N}$ witness that $\mathcal{G}$ has the EPP for finite host graphs. We show that the same function witnesses that $\mathcal{G}$ has the EPP. Let $k \in \mathbb{N}$ and let Γ be any infinite graph that does not contains k disjoint copies of graphs from $\mathcal{G}$. We have to find a set $X \subseteq V(\Gamma)$ of size at most $f(k)$ such that $\Gamma - X$ does not contain any copy of a graph from $\mathcal{G}$. This is achieved by a standard compactness argument, for which we use the Compactness Principle (see Section 3.2.4). We provide the details in the following.

Let $\mathcal{F}$ be the set of all finite subsets of $V(\Gamma)$ and let $U \in \mathcal{F}$. We call a function $g : U \to \{0,1\}$ admissible if $X' := g^{-1}\{1\}$ has size at most $f(k)$ and $\Gamma[U] - X'$ does not contain any copy of a graph from $\mathcal{G}$. We show that every finite $\mathcal{U} \subseteq \mathcal{F}$ is compatible. Indeed, since $U^* := \bigcup \mathcal{U} \subseteq V(\Gamma)$ is finite and since $\mathcal{G}$ has the EPP for finite host graphs, there is a set $X^* \subseteq U^*$ of size at most $f(k)$ such that $\Gamma[U^*] - X^*$ does not contain any copy of a graph from $\mathcal{G}$. Then the characteristic function $g^* : V(\Gamma) \to \{0,1\}$ of X^* witnesses that $\mathcal{U}$ is compatible. By the Compactness Principle (Theorem 3.2.2), also $\mathcal{F}$ is compatible; let $h : V(\Gamma) \to \{0,1\}$ witness compatibility of $\mathcal{F}$. It is straightforward to check that $X := h^{-1}\{1\}$ has size at most $f(k)$ and that $\Gamma - X$ does not contain any copy of a member of $\mathcal{G}$. □

Proposition 3.3.2. *Every finite graph G has the EPP.*

Proof. Let $f : \mathbb{N} \to \mathbb{N}, k \mapsto |G| \cdot (k-1)$. Let $k \in \mathbb{N}$ and consider any graph Γ such that $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ of size at most $f(k)$. We recursively find k disjoint copies of G in Γ. Having found disjoint copies $G_1, \dots, G_i$ for $i < k$, we find another disjoint copy by choosing $X := V(G_1) \cup \dots \cup V(G_i)$, which is a set of size at most $f(k)$. □

Proposition 3.3.3. *For every infinite cardinal κ, every graph G of size less than κ has the EPP.*

Proof. Let Γ be any graph such that $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ of size less than κ. We recursively construct a family $(G_\alpha)_{\alpha<\kappa}$ of pairwise disjoint copies of G in Γ. If $\alpha < \kappa$ and we have defined G_β for all $\beta < \alpha$, then we find G_α by choosing $X := \bigcup\{V(G_\beta) : \beta < \alpha\}$, which is a set of size less than κ. □

Proposition 3.3.4. *The ray does not have the EPP or the $\aleph_0$-EPP.*

Proof. Consider a ray as host graph Γ. Then Γ does not contain contain 2 disjoint rays. However, for every finite $X \subseteq V(\Gamma)$ the graph $\Gamma - X$ still contains a ray. Therefore, the ray does not have the EPP or the $\aleph_0$-EPP (the former holds because it is not possible to define $f(2)$). □

For Proposition 3.3.6, we need the following result by Halin [41, Theorem 2]:

Lemma 3.3.5. *Let G be a connected infinite locally finite graph such that the maximum number of disjoint rays in G is $n \in \mathbb{N}$. Then there is a tree-decomposition $(R, \mathcal{V})$ of G into finite parts where R is a ray such that all adhesion sets have size n and are pairwise disjoint.*

Proposition 3.3.6. *Let G be a connected infinite locally finite graph that does not contain infinitely many disjoint rays. Then G does not have the EPP or the $\aleph_0$-EPP.*

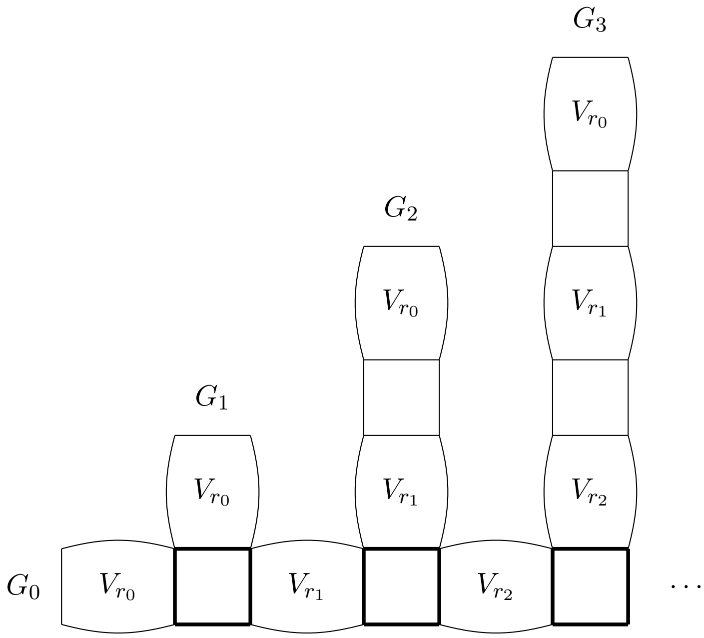

Figure 3.1: The graph Γ from the proof of Proposition 3.3.6. Adhesion sets in the graphs G_i are depicted by squares and each part V_{r_i} consists of the respective labelled area together with its adjacent squares. Note that every ray in Γ meets all but finitely many of the thickly drawn adhesion sets.

Proof. Halin showed in [38] that the ray is ubiquitous. Therefore, since G does not contain infinitely many disjoint rays, there is a maximum number $n \in \mathbb{N}$ of disjoint rays in G.

Consider the tree-decomposition $(R, \mathcal{V})$ of G from Lemma 3.3.5 where $R = r_0 r_1 r_2 \dots$. Let $G_0, G_1, G_2, \dots$ be infinitely many copies of G such that:

- For all $j < k \in \mathbb{N}$ the sets of vertices of G_j and G_k corresponding to $V(G) \setminus \bigcup\{V_{r_i} : i \geq k\}$ are disjoint.
- For all $j < k \leq i \in \mathbb{N}$ and all $v \in V_{r_i}$, the vertex of G_j corresponding to v coincides with the vertex of G_k corresponding to v.

Let $\Gamma := \bigcup_{k \in \mathbb{N}} G_k$ (see Figure 3.1). Since $\bigcap_{k \in \mathbb{N}} G_k = \emptyset$, for every finite $X \subseteq V(\Gamma)$ there is $k \in \mathbb{N}$ such that G_k avoids X.

Furthermore, it is clear from Figure 3.1 that for every ray S in Γ there is $i \in \mathbb{N}$ such that for all $j \geq i$, the ray S meets the set of vertices of G_0 corresponding to the adhesion set $V_{r_j} \cap V_{r_{j+1}}$. As all adhesion sets have size n, the graph Γ cannot contain more than n disjoint rays. Therefore, Γ does not contain two disjoint copies of G. Since we have also seen that $\Gamma - X$ contains a copy of G for all finite $X \subseteq V(\Gamma)$, it follows that G does not have the EPP or the $\aleph_0$-EPP. $\square$

Proposition 3.3.7. *$\mathcal{M}(K_{1,\aleph_0})$ does not have the EPP or the $\aleph_0$-EPP.*

Proof. Let Γ be the infinite comb (Figure 3.2). Then Γ does not contain 2 disjoint copies of $K_{1,\aleph_0}$ as minors, but for every finite set $X \subseteq V(\Gamma)$, the graph $\Gamma - X$ contains $K_{1,\aleph_0}$ as a minor. Thus $\mathcal{M}(K_{1,\aleph_0})$ does not have the EPP or the $\aleph_0$-EPP. □

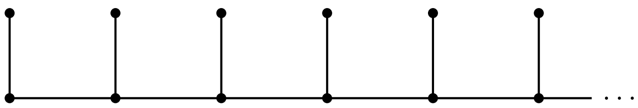

Figure 3.2: The infinite comb.

Proposition 3.3.8. *If a class of graphs has the $\aleph_0$-EPP, then it is ubiquitous.*

Proof. Let $\mathcal{G}$ be any class of graphs that has the $\aleph_0$-EPP. To see that $\mathcal{G}$ is ubiquitous, let Γ be any graph that contains n disjoint copies of graphs from $\mathcal{G}$ for all $n \in \mathbb{N}$. Then $\Gamma - X$ contains a copy of G for every finite set $X \subseteq V(G)$, and it follows from the $\aleph_0$-EPP that Γ contains infinitely many disjoint copies of graphs from $\mathcal{G}$. □

3.4 Constructing counterexamples to Seymour's Self-Minor Conjecture

A graph H is a *proper minor* of a graph G if H is a minor of some proper subgraph of G. In this section, we prove:

Theorem 3.1.12. *For every uncountable cardinal κ there is a graph of size κ which is not a proper minor of itself.*

Our construction relies on the existence of large minor-antichains of infinite graphs. Their existence is proved by Komjáth [48] (also see [66]):

Theorem 3.4.1. *For every uncountable cardinal κ, there are 2^κ graphs of size κ none of them being a minor of another.*

Let us briefly look at the context of Theorem 3.1.12. Seymour's Self-Minor Conjecture states that every infinite graph is a proper minor of itself. While it is still open whether the conjecture holds for countable graphs, Oporowski [61] disproved Seymour's conjecture by constructing a continuum-sized graph which is not a proper minor of itself. In Theorem 3.1.12 we show more generally that there exist counterexamples of all uncountable cardinalities. It should be noted that Oporowski's result (1990) is older than Komjáth's Theorem 3.4.1 (1995), and instead relies on a weaker result by Thomas [76] from 1988. Our construction is different from Oporowski's and slightly simpler, but some of the ideas and tools we use are the same:

Lemma 3.4.2 ([61], Lemma 1)**.** *Let G' be obtained from G by adding a new vertex adjacent to every vertex of G, and let H' be obtained from H similarly. Then H is a minor of G if and only if H' is a minor of G'.*

A *block* of a graph is a maximal connected subgraph that has no cutvertex.

Lemma 3.4.3 ([61], Lemma 3)**.** *Let H be a minor of G with branch sets $\{V_h : h \in V(H)\}$. Then for every block A of H there is a block B of G such that A is a minor of B with branch sets $\{V_h \cap V(B) : h \in V(A)\}$.*

Proof of Theorem 3.1.12. By Theorem 3.4.1 there is a κ-sized minor-antichain of graphs of size κ. For every graph in this antichain, add two vertices which are adjacent to all vertices of this graph and to each other. The resulting set $\mathcal{H}$ of graphs is again a minor-antichain by Lemma 3.4.2 and all graphs in $\mathcal{H}$ are 2-connected. Suppose without loss of generality that the graphs in $\mathcal{H}$ are pairwise disjoint.

Let T be a κ-regular tree and fix a bijection $f : V(T) \to \mathcal{H}$. Moreover, fix for every $t \in V(T)$ a bijection $g_t : N_T(t) \to V(f(t))$, which is possible since $N_T(t)$ and $V(f(t))$ are both sets of size κ. We construct a graph G from $\bigcup \mathcal{H}$ by identifying for every edge $tu \in E(T)$ the vertices $g_t(u)$ and $g_u(t)$. Note that the elements of $\mathcal{H}$ are, up to isomorphism, precisely the blocks of G, and every $x \in V(G)$ is contained in precisely two blocks of G.

Now suppose that G is a minor of itself with branch sets $\{V_x : x \in V(G)\}$. To show that G is not a proper minor of itself, we prove the stronger assertion that $V_x = \{x\}$ for all $x \in V(G)$.

By Lemma 3.4.3, for every block A of G there is a block B of G such that A is a minor of B with branch sets $\{V_x \cap V(B) : x \in V(A)\}$. Since $\mathcal{H}$ is a minor-antichain, we must have $A = B$. In particular, $V_x \cap V(A) \neq \varnothing$ for all $x \in V(A)$.

Consider any vertex $x \in V(G)$ and let A, A' be the two blocks of G containing x. As observed above, V_x meets both A and A'. Since V_x is connected and x separates A from A' in G, we must have $x \in V_x$. Since the same holds for all $x \in V(G)$ and since distinct branch sets are disjoint, it follows that $V_x = \{x\}$. □

3.5 Graphs without the κ-EPP

In this section, we prove:

Theorem 3.1.11. *For every uncountable cardinal κ there is a graph that does not have the κ-EPP.*

We will obtain such a graph by applying the following operation to any κ-sized graph which is not a proper subgraph of itself. For a graph G, we write $G^\vee$ for the graph obtained from G by adding two new vertices v_1, v_2 for each $v \in V(G)$ and adding edges vv_1 and vv_2. Note that every vertex of $G^\vee$ has either degree 1 or degree at least 3, and the set of vertices of degree at least 3 of $G^\vee$ is $V(G)$.

The following lemma is straightforward and we omit the proof:

Lemma 3.5.1. *A graph H is a subgraph of a graph G if and only if $H^\vee$ is a subgraph of $G^\vee$.*

Theorem 3.5.2. *Let κ be an infinite cardinal and let G be a graph of size at least κ which is not a proper subgraph of itself. Then $G^\vee$ does not have the κ-EPP.*

Proof. Consider a set $\mathcal{S}$ consisting of κ copies of $G^\vee$ such that:

(i) for all $H \neq H' \in \mathcal{S}$, the set of vertices of degree at least 3 of H is disjoint from the set of vertices of degree at least 3 of H',

(ii) for all $H \neq H' \in \mathcal{S}$, a degree 1 vertex of H coincides with a degree 1 vertex of H', and

(iii) every vertex of $\bigcup \mathcal{S}$ is contained in at most two distinct graphs from $\mathcal{S}$.

Since $G^\vee$ has κ vertices of degree 1, such a set $\mathcal{S}$ can easily be obtained by beginning with κ disjoint copies of $G^\vee$ and recursively identifying pairs of degree 1 vertices.

We show that $\Gamma := \bigcup \mathcal{S}$ witnesses that $G^\vee$ does not have the κ-EPP. Indeed, by (iii), $\Gamma - X$ contains a copy of $G^\vee$ for all $X \subseteq V(\Gamma)$ of size less than κ. It is left to show that Γ does not contain two disjoint copies of $G^\vee$. Note that by (i) and (iii), every vertex of Γ that is contained in more than one graph from $\mathcal{S}$ has degree 2. Therefore, since $G^\vee$ has no degree 2 vertices, any copy of $G^\vee$ in Γ is completely contained in some member of $\mathcal{S}$. As $G^\vee$ is not a proper subgraph of itself by Lemma 3.5.1, any copy of $G^\vee$ in Γ must even coincides with a member of $\mathcal{S}$. Thus Γ does not contain two disjoint copies of $G^\vee$ by (ii). □

Proof of Theorem 3.1.11. By Theorem 3.1.12, there is a graph G of size κ which is not a proper minor of itself and in particular not a proper subgraph of itself. Then $G^\vee$ does not have the κ-EPP by Theorem 3.5.2. □

3.6 Trees without the κ-EPP

Using a similar strategy as in the proof of Theorem 3.1.11 given in the previous two sections, we show:

Theorem 3.1.14. *It is consistent with ZFC that for every uncountable cardinal κ there is a tree that does not have the κ-EPP.*

The set-theoretic assumption we need is that there exist no regular limit cardinals. This assumption is used in Theorem 3.6.3, in which we construct a 2^κ-sized subgraph-antichain of κ-sized trees. Using the graphs from this antichain as building blocks, we construct a tree T of size κ which is not a proper subgraph of itself in Theorem 3.1.15. Finally, we deduce that $T^\vee$ does not have the κ-EPP. We do not know if any of these theorems also hold in ZFC (see Problems 3.1.13, 3.1.16, 3.1.17 and 3.1.18).

We begin by proving two lemmas that are needed for Theorem 3.6.3. Our proof of Theorem 3.6.3, including the two lemmas, is based on ideas from Komjáth's proof of Theorem 3.4.1 [48].

Lemma 3.6.1. *Let S be any set of infinite cardinality κ. Then there is a 2^κ-sized set $\mathcal{P}$ of κ-sized subsets of S such that no element of $\mathcal{P}$ is a subset of another.*

Proof. Let P be a partition of S into 2-element subsets and note that $|P| = |S|$ since S is infinite. Then the set $\mathcal{P}$ of all κ-sized subsets of S that contain exactly one of the two elements of all sets in P is as required. $\square$

Lemma 3.6.2. *Let κ be any infinite cardinal. If there exists a κ-sized subgraph-antichain of (rayless) trees of size at most κ, then there exists a 2^κ-sized subgraph-antichain of (rayless) trees of size κ.*

Proof. Let $\mathcal{U}$ be a κ-sized subgraph-antichain of pairwise disjoint trees of size at most κ and let $\mathcal{S}$ be a 2^κ-sized $\subseteq$-antichain of κ-sized subsets of $\mathcal{U}$, which exists by Lemma 3.6.1. For every $S \in \mathcal{S}$, let T_S be a tree obtained from $\bigcup S$ by adding a new vertex v_S and joining it to one vertex of each tree in S. We claim that $\mathcal{T} := \{T_S : S \in \mathcal{S}\}$ is the desired subgraph-antichain.

Clearly, the trees in $\mathcal{T}$ have size κ, and if the trees in $\mathcal{U}$ are rayless, then the trees in $\mathcal{T}$ are rayless as well. Finally, let $S \neq S' \in \mathcal{S}$ and suppose for a contradiction that there is a subgraph embedding $f : T_S \to T_{S'}$. We cannot have $f(v_S) = v_{S'}$ since then f would have to send the trees from S injectively to trees from S', but $S \not\subseteq S'$ and $\mathcal{U}$ is a subgraph-antichain. Hence $f(v_S) \in U$ for some tree $U \in S'$. But since the vertex $v_{S'}$ separates U from all other elements of S' in $T_{S'}$, it follows that f embeds all but at most one tree from S into U, contradicting that $\mathcal{U}$ is a subgraph-antichain. $\square$

Theorem 3.6.3. *It is consistent with ZFC that for every infinite cardinal κ, there are 2^κ rayless trees of size κ none of them being a subgraph of another.*

Proof. We assume that there exist no regular limit cardinals, which are also known as weakly inaccessible cardinals and whose existence is not provable in ZFC (see [44]). Under this assumption, we prove the theorem by induction on κ.

To find a $2^{\aleph_0}$-sized antichain of $\aleph_0$-sized rayless trees, apply Lemma 3.6.2 to the set of double-ended forks shown in Figure 3.3.

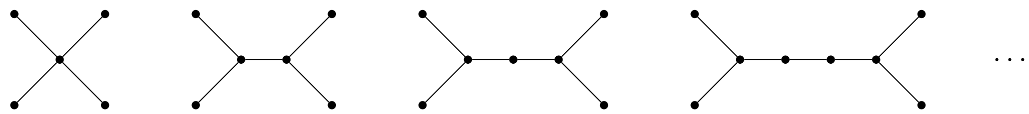

Figure 3.3: A countable subgraph-antichain of finite trees.

Now let $\kappa > \aleph_0$ and suppose that a 2^κ-sized antichain $\mathcal{T}$ of κ-sized rayless trees is given. Since $|\mathcal{T}| = 2^\kappa \geq \kappa^+$, there exists a $2^{(\kappa^+)}$-sized antichain of κ^+-sized rayless trees by Lemma 3.6.2.

Finally, suppose that κ is a limit cardinal and that the theorem holds for all cardinals less than κ. By the assumption that there are no regular limit cardinals, κ is singular. Consider an increasing cofinal sequence of cardinals $(\kappa_\alpha : \alpha < \mathrm{cf}(\kappa))$ in κ with $\mathrm{cf}(\kappa) < \kappa_0$. Let $\{T_\alpha : \alpha < \mathrm{cf}(\kappa)\}$ be a subgraph-antichain of $\mathrm{cf}(\kappa)$-sized rayless trees, which exists by the induction hypothesis. Additionally, consider for every $\alpha < \mathrm{cf}(\kappa)$ a subgraph-antichain $\{T_{\alpha,\beta} : \beta < \kappa_\alpha\}$ of κ_α-sized rayless trees, which exists again by the induction hypothesis.

For all $\alpha < \mathrm{cf}(\kappa)$ and $\beta < \kappa_\alpha$, let $T^*_{\alpha,\beta}$ be a tree obtained from the disjoint union of a κ-star with root $r_{\alpha,\beta}$ and $T_{\alpha,\beta}$ and T_α by adding an edge between $r_{\alpha,\beta}$ and an arbitrary vertex of $T_{\alpha,\beta}$ and an edge between $r_{\alpha,\beta}$ and an arbitrary vertex of T_α (see Figure 3.4). We show that $\{T^*_{\alpha,\beta} : \alpha < \mathrm{cf}(\kappa), \beta < \kappa_\alpha\}$ is a subgraph-antichain. Since this antichain has size κ and consists of rayless trees of size κ, the assertion of the theorem then follows from Lemma 3.6.2.

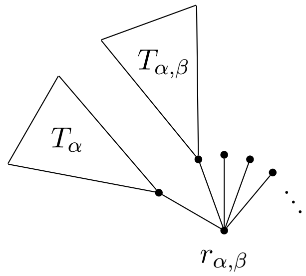

Figure 3.4: The tree $T^*_{\alpha,\beta}$ from the proof of Theorem 3.6.3.

Indeed, consider any $\alpha, \alpha' < \mathrm{cf}(\kappa)$ and $\beta < \kappa_\alpha$ and $\beta' < \kappa_{\alpha'}$ with $(\alpha, \beta) \neq (\alpha', \beta')$ and suppose for a contradiction that there is a subgraph embedding $f : T^*_{\alpha,\beta} \to T^*_{\alpha',\beta'}$. Note that f sends $r_{\alpha,\beta}$ to $r_{\alpha',\beta'}$, since these are the unique vertices of degree κ. Since $T_{\alpha,\beta}$ and $T_{\alpha',\beta'}$ are the only components of $T^*_{\alpha,\beta} - r_{\alpha,\beta}$ and $T^*_{\alpha',\beta'} - r_{\alpha',\beta'}$ of size greater than $\mathrm{cf}(\kappa)$, the map f sends $T_{\alpha,\beta}$ into $T_{\alpha',\beta'}$, and consequently also T_α into $T_{\alpha'}$. However, if $\alpha \neq \alpha'$, then T_α is not a subgraph of $T_{\alpha'}$, and if $\alpha = \alpha'$ but $\beta \neq \beta'$, then $T_{\alpha,\beta}$ is not a subgraph of $T_{\alpha',\beta'}$, a contradiction. □

Using a similar construction as in the proof of Theorem 3.1.12, we prove:

Theorem 3.1.15. *It is consistent with ZFC that for every uncountable cardinal κ there is a tree of size κ which is not a proper subgraph of itself.*

Proof. We assume that there exists a κ-sized subgraph-antichain $\mathcal{U}'$ of κ-sized trees, which is consistent with ZFC by Theorem 3.6.3. Then also $\mathcal{U} := \{U^\vee : U \in \mathcal{U}'\}$ is an antichain by Lemma 3.5.1. Without loss of generality the elements of $\mathcal{U}$ are pairwise disjoint. Note that

(1) no tree in $\mathcal{U}$ contains two adjacent vertices of degree at most 2.

Let T be a κ-regular tree and fix a bijection $f : V(T) \to \mathcal{U}$. Moreover, fix for every $t \in V(T)$ a bijection $g_t : N_T(t) \to L(f(t))$, where $L(f(t))$ denotes the (κ-sized) set of leaves of the tree $f(t)$. We construct a tree U^* from $\bigcup \mathcal{U}$ by adding a $g_t(u)$–$g_u(t)$ edge for every edge $tu \in E(T)$. Thus for every $U \in \mathcal{U}$ and every leaf ℓ of U, there is exactly one edge in the tree U^* connecting ℓ to some leaf of another tree from $\mathcal{U}$.

We call the edges of $E(U^*) \setminus E(\bigcup \mathcal{U})$ *new edges* of U^*. Since new edges only connect leaves of trees from $\mathcal{U}$,

(2) all endvertices of new edges have degree 2 in U^*.

Furthermore, since every vertex of U^* is either a leaf of some $U \in \mathcal{U}$ or separates two leaves of some $U \in \mathcal{U}$,

(3) every vertex of U^* separates two trees from $\mathcal{U}$ in U^*.

Let h be any subgraph embedding of U^* into U^* and consider any $U \in \mathcal{U}$. By (1) and (2), the tree $h(U)$ does not contain any new edges of U^*. Hence there is $U' \in \mathcal{U}$ such that $h(U)$ is completely contained in U'. Since $\mathcal{U}$ is an antichain, we have $U' = U$. In particular, the image of h intersects $U \in \mathcal{U}$, and by the arbitrary choice of U the same is true for all elements of $\mathcal{U}$. Thus by (3), every vertex of U^* must be contained in the image of h, which shows that h is not a proper subgraph embedding. □

Proof of Theorem 3.1.14. Combine Theorems 3.1.15 and 3.5.2. □

3.7 Proof of the main theorems

In this section we prove Theorems 3.1.4, 3.1.8, and 3.1.9. Some of the definitions and lemmas that we need can be found in a similar or identical form in [10], namely Definitions and Lemmas 3.7.3–3.7.8, Definition 3.7.10, Lemma 3.7.11, and Lemma 3.7.13. On the other hand, some of the more involved auxiliary results in this section are new. Lemma 3.7.9 is a strengthening of [10, Lemma 7.3] and its proof requires new ideas, and Lemma 3.7.12 is new.

3.7.1 κ-closure

In this subsection, we prove Lemma 3.7.6 and give the definitions necessary to formulate it. Let G be a graph with a tree-decomposition $(T, \mathcal{V})$.

Definition 3.7.1 ($G(U), \mathring{G}(U)$)**.** Given a subtree U of T, we write $G(U)$ for the subgraph of G induced by $\bigcup\{V_u : u \in V(U)\}$. If the unique T-minimal node of U is a successor of a node t of T, then we write $\mathring{G}(U)$ for the subgraph of G induced by $\bigcup\{V_u : u \in V(U)\} \setminus V_t$. If U consists of a single vertex u, we also write $G(u)$ instead of $G(U)$.

Definition 3.7.2 (hinged)**.** Let $t \in V(T)$ and $s, s' \in \mathrm{succ}_T(t)$. We call a subgraph embedding $f : \mathring{G}(\mathrm{br}_T(s)) \to \mathring{G}(\mathrm{br}_T(s'))$ *hinged (at t)* if any vertex v of $\mathring{G}(\mathrm{br}_T(s))$ has the same neighbourhood in $G(t)$ as $f(v)$.

Definition 3.7.3 (κ-embeddable)**.** Let $(Q, \leq)$ be quasi-ordered. We say that an element $q \in Q$ is *κ-embeddable* in Q with respect to $\leq$ if there exist κ distinct $q' \in Q$ with $q \leq q'$.

Lemma 3.7.4 ([10])**.** *For any well-quasi-ordered set $(Q, \leq)$ and any infinite cardinal κ, the number of elements of Q which are not κ-embeddable in Q with respect to $\leq$ is less than κ.*

Definition 3.7.5 (κ-closed, κ-closure)**.** Let U be a subtree of T and U' a subtree of U. We say that U' is *κ-closed* in U if for all $t \in U'$ and all $s \in \mathrm{succ}_U(t) \setminus V(U')$, the graph $\mathring{G}(\mathrm{br}_T(s))$ is κ-embeddable in $\{\mathring{G}(\mathrm{br}_T(s')) : s' \in \mathrm{succ}_T(t)\}$ with respect to subgraph embeddings hinged at t. Note that if U' is κ-closed in U, then U' is also κ-closed in U'' for any tree U'' with $U' \subseteq U'' \subseteq U$.

The *κ-closure* of U' in U is the smallest subtree of U that is κ-closed in U and contains U'. Such a subtree always exists since the set of all κ-closed subtrees of U with U' as a subtree contains U and is closed under arbitrary intersections. See Figure 3.5 for an example.

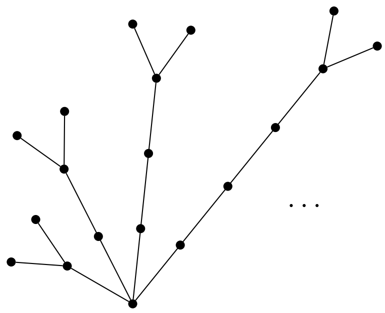

Figure 3.5: Consider the tree S from the figure rooted in its unique vertex of infinite degree with tree-decomposition $(S, \mathcal{W})$ where $W_s = \{t \in V(S) : t \leq s\}$ for all $s \in V(S)$. Then the $\aleph_0$-closure of the root of S in S is S itself. However, when we delete all leaves of S, then $\aleph_0$-closure of the root of S in S is just the root of S.

Lemma 3.7.6. *Let G be a graph with a tree-decomposition $(T, \mathcal{V})$ into finite parts such that the set of induced subgraphs of G is lwqo. Let κ be an infinite cardinal, let U be a subtree of T, and let U' be a subtree of U with $|U'| < \kappa$. If either*

- *$\kappa = \aleph_0$ and T is rayless, or*
- *κ is regular and uncountable,*

then also the κ-closure of U' in U has size less than κ.

Proof. First of all, note that since the parts of $(T, \mathcal{V})$ are finite and the set of induced subgraphs of G is lwqo, for every $t \in V(T)$ the set $\mathcal{S} := \{\mathring{G}(\mathrm{br}_T(s)) : s \in \mathrm{succ}_T(t)\}$ is wqo by hinged subgraph embeddings. Indeed, label the vertices of all graphs from $\mathcal{S}$ with their neighbourhood in the finite set V_t. Since every label-preserving subgraph embedding of one graph from $\mathcal{S}$ into another is hinged and since the elements of $\mathcal{S}$ are lwqo, they are also wqo by hinged subgraph embeddings.

Let $\overline{U'}$ denote the κ-closure of U' in U. Then every vertex u of $\overline{U'}$ has degree less than κ. Indeed, due to its minimality, $\overline{U'}$ contains precisely all successors s of u in U for which $\mathring{G}(\mathrm{br}_T(s))$ is not κ-embeddable in $\{\mathring{G}(\mathrm{br}_T(s')) : s' \in \mathrm{succ}_T(u)\}$ with respect to hinged subgraph embeddings. Since this set is lwqo, u has less than κ children in $\overline{U'}$ by Lemma 3.7.4.

Finally, we deduce that $|\overline{U'}| < \kappa$. If κ is regular and uncountable, then $|\overline{U'}| < \kappa$ since every tree on κ nodes contains a vertex of degree κ. Now suppose that $\kappa = \aleph_0$ and that T is rayless. Then $\overline{U'}$ is a rayless, locally finite graph and hence $|\overline{U'}| < \aleph_0$ by König's infinity lemma (see [26, Proposition 8.2.1]). □

3.7.2 Hordes

For notational convenience, we fix for this subsection:

- a graph G such that the set of induced subgraphs of G is lwqo, and
- a tree-decomposition $(T, \mathcal{V})$ of G into finite parts.

Definition 3.7.7 (suitable)**.** Let U be a subtree of T and let Γ be any graph. We call a subgraph embedding $f : G(U) \to \Gamma$ *suitable at a node* $u \in U$ if there is a subgraph embedding $g : G(\mathrm{br}_T(u)) \to \Gamma$ such that f and g coincide on $G(u)$. We say that f is *suitable* if f is suitable at every $u \in U$.

Definition 3.7.8 (hordes)**.** Let U be a subtree of T and let Γ be any graph. We call $(f_i)_{i \in I}$ a *U-horde* (in Γ) if:

- f_i is a subgraph embedding of $G(U)$ into Γ for all $i \in I$,
- the images of f_i for $i \in I$ are pairwise disjoint, and
- f_i is suitable for all $i \in I$.

Note that if $(f_i : i \in I)$ is a T-horde in Γ, then Γ contains $|I|$ disjoint copies of G. We say that a U-horde $(f_i : i \in I)$ *extends* a U'-horde $(f'_i : i \in J)$ if U' is a subtree of U and $J \subseteq I$ and $f_i \restriction G(U') = f'_i$ for all $i \in J$.

An ordinal-indexed sequence $(A_\alpha : \alpha \leq \gamma)$ of sets (or trees) is called *increasing* if $A_\alpha \subseteq A_\beta$ for all $\alpha < \beta \leq \gamma$, and *continuous* if $A_\beta = \bigcup\{A_\alpha : \alpha < \beta\}$ for all limit ordinals $\beta \leq \gamma$.

Lemma 3.7.9 (regular hordes)**.** *Let κ be an infinite cardinal such that one of the following holds:*

- *$\kappa = \aleph_0$ and T is rayless, or*
- *κ is regular and uncountable.*

Let U be a rooted subtree of T with $|U| \leq \kappa$ and let U' be a κ-closed rooted subtree of U with $|U'| < \kappa$. Furthermore, let I be any set with $|I| \leq \kappa$ and let $J \subseteq I$ with $|J| < \kappa$. Let Γ be any graph and suppose that one of the following holds:

(1) $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ with $|X| < \kappa$, or

(2) $I = J$.

Then any U'-horde $(f'_i : i \in J)$ in Γ can be extended to a U-horde $(f_i : i \in I)$ in Γ.

We will refer to this lemma as either Lemma 3.7.9 (1) or Lemma 3.7.9 (2), depending on which of the two conditions is used. Note that Lemma 3.7.9 (1) implies that if $|T| \leq \kappa$ for T and κ as in the lemma, then G has the κ-EPP. The more general statement of Lemma 3.7.9 will be useful in the proofs of later lemmas.

Proof. Let $\mu := |V(U) \times I| \leq \kappa$ and enumerate $V(U) \times I = \{(u_\alpha, i_\alpha) : \alpha < \mu\}$ in such a way that for all $\alpha \leq \mu$ and $i \in I$, the set $\{u_\beta : \beta < \alpha, i_\beta = i\}$ induces a rooted subtree S_i^α of U. Let U_i^α be the κ-closure of $S_i^\alpha \cup U'$ in U for all $i \in J$ and let U_i^α be the κ-closure of S_i^α in U for all $i \in I \setminus J$. Since $|S_i^\alpha| < \kappa$ and $|U'| < \kappa$, we have $|U_i^\alpha| < \kappa$ in both cases by Lemma 3.7.6. Also note that for all $\alpha < \mu$, there are less than κ many $i \in I$ for which $U_i^\alpha \neq \varnothing$.

First, we establish that the sequence $(U_i^\alpha : \alpha \leq \mu)$ is continuous for all $i \in I$, i.e. the tree $U'^\alpha_i := \bigcup\{U_i^\beta : \beta < \alpha\}$ coincides with U_i^α for all limit ordinals $\alpha \leq \mu$. We assume that $i \in I \setminus J$; the proof for $i \in J$ is similar. To see that U'^α_i is κ-closed in U, consider nodes $t \in U'^\alpha_i$ and $s \in \mathrm{succ}_U(t) \setminus U'^\alpha_i$. Then there is $\beta < \alpha$ such that $t \in U_i^\beta$. Since U_i^β is κ-closed and $s \in \mathrm{succ}_U(t) \setminus U_i^\beta$, the graph $\mathring{G}(\mathrm{br}_T(s))$ is κ-embeddable in $\{\mathring{G}(\mathrm{br}_T(s')) : s' \in \mathrm{succ}_T(t)\}$ with respect to subgraph embeddings hinged at t. This shows that also U'^α_i is κ-closed in U. Next, we show that U'^α_i is the smallest κ-closed tree containing S_i^α. If there was a κ-closed tree S containing S_i^α with $U'^\alpha_i \not\subseteq S$, then there would be a $\beta < \alpha$ such that $U_i^\beta \not\subseteq S$. Since the intersection of κ-closed trees is again κ-closed, $S \cap U_i^\beta$ would be a κ-closed tree containing S_i^β, contradicting the minimality of U_i^β. Thus U'^α_i is the smallest κ-closed tree containing S_i^α, which proves that $U'^\alpha_i = U_i^\alpha$.

Now we recursively define subgraph embeddings $f_i^\alpha : G(U_i^\alpha) \to \Gamma$ for $i \in I$ and $\alpha \leq \mu$ such that:

- $f_i^\alpha \restriction G(U_i^\beta) = f_i^\beta$ for all $i \in I$ and $\beta < \alpha \leq \mu$,
- $f_i^\alpha \restriction G(U') = f_i'$ for all $i \in J$ and $\alpha \leq \mu$,
- the images of f_i^α and f_j^α are disjoint for all $i \neq j \in I$ and $\alpha \leq \mu$,
- f_i^α is suitable for all $i \in I$ and $\alpha \leq \mu$.

Then $(f_i^\mu : i \in I)$ will be the desired U-horde extending $(f_i' : i \in J)$.

We have $S_i^0 = \varnothing$ and thus $U_i^0 = \varnothing$ for all $i \in I \setminus J$, and we have $U_i^0 = U'$ for all $i \in J$ since U' is κ-closed. We let f_i^0 be the empty map for all $i \in I \setminus J$ and define $f_i^0 := f_i'$ for all $i \in J$. Now consider any $\alpha \leq \mu$ and suppose that we have already defined f_i^β for all $\beta < \alpha$ and $i \in I$. If α is a limit, then for all $i \in I$, let $f_i^\alpha : G(U_i^\alpha) \to \Gamma$ be the subgraph embedding extending f_i^β for all $\beta < \alpha$, which exists and is unique since $U_i^\alpha = \bigcup\{U_i^\beta : \beta < \alpha\}$. It is straightforward to check that f_i^α satisfies the four properties listed above.

Finally, suppose that $\alpha = \beta + 1$ is a successor and that f_i^β has been defined for all $i \in I$. For all $i \neq \ell := i_\beta$, we have $S_i^\alpha = S_i^\beta$ and thus $U_i^\alpha = U_i^\beta$ and we define $f_i^\alpha := f_i^\beta$. It is left to define f_ℓ^α.

First suppose that $U_\ell^\beta = \varnothing$. Note that U' is non-empty since it contains the root of T, so as $U_\ell^\beta = \varnothing$ we have $\ell \in I \setminus J$. In particular we have $I \neq J$ and thus we may assume that (1) holds. Consider the set $X := \bigcup\{f_i^\beta(G(U_i^\beta)) : i \in I\}$. Since $|U_i^\beta| < \kappa$ for all $i \in I$, since the parts of $(T, \mathcal{V})$ are finite, and since there are less than κ many $i \in I$ for which $U_i^\beta \neq \varnothing$, it follows from regularity of κ that $|X| < \kappa$. Therefore, by (1) there is a subgraph embedding $f : G \to \Gamma$ such that $f(G)$ avoids X. We define $f_\ell^\alpha := f \restriction G(U_\ell^\alpha)$, which clearly satisfies all prerequisites.

Now suppose that $U_\ell^\beta \neq \emptyset$. If $u_\beta \in U_\ell^\beta$, then $U_\ell^\alpha = U_\ell^\beta$ and we define $f_\ell^\alpha := f_\ell^\beta$. Otherwise, u_β must be an immediate successor of a node t of $S_\ell^\beta \subseteq U_\ell^\beta$. Note that by minimality of the κ-closure, we have $U_\ell^\alpha = U_\ell^\beta \cup \mathrm{br}_{U_\ell^\alpha}(u_\beta)$. We set $f_\ell^\alpha(v) := f_\ell^\beta(v)$ for all $v \in G(U_\ell^\beta)$, so it remains to define f_ℓ^α for the vertices of $\mathring{G}(\mathrm{br}_{U_\ell^\alpha}(u_\beta))$.

Since f_ℓ^β is suitable, there is a subgraph embedding $g : G(\mathrm{br}_T(t)) \to \Gamma$ such that g and f_ℓ^β coincide on $G(t)$. Moreover, since U_ℓ^β is κ-closed, the graph $\mathring{G}(\mathrm{br}_T(u_\beta))$ is κ-embeddable in $\{\mathring{G}(\mathrm{br}_T(s)) : s \in \mathrm{succ}_T(t)\}$ with respect to subgraph embeddings hinged at t. This means that there exists a κ-sized set $B \subseteq \mathrm{succ}_T(t)$ such that for all $b \in B$ there is a hinged subgraph embedding $h_b : \mathring{G}(\mathrm{br}_T(u_\beta)) \to \mathring{G}(\mathrm{br}_T(b))$. We want to find an element $b^* \in B$ such that the image of $g \circ h_{b^*} : \mathring{G}(\mathrm{br}_T(u_\beta)) \to \Gamma$ is disjoint from the set $\bigcup\{f_i^\beta(G(U_i^\beta)) : i \in I\}$, which has size less than κ as we have seen before. As the graphs $\mathring{G}(\mathrm{br}_T(b))$ for $b \in B$ are pairwise disjoint (which follows from the definition of a tree-decomposition) and $|B| = \kappa$, we can indeed find a $b^* \in B$ which is as required. We define $f_\ell^\alpha(v) := g(h_{b^*}(v))$ for all $v \in \mathring{G}(\mathrm{br}_{U_\ell^\alpha}(u_\beta))$, which completes the construction.

To see that f_ℓ^α is a subgraph embedding, first note that the set $f_\ell^\alpha(\mathring{G}(\mathrm{br}_{U_\ell^\alpha}(u_\beta))) = g(h_{b^*}(\mathring{G}(\mathrm{br}_{U_\ell^\alpha}(u_\beta))))$ is disjoint from $f_\ell^\alpha(G(U_\ell^\beta))$ because $f_\ell^\alpha(G(U_\ell^\beta)) = f_\ell^\beta(G(U_\ell^\beta)) \subseteq \bigcup\{f_i^\beta(G(U_i^\beta)) : i \in I\}$. So $f_\ell^\alpha \restriction G(U_\ell^\beta)$ and $f_\ell^\alpha \restriction \mathring{G}(\mathrm{br}_{U_\ell^\alpha}(u_\beta))$ are subgraph embeddings with disjoint images. Thus it remains to show that for every edge vw of G with $v \in G(U_\ell^\beta)$ and $w \in \mathring{G}(\mathrm{br}_{U_\ell^\alpha}(u_\beta))$, there is a $f_\ell^\alpha(v)$–$f_\ell^\alpha(w)$ edge[1] in Γ. Since $v \in V_t$ by the definition of a tree-decomposition and since h_{b*} is hinged at t, there is an edge in G between v and $h_{b^*}(w)$. Since g is a subgraph embedding, there is also an edge in Γ between $g(v) = f_\ell^\beta(v) = f_\ell^\alpha(v)$ and $g(h_{b^*}(w)) = f_\ell^\alpha(w)$.

To see that f_ℓ^α is suitable, consider any $u \in U_\ell^\alpha$. If $u \in U_\ell^\beta$, then f_ℓ^α is suitable at u since f_ℓ^β is. Otherwise, u is contained in $\mathrm{br}_{U_\ell^\alpha}(u_\beta)$ and suitability at u is witnessed by $g \circ h_{b^*} \restriction G(\mathrm{br}_T(u))$. □

To deal with trees of singular cardinality, we use the following definition and lemma from [10]:

Definition 3.7.10 (U-representation)**.** Let U be any tree and write $\eta := \mathrm{cf}(|U|)$. A *U-representation* is an increasing continuous sequence $(U_\alpha : \alpha \leq \eta)$ of subtrees of U such that:

- $U_\eta = U$,
- $|U_\alpha| < |U|$ for all $\alpha < \eta$, and

[1] In the proof of Theorem 3.8.1, we instead need to find $f_\ell^\alpha(v)$–$f_\ell^\alpha(w)$ paths in Γ such that their sets of inner vertices for distinct edges vw are disjoint from each other and from $G(U_\ell^\beta)$ and $\mathring{G}(\mathrm{br}_T(u_\beta))$. The arguments to achieve this are similar to the arguments used in this proof.

- U_α is $|U_\alpha|^+$-closed in U for all $\alpha < \eta$.

We say that the U-representation *extends* a subtree U' of U if $U' \subseteq U_0$.

Lemma 3.7.11 ([10] Lemma 7.5)**.** *For every tree U of singular cardinality and every subtree U' of U with $|U'| < |U|$, there is a U-representation extending U'.*

Let κ be a regular uncountable cardinal and Γ a graph such that $\Gamma - X$ contains a copy of G for all $X \subseteq \Gamma$ of size less than κ. In Lemma 3.7.9 (1), we have seen that for every set I of size at most κ and every rooted subtree U of T of size at most κ, there is a U-horde $(f_i : i \in I)$ in Γ. In the following lemma, we show that the same is true for all (not necessarily regular) uncountable cardinals κ. To make the inductive proof of the lemma work, we prove the following stronger statement:

Lemma 3.7.12 (uncountable hordes)**.** *Let μ and κ be uncountable cardinals such that μ is regular and $\mu \leq \kappa$. Let Γ be any graph such that $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ with $|X| < \kappa$. Let U be a rooted subtree of T with $|U| \leq \kappa$ and let U' be a μ-closed rooted subtree of U with $|U'| < \mu$. Let I be any set with $|I| \leq \kappa$ and let J be a subset of I with $|J| < \mu$. Then any U'-horde $(f'_i : i \in J)$ in Γ can be extended to a U-horde $(f_i : i \in I)$.*

Proof. We prove the lemma by induction on $\max(|U|, |I|)$ and distinguish between four cases.

Case 1. $\max(|U|, |I|) \leq \mu$ (base case).

We can extend $(f'_i : i \in J)$ to a U-horde $(f_i : i \in I)$ by Lemma 3.7.9 (1) (where we insert μ for the variable κ from Lemma 3.7.9) since

- μ is regular and uncountable,
- U' is μ-closed in U,
- $|U'| < \mu$,
- $|U|, |I| \leq \mu$,
- $|J| < \mu$, and
- $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ of size less than μ (as $\mu \leq \kappa$).

Case 2. $\max(|U|, |I|) > \mu$ and $\max(|U|, |I|)$ is regular.

Let U'' be $\max(|U|, |I|)$-closure of U' in U. Since $|U'| < \mu < \max(|U|, |I|)$ and $\max(|U|, |I|)$ is regular and uncountable, also $|U''| < \max(|U|, |I|)$ by Lemma 3.7.6. Moreover, we have $|J| < \mu < \max(|U|, |I|)$ and thus $\max(|U''|, |J|) < \max(|U|, |I|)$. Therefore, we can apply the induction hypothesis to extend $(f'_i : i \in J)$ to a U''-horde $(f''_i : i \in J)$. We can further extend $(f''_i : i \in J)$ to a U-horde $(f_i : i \in I)$ by Lemma 3.7.9 (1) (where we insert $\max(|U|, |I|)$ for the variable κ from Lemma 3.7.9) since

- $\max(|U|, |I|)$ is regular and uncountable (as $\max(|U|, |I|) > \mu$),
- U'' is $\max(|U|, |I|)$-closed in U,
- $|U''| < \max(|U|, |I|)$,
- $|J| < \max(|U|, |I|)$,
- $|U|, |I| \leq \max(|U|, |I|)$, and
- $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ of size less than $\max(|U|, |I|)$ (as $\max(|U|, |I|) \leq \kappa$).

Case 3. $\max(|U|, |I|) > \mu$, $\max(|U|, |I|)$ is singular, and $|U| < |I|$.

Write $\eta := \mathrm{cf}(|I|)$. Since $|I| = \max(|U|, |I|)$ is singular and therefore a limit cardinal and since $|J| < \mu < \max(|U|, |I|) = |I|$, we can find an increasing continuous sequence $(I_\alpha : \alpha \leq \eta)$ of infinite subsets of I such that:

- $|U| \leq |I_0|$,
- $J \subseteq I_0$,
- $I_\eta = I$, and
- $|I_\alpha| < |I|$ for all $\alpha < \eta$.

We recursively construct U-hordes $(f_i : i \in I_\alpha)$ extending $(f'_i : i \in J)$ and each other. To start the recursion, we extend $(f'_i : i \in J)$ to a U-horde $(f_i : i \in I_0)$ by the induction hypothesis, which is possible since $\max(|U|, |I_0|) = |I_0| < |I| = \max(|U|, |I|)$. For successor steps, let $\alpha < \eta$ and suppose that the U-horde $(f_i : i \in I_\alpha)$ has already been defined. We have $\max(U, I_{\alpha+1}) = |I_{\alpha+1}| < |I| = \max(|U|, |I|)$. Therefore, we can extend $(f_i : i \in I_\alpha)$ to a U-horde $(f_i : i \in I_{\alpha+1})$ by the induction hypothesis applied with $\mu = |I_\alpha|^+$ since

- $|I_\alpha|^+$ is regular and uncountable (as I_α is infinite),
- $|I_\alpha|^+ \leq |I| \leq \kappa$,
- U is $|I_\alpha|^+$-closed in U, and
- $|U| \leq |I_0| \leq |I_\alpha| < |I_\alpha|^+$.

If $\alpha \leq \eta$ is a limit, let $(f_i : i \in I_\alpha)$ be the unique U-horde that extends $(f_i : i \in I_\beta)$ for all $\beta < \alpha$, which exists by continuity of the sequence $(I_\alpha : \alpha \leq \eta)$.

Case 4. $\max(|U|, |I|) > \mu$, $\max(|U|, |I|)$ is singular, and $|U| \geq |I|$.

Write $\eta := \mathrm{cf}(|U|)$. Since $|U'| < \mu < \max(|U|, |I|) = |U|$, there is a U-presentation $(U_\alpha : \alpha \leq \eta)$ extending U' by Lemma 3.7.11. By considering a subsequence of the U-representation, we assume without loss of generality that U_0 is infinite. Since $|J| < \mu < \max(|U|, |I|) = |U|$, we may also assume that $|J| \leq |U_0|$. Since $|I| \leq |U|$, we can find a (not necessarily strictly) increasing continuous sequence $(I_\alpha : \alpha \leq \eta)$ of subsets of I such that

- $I_0 = J$,
- $I_\eta = I$, and
- $|I_\alpha| \leq |U_\alpha|$ for all $\alpha \leq \eta$.

We recursively construct U_α-hordes $(f_i^\alpha : i \in I_\alpha)$ extending $(f_i' : i \in J)$ and each other. To start the recursion, we extend $(f_i' : i \in J)$ to a U_0-horde $(f_i^0 : i \in I_0)$ by the induction hypothesis, which is possible since $\max(|U_0|, |I_0|) = |U_0| < |U| = \max(|U|, |I|)$. For successor steps, let $\alpha < \eta$ and suppose that the U_α-horde $(f_i^\alpha : i \in I_\alpha)$ has already been defined. We have $\max(|U_{\alpha+1}|, |I_{\alpha+1}|) = |U_{\alpha+1}| < |U| = \max(|U|, |I|)$. Therefore, we can extend $(f_i^\alpha : i \in I_\alpha)$ to a $U_{\alpha+1}$-horde $(f_i^{\alpha+1} : i \in I_{\alpha+1})$ by the induction hypothesis applied with $\mu = |U_\alpha|^+$ since

- $|U_\alpha|^+$ is regular and uncountable (as $|U_\alpha|$ is infinite),
- $|U_\alpha|^+ \leq |U| \leq \kappa$,
- $|I_{\alpha+1}| \leq |U_{\alpha+1}| \leq |U| \leq \kappa$,
- U_α is $|U_\alpha|^+$-closed in U and thus in $U_{\alpha+1}$ by the definition of a U-representation,
- $|U_\alpha| < |U_\alpha|^+$, and
- $|I_\alpha| \leq |U_\alpha| < |U_\alpha|^+$.

If $\alpha \leq \eta$ is a limit, then let $(f_i^\alpha : i \in I_\alpha)$ be the unique U_α-horde extending $(f_i^\beta : i \in I_\beta)$ for all $\beta < \alpha$, which exists by continuity of $(I_\alpha : \alpha \leq \eta)$ and $(U_\alpha : \alpha \leq \eta)$. □

Our final lemma allows us to extend certain U-hordes, which we will find using Lemma 3.7.9 or Lemma 3.7.12, to T-hordes:

Lemma 3.7.13 (for extending hordes)**.** *Let μ be an infinite cardinal such that one of the following holds:*

- *$\mu = \aleph_0$ and T is rayless, or*
- *μ is regular and uncountable.*

Let U be a rooted subtree of T, and U' a μ-closed rooted subtree of U of size less than μ. Then any U'-horde $(f_i' : i \in I)$ with $|I| < \mu$ in a graph Γ can be extended to a U-horde $(f_i : i \in I)$ in Γ.

Proof. By induction on $|U|$.

Case 1. $|U| \leq \mu$ (base case).

Then the assertion follows from Lemma 3.7.9 (2).

Case 2. $|U| > \mu$ and $|U|$ is regular.

Then $|U'| < \mu < |U|$, so also the $|U|$-closure U'' of U' in U has size less than $|U|$ by Lemma 3.7.6. Hence by the induction hypothesis, we can extend $(f'_i : i \in I)$ to a U''-horde $(f''_i : i \in I)$ in Γ. We can further extend $(f''_i : i \in I)$ to a U-horde $(f_i : i \in I)$ in Γ by Lemma 3.7.9 (2) since

- $|U|$ is regular and uncountable,
- U'' is $|U|$-closed in U,
- $|U''| < |U|$, and
- $|J| \leq |I| < \mu < |U|$.

Case 3. $|U| > \mu$ and $|U|$ is singular.

Write $\eta := \mathrm{cf}(|U|)$. Since $|U'| < \mu < |U|$, there is a U-representation $(U_\alpha : \alpha \leq \eta)$ extending U' by Lemma 3.7.11. Without loss of generality, suppose that $\mu \leq |U_0|$. We recursively construct U_α-hordes $(f_i^\alpha : i \in I)$ for $\alpha < \eta$ extending $(f'_i : i \in I)$ and each other. To start the recursion, we extend $(f'_i : i \in I)$ to a U_0-horde $(f_i^0 : i \in I)$ by the induction hypothesis. For successor steps, let $\alpha < \eta$ and extend the U_α-horde $(f_i^\alpha : i \in I)$ to a $U_{\alpha+1}$-horde $(f_i^{\alpha+1} : i \in I)$ by the induction hypothesis applied with $\mu = |U_\alpha|^+$. This is possible because

- $|U_\alpha|^+$ is regular and uncountable,
- U_α is $|U_\alpha|^+$-closed in U and thus in $U_{\alpha+1}$ by the definition of a U-representation,
- $|U_\alpha| < |U_\alpha|^+$, and
- $|I| < \mu \leq |U_0| < |U_\alpha|^+$.

If α is a limit, then for all $i \in I$ let $f_i^\alpha : G(U_\alpha) \to \Gamma$ be the unique function extending f_i^β for all $\beta < \alpha$, which completes the construction. □

3.7.3 Deducing the main theorems and their corollaries

Theorem 3.1.9. *Let G be any graph admitting a tree-decomposition into finite parts such that the set of induced subgraphs of G is lwqo. Then G has the κ-EPP for every uncountable cardinal κ.*

Proof. Let $(T, \mathcal{V})$ be a tree-decomposition of G into finite parts, let κ be any uncountable cardinal, and let Γ be a graph such that $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ of size less than κ. Moreover, fix any set I of size κ. We show that Γ contains a T-horde $(f_i : i \in I)$; then the graphs $f_i(G)$ for $i \in I$ are κ disjoint copies of G in Γ and the proof is complete.

Let U denote the κ^+-closure in T of the root of T and note that $|U| < \kappa^+$ by Lemma 3.7.6. By Lemma 3.7.12 (with $J = \varnothing$) there is a U-horde $(f'_i : i \in I)$ in Γ. By Lemma 3.7.13 (with $\mu = \kappa^+$) this horde can be extended to a T-horde $(f_i : i \in I)$ in Γ. □

Theorem 3.1.8. *Let G be any rayless graph such that the set of induced subgraphs of G is lwqo. Then G has the $\aleph_0$-EPP.*

Proof. Let Γ be a graph such that $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ of size less than $\aleph_0$. Moreover, fix any set I of size $\aleph_0$. Since G is rayless, by Lemma 3.2.1 there is a tree-decomposition $(T, \mathcal{V})$ of G into finite parts such that T is rayless. We show that Γ contains a T-horde $(f_i : i \in I)$; then the graphs $f_i(G)$ for $i \in I$ are $\aleph_0$ disjoint copies of G in Γ and the proof is complete.

Let U denote the $\aleph_1$-closure in T of the root of T and note that $|U| < \aleph_1$ by Lemma 3.7.6. By Lemma 3.7.9 (1) (with $\kappa = \aleph_0$ and $J = \varnothing$) there is a U-horde $(f'_i : i \in I)$ in Γ. By Lemma 3.7.13 (with $\mu = \aleph_1$) this horde can be extended to a T-horde $(f_i : i \in I)$ in Γ. □

Theorem 3.1.4. *Let G be any rayless graph such that the set of induced subgraphs of G is lwqo. Then G has the EPP.*

Proof. By Lemma 3.2.1 there is a tree-decomposition $(T, \mathcal{V})$ of G into finite parts such that T is rayless. Let U denote the $\aleph_0$-closure in T of the root of T and note that U is finite by Lemma 3.7.6. Consider the function $f : \mathbb{N} \to \mathbb{N}, k \mapsto |G(U)| \cdot (k-1)$, let $k \in \mathbb{N}$, and let Γ be a graph such that $\Gamma - X$ contains a copy of G for all $X \subseteq V(\Gamma)$ of size at most $f(k)$. We show that Γ contains a T-horde $(f_i : i \in \{1, \dots, k\})$; then Γ contains k disjoint copies of G and the proof is complete.

We begin by finding a U-horde $(f'_i : i \in \{1, \dots, k\})$ in Γ. We define f'_i for $i \leq k$ recursively, so let $i \leq k$ and suppose that $f'_j : G(U) \to \Gamma$ has been defined for all $j < i$. Since the set $X := \bigcup \{f_j(G(U)) : j < i\}$ has size at most $f(k)$, there is a subgraph embedding $g : G \to \Gamma$ whose image avoids X. We set $f'_i := g \restriction G(U)$, completing the recursive construction.

By Lemma 3.7.13 (with $\mu = \aleph_0$), we can extend the U-horde $(f'_i : i \in \{1, \dots, k\})$ to the desired T-horde $(f_i : i \in \{1, \dots, k\})$ in Γ. □

Finally, we deduce Corollaries 3.1.5 and 3.1.10 using the following result of Jia [45]:

Theorem 3.7.14. *For every $n \in \mathbb{N}$, the class of all graphs that do not contain a path of length n as a subgraph is lwqo.*

Corollary 3.1.5. *Let G be a graph that does not contain a path of length n as a subgraph for some $n \in \mathbb{N}$. Then G has the EPP.*

Proof. Since no (induced) subgraph of G contains a path of length n as a subgraph, the set of induced subgraphs of G is lwqo by Theorem 3.7.14. Thus G has the EPP by Theorem 3.1.4. □

Corollary 3.1.10. *Let G be a graph that does not contain a path of length n as a subgraph for some $n \in \mathbb{N}$. Then G has the κ-EPP for every infinite cardinal κ.*

Proof. Since no (induced) subgraph of G contains a path of length n as a subgraph, the set of induced subgraphs of G is lwqo by Theorem 3.7.14. If $\kappa = \aleph_0$, then G has the κ-EPP by Theorem 3.1.8. If κ is uncountable, we apply Lemma 3.2.1 to find a tree-decomposition of G into finite parts. Then G has the κ-EPP by Lemma 3.1.9. □

3.8 Classes defined by topological minors

A *topological minor embedding* of a graph H into a graph G is a map f with domain $V(H) \cup E(H)$ that maps every vertex of H to a vertex of G and every edge of H to a path in G such that:

- $f \restriction V(H)$ is injective,
- $f(vw)$ is an $f(v)$–$f(w)$ path in G for every edge $vw \in E(H)$,
- for all $e \in E(H)$, the set of inner vertices of the path $f(e)$ does not contain any vertex $f(v)$ for $v \in V(H)$ or any inner vertex of a path $f(e')$ for $e \neq e' \in E(H)$.

Let L be any finite set and suppose that the vertices of the graphs H and G are labelled by elements of L. We say that f is *label-preserving* if every vertex v of H has the same label as $f(v)$. Furthermore, we say that a class $\mathcal{G}$ of graphs is *lwqo by topological minors*, if for every finite set L, the class of all graphs from $\mathcal{G}$ with vertices labelled by elements of L is wqo by label-preserving topological minor embeddings.

With this definition, we can state the following version of Theorems 3.1.8 and 3.1.9 for topological minors:

Theorem 3.8.1. *Let G be any graph such that the set of induced subgraphs of G is lwqo by topological minors.*

- *If G is rayless, then $\mathcal{T}(G)$ has the $\aleph_0$-EPP.*
- *If G admits a tree-decomposition into finite parts, then $\mathcal{T}(G)$ has the κ-EPP for every uncountable cardinal κ.*

Proof outline. The proof is essentially the same as the proof of Theorems 3.1.8 and 3.1.9 given in Section 3.7, with the difference that all subgraph embeddings must be replaced by topological minor embeddings throughout Section 3.7. Additionally, some statements and some notations concerning subgraph embeddings must be transferred to the setting of topological minor embeddings in a natural way. The notational changes include:

- Let f be a topological minor embedding of a graph G into a graph H and let G' be a subgraph of G. We write $f(G')$ for the subgraph of H consisting of all vertices $f(v)$ for $v \in V(G')$ and all paths $f(e)$ for $e \in E(G')$. The *image* of f is the subgraph $f(G)$ of H.
- Let f be a topological minor embedding of a graph G into a graph H and f' a topological minor embedding of H into a graph I. The *concatenation* $f' \circ f$ is the topological minor embedding g of G into I such that:
  - $g(v) = f'(f(v))$ for all $v \in V(G)$, and
  - $g(e)$ for $e \in E(G)$ is the path in I obtained by concatenating all paths $f'(e')$ for $e' \in E(f(e))$. □

Finally, we deduce Corollary 3.1.19. We need the following theorem by Laver [58]:

Theorem 3.8.2. *The class of all trees is lwqo by topological minors.*

Corollary 3.1.19. *$\mathcal{T}(T)$ has the κ-EPP for every uncountable cardinal κ and every tree T. If T is rayless, then $\mathcal{T}(T)$ also has the $\aleph_0$-EPP.*

Proof. Combine Theorems 3.8.1 and 3.8.2. □

Chapter 4

Ubiquity of oriented rays

4.1 Introduction

A (di)graph H is called $\trianglelefteq$-*ubiquitous* for a binary (di)graph relation $\trianglelefteq$ if any (di)graph G that contains k disjoint copies of H for every $k \in \mathbb{N}$ also contains infinitely many disjoint copies of H with respect to $\trianglelefteq$. Possible relations for $\trianglelefteq$ are e.g. the subgraph, topological minor or minor relation for graphs or the subdigraph relation for digraphs.

Halin started the investigation of ubiquity in graphs with his landmark result that rays are subgraph-ubiquitous in [38]. Andreae conjectured that every locally finite connected graph is minor-ubiquitous after studying minor-ubiquity in [2, 3]. Noteworthy progress towards this conjecture was recently achieved by Bowler, Elbracht, Erde, Gollin, Heuer, Pitz and Teegen in a series of papers [10–12], in which they proved, among several other results, that all trees are topological-minor-ubiquitous. Throughout the years several results proving and disproving the ubiquity of certain graphs have been published, including results concerning different notions of ubiquity as in [9, 55].

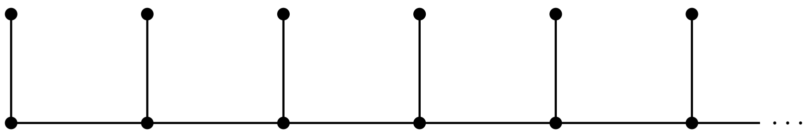

Figure 4.1: The comb, a graph that is not subgraph-ubiquitous.

An example for a graph that is not subgraph-ubiquitous is the comb [2] (see Figure 4.1). For graphs that are not topological-minor-ubiquitous, see [3]. Very recently, Carmesin provided an example of a locally finite graph that is not minor-ubiquitous in [20]. We remark that this does not contradict Andreae's conjecture as the graph is not connected.

In [9] Bowler, Carmesin and Pott first suggested the topic of ubiquity in digraphs by asking whether any digraph containing arbitrarily many edge-disjoint directed double rays also contains infinitely many of them.

We take on the quest of investigating ubiquity in digraphs. We characterise which *oriented rays*, i.e. digraphs whose underlying undirected graphs are rays, are ubiquitous regarding the subdigraph relation. Whenever we write *ubiquitous* without specifying the relation, we refer to the subdigraph relation. Furthermore, we call a vertex of an oriented ray a *turn* if it has in-degree or out-degree 2. Our main result reads as follows:

Theorem 4.1.1. *An oriented ray is ubiquitous if and only if it has a finite number of turns.*

We develop novel methods that enhance the common techniques in the field of ubiquity theory. For the backward implication of Theorem 4.1.1 (see Section 4.3) we follow the proof for Halin's ray ubiquity result in undirected graphs given in [26, Theorem 8.2.5 (i)]. To make this possible we require sets of arbitrarily many disjoint copies of an oriented ray that have an additional property, they must be forked. The existence of such forked sets, which we prove in Lemma 4.3.3, is our key contribution in the proof of the forward implication.

For the forward implication of Theorem 4.1.1 (see Section 4.4) we construct a counterexample from infinitely many disjoint copies of an oriented ray by identifying vertices. In the proof of Theorem 4.4.4 we extend the common technique of identifying vertices by a recursive choice of the vertices which will be identified.

Further results and open problems on ubiquity in digraphs are presented in Chapter 5.

4.2 Preliminaries

For $n \in \mathbb{N}$ we denote $[n] := \{1, \dots, n\}$. From now on we write *ray* instead of oriented ray. Similarly, a *path* is a digraph whose underlying undirected graph is a path. We call a path a *dipath* if all its arcs are consistently oriented, i.e. each vertex has in- and out-degree at most 1. We say that an arc of a ray R is *in-oriented* if it is directed towards the unique vertex of R with undirected degree 1 and *out-oriented* otherwise. An *in-ray* is a ray in which all arcs are in-oriented and an *out-ray* is a ray in which all arcs are out-oriented. A maximal dipath contained in a ray is called a *phase*. We call a phase of a ray *in-oriented* if its arcs are in-oriented and *out-oriented* otherwise.

A ray with infinitely many turns can be represented by a sequence of natural numbers where the n-th term of the sequence represents the length of the n-th phase. We call this the *representing sequence* of the ray. The representing sequence is *bounded* if there is $b \in \mathbb{N}$ such that all elements of this sequence are contained in $[b]$. Otherwise the representing sequence is called *unbounded*.

For a digraph D and vertices $v, w \in D$, we write $d_D(v, w)$ for the distance between v and w in the underlying undirected graph. Let $R \subseteq D$ be a ray. For any $v \in R$ we write vR for the tail of R starting in the vertex v, and for any $a \in A(R)$ we write aR for the tail of R starting with the arc a. For $w \in D$ we say that $v \in R$ lies *beyond* w *on* R if $w \notin V(vR)$. Two rays R_1, R_2 *traverse an arc* $uv \in A(R_1) \cap A(R_2)$ *in the same direction* if either v lies beyond u on R_1 and R_2 or u lies beyond v on R_1 and R_2. Otherwise R_1 and R_2 *traverse* uv *in opposite directions.*

Let $\overleftarrow{D}$ be the digraph obtained from D by changing the orientation of every arc. To warm up with the definition of ubiquity, we prove the following simple lemma:

Lemma 4.2.1. *A digraph H is ubiquitous if and only if $\overleftarrow{H}$ is ubiquitous.*

Proof. It suffices to show that $\overleftarrow{H}$ is ubiquitous if H is ubiquitous. Let D be any digraph containing arbitrarily many disjoint copies of $\overleftarrow{H}$. Hence $\overleftarrow{D}$ contains arbitrarily many disjoint copies of $\overleftarrow{\overleftarrow{H}} = H$. Then $\overleftarrow{D}$ also contains infinitely many disjoint copies of H since H is ubiquitous. Therefore $\overleftarrow{\overleftarrow{D}} = D$ contains infinitely many disjoint copies of $\overleftarrow{H}$, which proves that $\overleftarrow{H}$ is ubiquitous. □

Corresponding to [10, Definition 5.1], for digraphs D and R we call a collection $\mathcal{F}$ of finite sets of disjoint copies of R in D an *R-tribe in D*. If R is clear from the context we may also just say *tribe* instead, similarly for the containing graph D. Furthermore, for an R-tribe $\mathcal{F}$ in D we call $F \in \mathcal{F}$ a *layer of $\mathcal{F}$*, any element of F a *member of $\mathcal{F}$* and say that $\mathcal{F}$ is *thick* if for each $n \in \mathbb{N}$ there is a layer F of $\mathcal{F}$ with $|F| \geq n$. Note that if D contains arbitrarily many disjoint copies of R, then D contains a thick R-tribe. A tribe $\mathcal{F}'$ in D is an *(R-)subtribe*[1] of an R-tribe $\mathcal{F}$ in D if every layer of $\mathcal{F}'$ is a subset of a layer of $\mathcal{F}$.

Whenever we consider a copy R' of a digraph R we implicitly fix an isomorphism $\varphi\colon R \to R'$ and for a subdigraph $\hat{R} \subseteq R$ we write in short $\hat{R}'$ for $\varphi(\hat{R})$. With this, we say that an R-tribe $\mathcal{F}$ is *forked at* $\hat{R}$ if $\hat{R}' \cap R'' = \varnothing$ for any two distinct members R', R'' of $\mathcal{F}$.

4.3 Positive results

In this section, we prove

Theorem 4.3.1. *A ray is ubiquitous if it has a finite number of turns.*

The proof of Theorem 4.3.1 will be an easy consequence of Theorem 4.3.2 together with Lemma 4.3.3 below. The following theorem is a variant of Halin's ray ubiquity result for digraphs with an additional restriction on the start vertices of the rays. The proof given here is derived from the proof of Halin's result in [26, Theorem 8.2.5 (i)].

Theorem 4.3.2. *Let D be a digraph and R an out-ray. If there exists a thick R-tribe F in D and $X \subseteq V(D)$ such that each member of F has its first vertex in X, then there are infinitely many disjoint out-rays in D whose first vertices are contained in X.*

Proof. We will recursively fix for every $n \in \mathbb{N}^+$ a set $\mathcal{R}^n = \{R_1^n, \dots, R_n^n\}$ of pairwise disjoint out-rays in D and a set of vertices $\{u_1^n, \dots, u_n^n\}$ such that for every $k \in [n]$:

- R_k^n has its first vertex in X,
- $u_k^n \in R_k^n$, and
- $R_k^n u_k^n \subsetneq R_k^{n+1} u_k^{n+1}$.

Then $\{\bigcup_{n \geq k} R_k^n u_k^n \colon k \in \mathbb{N}\}$ is an infinite set of pairwise disjoint out-rays where each of the rays has its first vertex in X.

[1] Note that this definition of subtribe corresponds to the notion of flat subtribe in [10]. The definition of subtribe in [10] however is different and more general.

For $n = 1$ we pick one ray $R_1^1 \in \bigcup \mathcal{F}$, set $\mathcal{R}^1 := \{R_1^1\}$ and pick $u_1^1 \in R_1^1$ arbitrarily. This satisfies the required properties.

Now let $\ell \geq 1$ and suppose that for all $i \in [\ell]$ there are sets $\mathcal{R}^i$ and $\{u_1^i, \dots, u_i^i\}$ subject to the conditions above. Consider a layer F of $\mathcal{F}$ of size at least $|\bigcup_{i\in[\ell]} R_i^\ell u_i^\ell| + \ell^2 + 1$. First, we delete from F every ray that meets a path $R_i^\ell u_i^\ell$ for some $i \in [\ell]$. Then there are still at least $\ell^2 + 1$ rays left in F. Next, we repeatedly check whether there is a ray $R_i^\ell \in \mathcal{R}^\ell$ for which $R_i^{\ell+1}$ has not yet been defined and that meets at most ℓ of the remaining elements in F. If that is the case, we set $R_i^{\ell+1} := R_i^\ell$, choose a vertex $u_i^{\ell+1}$ beyond u_i^ℓ on R_i^ℓ arbitrarily, and delete the at most ℓ many rays from F that have non-empty intersection with $R_i^{\ell+1}$. Suppose that after $m \leq \ell$ many steps, every ray in $\mathcal{R}^\ell$ meets either none or more than ℓ many rays from the reduced F, which we will refer to as F'.

Consider the $(\ell - m)$-sized subset $J \subseteq [\ell]$ containing all $j \in [\ell]$ for which $R_j^{\ell+1}$ has not yet been defined. Then any ray R_j^ℓ with $j \in J$ meets more than ℓ rays from F'. We deleted at most $|\bigcup_{i\in[\ell]} R_i^\ell u_i^\ell|$ rays from F in the first step and at most $m\ell$ in the second step, thus F' has size at least

$$|\bigcup_{i\in[\ell]} R_i^\ell u_i^\ell| + \ell^2 + 1 - |\bigcup_{i\in[\ell]} R_i^\ell u_i^\ell| - m\ell = (\ell - m)\ell + 1.$$

For any ray R_j^ℓ with $j \in J$ we fix the vertex $c_j \in R_j^\ell$ which is the first intersection of R_j^ℓ with the ℓ-th ray from F' that it meets. Note that c_j lies beyond u_j^ℓ on R_j^ℓ. Then $\bigcup_{j\in J} R_j^\ell c_j$ meets at most $|J|\ell = (\ell - m)\ell$ rays from F'. Therefore, there is at least one ray left in F' that is disjoint from $\bigcup_{j\in J} R_j^\ell c_j$ and we pick this ray as $R_{\ell+1}^{\ell+1}$. We choose an arbitrary vertex $u_{\ell+1}^{\ell+1} \in R_{\ell+1}^{\ell+1}$, define $F^* := F' \setminus \{R_{\ell+1}^{\ell+1}\}$, and write $F^* = \{S_i : i \in I\}$ for a suitable index set I.

Now for any ray $S_i \in F^*$ we choose a vertex w_i that lies beyond all vertices of $\bigcup_{j\in J} u_j^\ell R_j^\ell c_j$ on S_i. Consider the finite subdigraph

$$H := \bigcup_{j\in J} u_j^\ell R_j^\ell c_j \cup \bigcup_{i\in I} S_i w_i$$

of D. Additionally, we define $U := \{u_j^\ell : j \in J\}$ and $W := \{w_i : i \in I\}$. We show that for any set $Z \subseteq V(H)$ of fewer than $\ell - m$ vertices there is a U–W dipath in $H - Z$. Indeed, Z misses at least one dipath of the form $u_j^\ell R_j^\ell c_j$ and since there are $\ell \geq \ell - m$ many paths $S_i w_i$ with $u_j^\ell R_j^\ell c_j \cap S_i w_i \neq \emptyset$, at least one such path $S_i w_i$ avoids Z. Let v_j be the first vertex on $u_j^\ell R_j^\ell c_j$ which lies on $S_i w_i$; then $u_j^\ell R_j^\ell v_j S_i w_i$ is a dipath from u_j^ℓ to w_i: firstly, by the choice of v_j the underlying undirected graph of $u_j^\ell R_j^\ell v_j S_i w_i$ clearly is a path. Secondly, the dipath $u_j^\ell R_j^\ell v_j$ is directed from u_j^ℓ to v_j since $v_j \in u_j^\ell R_j^\ell$. Lastly, the dipath $v_j S_i w_i$ is directed from v_j to w_i since v_j lies in $S_i w_i$. Thus by Menger's theorem [5, Theorem 7.3.1], there is a set $\mathcal{P}$ of $\ell - m = |J|$ pairwise disjoint U–W dipaths in H.

For all $j \in J$, we write P_j for the dipath in $\mathcal{P}$ starting at u_j^ℓ. Let $h \colon J \to I$ such that $w_{h(j)}$ is the endvertex of P_j in W. Now we define

$$R_j^{\ell+1} := R_j^\ell u_j^\ell P_j w_{h(j)} S_{h(j)}$$

and $u_j^{\ell+1} := w_{h(j)}$, which clearly fulfils the required properties. □

Lemma 4.3.3. *Let D and H be digraphs and let $\hat{H} \subseteq H$ a finite subdigraph. If there exists a thick H-tribe $\mathcal{E}$ in D, then there is a thick H-subtribe $\mathcal{F}$ of $\mathcal{E}$ in D that is forked at $\hat{H}$.*

Proof. For all $n \in \mathbb{N}$, we recursively define a subset F_n of a layer of $\mathcal{E}$ containing at least n disjoint copies of H in D and a thick subtribe $\mathcal{E}_n$ of $\mathcal{E}$ such that

(i) the H-tribe $\mathcal{F}_n := \{F_0, \ldots, F_n\}$ is forked at $\hat{H}$,

(ii) for each $H_1 \in \bigcup \mathcal{E}_n$ and each $H_2 \in \bigcup \mathcal{F}_n$ the digraph $\hat{H}_1$ is disjoint from H_2 and the digraph $\hat{H}_2$ is disjoint from H_1.

In the end, $\{F_n : n \in \mathbb{N}\}$ will be a thick H-subtribe of $\mathcal{E}$ satisfying the lemma. For the first step we set $F_0 := \varnothing$ and $\mathcal{E}_0 := \mathcal{E}$. Now suppose that $\mathcal{F}_{n-1}$ and $\mathcal{E}_{n-1}$ are already defined. Set $h := |\hat{H}|$ and choose a layer L from $\mathcal{E}_{n-1}$ of size at least $h + n$. We will choose F_n as an n-element subset of L. Then $\mathcal{F}_n$ will be forked at $\hat{H}$ since (i) and (ii) hold for $\mathcal{E}_{n-1}$ and $\mathcal{F}_{n-1}$. Our task is to find a suitable subset F_n of L and a thick subtribe $\mathcal{E}_n$ of $\mathcal{E}$ such that $\bigcup \mathcal{F}_n$ and $\bigcup \mathcal{E}_n$ satisfy (ii). We begin by deleting from each layer $M \neq L$ of $\mathcal{E}_{n-1}$ any element that has non-empty intersection with some $H' \in L$ in its subdigraph $\hat{H}'$. Note that for every digraph $H' \in L$ there are at most $|\hat{H}'| = h$ many digraphs from M which meet $\hat{H}'$. Therefore we delete from every layer of $\mathcal{E}_{n-1}$ at most $h \cdot |L|$ elements and the resulting subtribe $\mathcal{C}$ of $\mathcal{E}_{n-1}$ is still a thick tribe in D.

Claim. For every $j \in \mathbb{N}$ there is a subset $L_j \subseteq L$ with $|L_j| = n$ and a subset C_j with $|C_j| \geq j$ of a layer of $\mathcal{C}$ such that for any $H_1 \in L_j$ and any $H_2 \in C_j$ the digraph H_1 is disjoint from $\hat{H}_2$ and H_2 is disjoint from $\hat{H}_1$.

Proof of the claim. Let $j \in \mathbb{N}$ and C a layer of $\mathcal{C}$ of size at least $j\binom{|L|}{n}$. By the construction of $\mathcal{C}$, we only need to find sets $L_j \subseteq L$ and $C_j \subseteq C$ such that no $H_1 \in L_j$ meets any $H_2 \in C_j$ in its subdigraph $\hat{H}_2$. For every $H' \in C$, at most $|\hat{H}'| = h$ elements of L meet $\hat{H}'$. Since $|L| \geq h + n$, we can choose for every $H' \in C$ a subset of n elements of L such that each of these does not meet $\hat{H}'$. This defines a map $\alpha \colon C \to \mathfrak{L} := \{L' \subseteq L \colon |L'| = n\}$. Since $|C| \geq j\binom{|L|}{n} = j|\mathfrak{L}|$, there is a set $L_j \in \mathfrak{L}$ with $|\alpha^{-1}(L_j)| \geq j$ by pigeon hole principle. Then L_j and $C_j := \alpha^{-1}(L_j)$ are as desired. □

Since L has only finitely many subsets, there is an infinite strictly increasing sequence $(j_k)_{k \in \mathbb{N}}$ in $\mathbb{N}$ such that the sets L_{j_k} coincide for all $k \in \mathbb{N}$. We choose this as the set F_n. By the claim, $\mathcal{E}_n := \{C_{j_k} : k \in \mathbb{N}\}$ is a thick subtribe of $\mathcal{E}_{n-1}$ satisfying (ii). This concludes the proof. □

Proof of Theorem 4.3.1. Let R be a ray with finitely many turns. Let D be a digraph and assume that D contains arbitrarily many disjoint copies of R, then D contains a thick R-tribe $\mathcal{E}$. We show that D contains infinitely many copies of R. Since all but finitely many arcs of R are oriented the same way, we may assume by Lemma 4.2.1 that all but finitely many arcs of R are out-oriented.

Let $\hat{R}$ be the (connected) subdigraph of R that consists precisely of all finite phases of R. By Lemma 4.3.3, there is a thick subtribe $\mathcal{F}$ of $\mathcal{E}$ that is forked at $\hat{R}$. Consider the set X which contains for any $R' \in \bigcup \mathcal{F}$ the first vertex of the out-ray $R' - \hat{R}'$, and the set $Y := \bigcup_{R' \in \bigcup \mathcal{F}} V(\hat{R}')$. By deleting $\hat{R}'$ from each member R' of $\mathcal{F}$, we obtain a thick $(R - \hat{R})$-tribe in $D - Y$. Hence, by Theorem 4.3.2 there exists an infinite family $(R_i)_{i \in \mathbb{N}}$ of disjoint out-rays in $D - Y$ such that each R_i starts in a vertex $r_i \in X$. By definition of X, for all $i \in \mathbb{N}$ there is a member S_i of $\mathcal{F}$ such that r_i is the first vertex of $S_i - \hat{S}_i$. Note that $\hat{S}_i$ and $\hat{S}_j$ are disjoint for $i \neq j$ since $\mathcal{F}$ is forked at $\hat{R}$. Finally, by combining each initial segment $\hat{S}_i$ with the out-ray R_i, we obtain infinitely many disjoint copies of R in D. □

4.4 Negative results

In this section, we construct for every oriented ray R with infinitely many turns a digraph D that contains arbitrarily but not infinitely many disjoint copies of R. The construction of D will differ depending on whether the representing sequence of R is bounded (see Theorem 4.4.2) or unbounded (see Theorem 4.4.4). However, the basic framework for the construction of D will be the same in both proofs as follows:

Let $(R(n,m))_{(n,m)\in I}$ be a family of pairwise disjoint copies of R, where

$$I := \{(n,m) \in \mathbb{N}^2 \colon n \leq m\}.$$

We let

$$D_{-1} := \bigcup_{(n,m)\in I} R(n,m),$$

$$J := \{((n^0,m^0),(n^1,m^1)) \in I^2 \colon m^0 < m^1\}$$

and fix an arbitrary sequence $((n_i^0,m_i^0),(n_i^1,m_i^1))_{i\in\mathbb{N}}$ in J which contains every element of J infinitely often. Further, let $(g_i^0,g_i^1)_{i\in\mathbb{N}}$ be a sequence of pairwise disjoint pairs of vertices of D_{-1} with $g_i^0 \in R(n_i^0,m_i^0)$ and $g_i^1 \in R(n_i^1,m_i^1)$ for all $i \in \mathbb{N}$ and let D be the digraph obtained from D_{-1} by identifying g_i^0 and g_i^1 for any $i \in \mathbb{N}$.

We can think of D as having the same arc set as D_{-1} and refer to the ray in D with arc set $E(R(n,m))$ as $R(n,m)$ for any $(n,m) \in I$.

Proposition 4.4.1. *Every digraph D, constructed as above, contains k disjoint copies of R for all $k \in \mathbb{N}$. Moreover, if any copy of R in D has a tail in $R(n,m)$ for some $(n,m) \in I$, then D does not contain infinitely many disjoint copies of R, and particular, R is non-ubiquitous.*

Proof. By the choice of I and J, we have identified infinitely many vertices of $R(n,m)$ and $R(n',m')$ if $m \neq m'$ and none otherwise in the construction of D. Hence D contains arbitrarily many disjoint copies of R as the rays $R(0,m), R(1,m), \ldots, R(m,m)$ are disjoint for all $m \in \mathbb{N}$.

Suppose for a contradiction that there is a family $\mathcal{R}$ of infinitely many disjoint copies of R in D. For any copy R' of R in $\mathcal{R}$, there is by assumption $(n,m) \in I$ such that a tail of R' coincides with a tail of $R(n,m)$. Since infinitely vertices of two rays $R(n,m)$ and $R(n',m')$ are identified if $m \neq m'$, there is a fixed $m^* \in \mathbb{N}$ such that each ray in $\mathcal{R}$ has a tail identical with a tail of $R(n, m^*)$ for some $n \in \mathbb{N}$. But $n \leq m^*$ by the definition of I. So tails of the rays in $\mathcal{R}$ are contained in finitely many rays, which contradicts that they are disjoint. □

In the proofs of Theorems 4.4.2 and 4.4.4 we will use different strategies to choose the sequences $(g_i^0, g_i^1)_{i \in \mathbb{N}}$ such that any copy of R in D has a tail in $R(n, m)$ for some $(n, m) \in I$. Then D contains arbitrarily but not infinitely many copies of R by Proposition 4.4.1 as desired.

Theorem 4.4.2. *All rays with a bounded representing sequence are non-ubiquitous.*

Proof. Let R be an arbitrary ray with a bounded representing sequence. Let c be the largest natural number that occurs infinitely often in its representing sequence. By Lemma 4.2.1, we may assume that infinitely many phases of length c in R are out-oriented. Further, let I, J, $((n_i^0, m_i^0), (n_i^1, m_i^1))_{i \in \mathbb{N}}$ and D_{-1} be as defined in the beginning of this section.

Now we define a sequence $(g_i^0, g_i^1)_{i \in \mathbb{N}}$ of pairwise disjoint pairs of vertices of D_{-1} recursively with $g_i^0 \in R(n_i^0, m_i^0)$ and $g_i^1 \in R(n_i^1, m_i^1)$ for all $i \in \mathbb{N}$. If (g_j^0, g_j^1) has been defined for all $j < i$, we pick for $\varepsilon \in \{0, 1\}$ the vertex g_i^ε beyond all vertices $g_0^0, g_0^1, \ldots, g_{i-1}^0, g_{i-1}^1$ on $R(n_i^\varepsilon, m_i^\varepsilon)$ with the following properties (see Figure 4.2):

(i) g_i^1 is a turn in $R(n_i^1, m_i^1)$ at the start of an out-oriented phase of length c, and

(ii) g_i^0 is a turn in $R(n_i^0, m_i^0)$ at the end of an out-oriented phase of length c with the property that $|R(n_i^0, m_i^0) g_i^0| > |R(n_i^1, m_i^1) g_i^1|$.

This is possible since $R(n_i^0, m_i^0)$ and $R(n_i^1, m_i^1)$ contain infinitely many out-oriented phases of length c. Let D be the digraph constructed from D_{-1} and $(g_i^0, g_i^1)_{i \in \mathbb{N}}$ as defined in the beginning of this section. We write g_i for the vertex $g_i^0 = g_i^1$ in D.

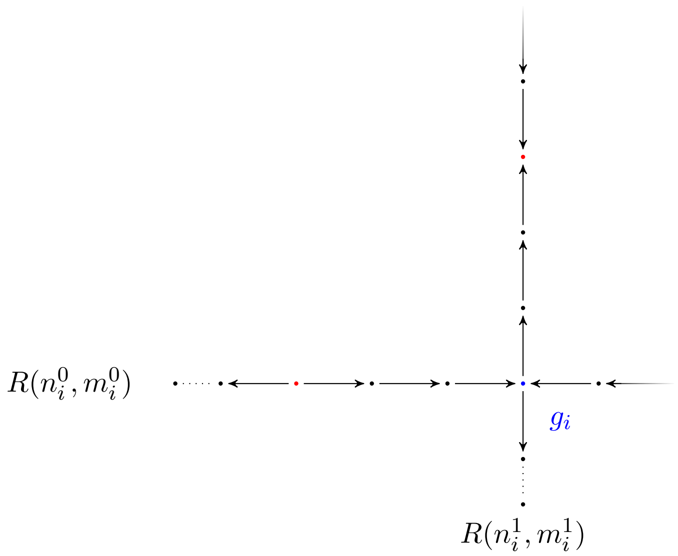

Figure 4.2: Example of a vertex $g_i^0 = g_i^1 = g_i$ in D for $c = 3$.

By Proposition 4.4.1, it suffices to prove that for any copy R' of R in D there is $(n, m) \in I$ such that a tail of R' coincides with a tail of $R(n, m)$. Begin by fixing an arbitrary copy R'

of R in D. Note that, by the construction of D, R' consists of (possibly infinite) segments contained in rays $R(n,m)$ for $(n,m) \in I$, and R' can switch from $R(n,m)$ to $R(n',m')$ with $(n,m) \neq (n',m')$ only at some identification vertex g_i. Since in the construction of D we only identified turns and c is the largest number which occurs infinitely often in the representing sequence of R, there is a tail R'' of R' whose representing sequence contains only numbers up to c and whose initial vertex is a turn of some ray $R(n^*,m^*)$ for $(n^*,m^*) \in I$. Thus each phase of any $R(n',m')$ for $(n',m') \in I$ is either completely traversed by R'' or all arcs of this phase are avoided by R''.

Let $i \in \mathbb{N}$ be arbitrary. By properties (i) and (ii), any ray in D traversing both the phase of $R(n_i^0,m_i^0)g_i$ incident with g_i and an arc of $R(n_i^1,m_i^1)$ incident with g_i contains a phase of length $> c$ (see Figure 4.2). Clearly, the same holds for a ray traversing both, the phase of $g_iR(n_i^1,m_i^1)$ incident with g_i and an arc of $R(n_i^0,m_i^0)$ incident with g_i. Recall from the construction of D that the vertex g_i has degree four. Thus if R'' contains g_i as an inner vertex, exactly one of the following properties holds:

(1) both arcs of R'' incident with g_i are contained in $R(n_i^0,m_i^0)$, or

(2) both arcs of R'' incident with g_i are contained in $R(n_i^1,m_i^1)$, or

(3) one arc of R'' incident with g_i is contained in $R(n_i^1,m_i^1)g_i$ and one is contained in $g_iR(n_i^0,m_i^0)$.

This feature restricts in which way R'' is embedded into D. First we observe:

Claim. For any $(n,m) \in I$ and any arc $a \in A(R(n,m)) \cap A(R'')$, the rays $R(n,m)$ and R'' traverse a in the same direction.

Proof of the claim. Suppose for a contradiction that the claim is false. Pick $(n,m) \in I$ and an arc $a \in A(R(n,m)) \cap A(R'')$ with $|R(n,m)a|$ minimal such that a contradicts this property (see Figure 4.3). Let u, v be the endvertices of a such that $|R(n,m)u| < |R(n,m)v|$ applies. By minimality of a, the other arc of R'' incident with u is not contained in $R(n,m)u$. Therefore $u = g_i$ for some $i \in \mathbb{N}$, i.e. there is $(n',m') \in I$ with $m \neq m'$ such that $u \in V(R(n,m)) \cap V(R(n',m'))$. Then (3) holds for R'' at the vertex u. As $a \in A(uR(n,m))$, the inequality $m < m'$ holds by definition of J and the other arc incident with u in R'' is contained in $R(n',m')u$. The rays R'' and $R(n',m')$ traverse this arc in opposite directions, but $|R(n',m')u| < |R(n,m)u|$ holds by property (ii) of the construction. This contradicts the minimality of $|R(n,m)a|$. □

Now let $m \in \mathbb{N}$ be the smallest number such that there is a pair $(n,m) \in I$ with $A(R'') \cap A(R(n,m)) \neq \varnothing$ and let a be an element of this intersection. We prove that aR'' coincides with $aR(n,m)$. Suppose not and let $(n',m') \neq (n,m) \in I$ and $u \in V(R(n,m)) \cap V(R(n',m'))$ such that u is incident with the first arc b of aR'' not contained in $A(R(n,m))$. Property (3) applies to u. By minimality of m we have $m < m'$ and thus $b \in A(R(n',m')u)$ by (3). Therefore aR'' and $R(n',m')$ traverse the arc b in opposite directions, which contradicts the claim above. Thus R' has a tail that coincides with a tail of $R(n,m)$. This completes the proof. □

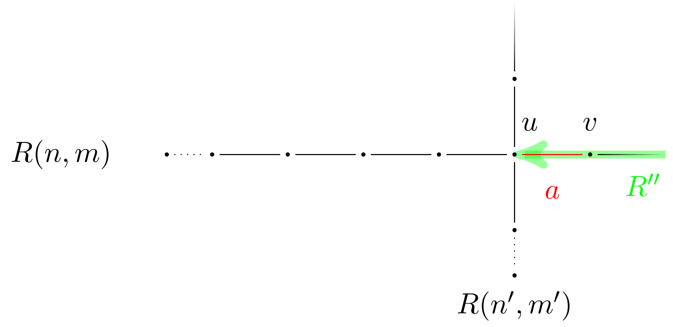

Figure 4.3: The arc a and the ray R'' in the proof of the claim for Theorem 4.4.2.

It is left to prove Theorem 4.4.4, for which we need the following Lemma:

Lemma 4.4.3. *Let R be a ray with an unbounded representing sequence. Then the tails vR and wR are non-isomorphic for all $v \neq w \in R$.*

Proof. Let R be a ray with an unbounded representing sequence. Suppose that there are $v \neq w \in R$ such that w lies beyond v on R and there is an isomorphism $\varphi\colon vR \to wR$. Clearly, we have $\varphi(v) = w$ and the paths $\varphi^n(v)R\varphi^{n+1}(v)$ are isomorphic for all $n \in \mathbb{N}$. Therefore, the representing sequence of vR is periodic. Thus the representing sequence of R is bounded, a contradiction. □

Theorem 4.4.4. *All rays with an unbounded representing sequence are non-ubiquitous.*

Proof. Let R be an arbitrary ray with an unbounded representing sequence. Further, let $I, J, ((n_i^0, m_i^0), (n_i^1, m_i^1))_{i\in\mathbb{N}}$ and D_{-1} be as defined in the beginning of this section.

We define a sequence $(g_i^0, g_i^1)_{i\in\mathbb{N}}$ of pairwise disjoint pairs of vertices of D_{-1} recursively with $g_i^0 \in R(n_i^0, m_i^0)$ and $g_i^1 \in R(n_i^1, m_i^1)$ for all $i \in \mathbb{N}$. We fix an enumeration $\{v_0, v_1, \ldots\}$ of $V(D_{-1})$. Denote by D_i the digraph obtained from D_{-1} by identifying the vertices g_ℓ^0 and g_ℓ^1 for all $\ell \leq i$ and write g_ℓ for the vertex $g_\ell^0 = g_\ell^1$. When a vertex v_j is identified with a vertex v_k in this process, we call the new identification vertex both v_j and v_k. We will make sure that the following holds for all $i \in \mathbb{N}$:

(i) Let $k \leq i$ and let S be a v_k–g_i path in D_i which is isomorphic to an initial segment of R. Then $|S| = |R(n_i^0, m_i^0)g_i|$ or $|S| = |R(n_i^1, m_i^1)g_i|$.

(ii) g_i is not a turn of $R(n_i^0, m_i^0)$ or $R(n_i^1, m_i^1)$ and thus g_i lies in the interior of phases M^0 of $R(n_i^0, m_i^0)$ and M^1 of $R(n_i^1, m_i^1)$. Let $t_0^\varepsilon, t_1^\varepsilon$ be the two turns which are endvertices of M^ε. Then the four numbers $d_{R(n_i^\varepsilon, m_i^\varepsilon)}(t_\delta^\varepsilon, g_i)$ for $\delta, \varepsilon \in \{0, 1\}$ are pairwise distinct. Furthermore, M^0 and M^1 do not contain any g_j with $j \neq i$.

Let us first derive from the existence of a sequence $(g_i^0, g_i^1)_{i\in\mathbb{N}}$ as above that the digraph D, constructed from D_{-1} and $(g_i^0, g_i^1)_{i\in\mathbb{N}}$ as in the beginning of this section, has the property that any copy of R in D has a tail that is contained in some $R(n, m)$. Then Proposition 4.4.1 ensures that D contains arbitrarily but not infinitely many disjoint copies of R.

We know that R' traverses infinitely many vertices g_j with $j \in \mathbb{N}$ (but R' does not necessarily swap from one ray of the form $R(n,m)$ to another at every g_j) because every ray $R(n,m)$ is glued together with other rays at infinitely many vertices in D. Suppose that R' starts in v_k and let g_i be the first vertex of R' which lies in $\{g_j : j \geq k\}$. Then the path $R'g_i$ is a subdigraph of D_i as it contains no vertex g_j with $j > i$. Hence by (i), there is a ray $R(n^*, m^*)$ containing g_i with $|R'g_i| = |R(n^*, m^*)g_i|$. The two rays R' and $R(n^*, m^*)$ are isomorphic as both are isomorphic to R, so the two initial segments of the same length are isomorphic, and so the tails g_iR' and $g_iR(n^*, m^*)$ are also isomorphic. Now assume for a contradiction that g_iR' exits $g_iR(n^*, m^*)$ and let $g_{i'}$ be the first vertex where g_iR' exists $g_iR(n^*, m^*)$. Note that $g_iR'g_{i'}$ and $g_iR(n^*, m^*)g_{i'}$ are identical. From this and the isomorphism found before of the two tails beginning at g_i, it follows that the tails $g_{i'}R'$ and $g_{i'}R(n^*, m^*)$ are also isomorphic. Thus, it follows from (ii) that R' has to exit $g_{i'}$ through the segment $g_{i'}R(n^*, m^*)t_1^*$, where t_1^* is the end-vertex beyond $g_{i'}$ of the phase containing $g_{i'}$ in $R(n^*, m^*)$, because this is the segment that has the correct length of the four segments of pairwise distinct lengths incident with $g_{i'}$. This means g_iR' does not exit $g_iR(n^*, m^*)$ at $g_{i'}$, a contradiction.

All that remains is to define (g_i^0, g_i^1) for all $i \in \mathbb{N}$. Suppose that (g_j^0, g_j^1) is already defined for all $j < i$. We write $R^\varepsilon := R(n_i^\varepsilon, m_i^\varepsilon)$ for $\varepsilon \in \{0, 1\}$. Our task is to specify suitable vertices $g_i^\varepsilon \in R^\varepsilon$. Let x be a vertex of R^0 that lies beyond all phases of R^0 that contain any vertex of $\{g_0, \ldots, g_{i-1}, v_0, \ldots, v_i\}$. Let $\mathcal{P}$ be the set of all $\{v_0, \ldots, v_i\}$–x paths in D_{i-1} which are isomorphic to initial segments of R. Then none of the vertices on R^0 beyond x are identified with other vertices, so

(∗) for any $y \in xR^0$, every $\{v_0, \ldots, v_k\}$–y path is of the form PxR^0y for some $P \in \mathcal{P}$.

The following claim implies that $\mathcal{P}$ is finite:

Claim. For all $j \in \mathbb{N}$ and for all vertices $v, w \in D_j$, the set of v–w paths in D_j is finite.

Proof of the claim. We use induction on j. The claim holds for $j = -1$ as D_{-1} is a disjoint union of rays. Now let $j \geq 0$ and consider an arbitrary v–w path P in D_j. If P does not use the vertex g_j, then P is also a v–w path in D_{j-1}, and there are only finitely many such paths by induction. Otherwise g_j lies on P, and P consists of a v–g_j path Q concatenated with a g_j–w path Q' in D_j. Since Q and Q' are also paths in D_{j-1}, there are only finitely many possibilities for Q and Q' by induction. □

We write $\mathcal{Q}$ for the subset of $\mathcal{P}$ consisting of all paths Q with $|Q| \neq |R^0x|$. Our next step is to find a vertex z^0 such that

(†) z^0 lies beyond x on R^0 and no path from $\mathcal{Q}$ can be extended to a $\{v_0, \ldots, v_i\}$–z^0 path in D_{i-1} which is isomorphic to an initial segment of R.

Let $Q \in \mathcal{Q}$. Since $|Q| \neq |R^0x|$, then $q \neq x$ where q is the end-vertex (other than the origin of R^0) of the initial segment to which Q is isomorphic in R^0. Then by Lemma 4.4.3, qR^0 and xR^0 are non-isomorphic. So there is an initial segment of xR^0 that is not isomorphic to an initial segment of qR^0. Let xR^0y_Q be such a segment, then QxR^0y_Q is not isomorphic to an initial segment of R^0, and hence also not to an initial segment of R. Let z^0 be the latest vertex on R^0 that is of the form y_Q for some $Q \in \mathcal{Q}$ if $\mathcal{Q}$ is non-empty or let $z^0 := x$ if $\mathcal{Q}$ is empty. Then by $(*)$, for every $Q \in \mathcal{Q}$ the path QxR^0z^0 contains QxR^0y_Q as a subpath. Thus it follows from the choice of y_Q that QxR^0z^0 cannot be isomorphic to an initial segment of R. This shows that our choice of z^0 satisfies $(\dagger)$. Similarly, define a vertex $z^1 \in R^1$.

Next, we show that D_i will satisfy (i) for every choice of vertices g_i^ε on R^ε that lie beyond z^ε on R^ε for $\varepsilon \in \{0, 1\}$. So suppose we have already fixed vertices g_i^ε as above and glued them together. Now consider any $k \leq i$ and any v_k–g_i path S in D_i which is isomorphic to an initial segment of R. Then S is also a v_k–g_i^0 or a v_k–g_i^1 path in D_{i-1}; suppose without loss of generality that the former holds. Let v^* be the last vertex of S that is contained in the set $\{g_0, \ldots, g_{i-1}, v_0, \ldots, v_i\}$. Since z^0 lies beyond v^* and g_i^0 lies beyond z^0 on R^0, it follows that S must contain z^0. Then Sx is contained in $\mathcal{P}$ but not in $\mathcal{Q}$ by $(\dagger)$ as Sz^0 is isomorphic to an initial segment of R. Therefore $|Sx| = |R^0x|$, so $|S| = |R^0g_i|$, which proves (i).

Finally, we further specify the choice of g_i^0 and g_i^1 so that (ii) holds. Recall that the representing sequence of R is unbounded. Therefore, we can find a phase M^0 of R^0 which is contained in z^0R^0 and has length at least 3. We choose g_i^0 as an interior vertex of M^0 so that g_i^0 has different distances to both endvertices of M^0. Next, find a phase M^1 of R^1 which is contained in z^1R^1 such that $|M^1| \geq 2|M^0| + 1$. Then (ii) is fulfilled for a vertex g_i^1 in M^1 which has distance $|M^0|$ to one endvertex of M^1 and hence distance $> |M^0|$ to its other endvertex. □

Chapter 5

On the ubiquity of oriented double rays

5.1 Introduction

In the previous chapter, we showed:

Theorem 5.1.1. *An oriented ray is ubiquitous if and only if it has a finite number of turns.*

In this chapter we extend the investigation of ubiquity in digraphs to the class of digraphs whose underlying undirected graphs are double rays. We call these digraphs *oriented double rays* and tackle the following problem:

Problem 5.1.2. *Which oriented double rays are ubiquitous?*

We define *turns* of oriented double rays in the same way as we defined turns of oriented rays, namely as vertices of in-degree or out-degree 2. It turns out that – in contrast to oriented rays – not only is it relevant whether the number of turns is finite or infinite, but also the parity of this number plays a role. The main result of this chapter reads as follows:

Theorem 5.1.3. *An oriented double ray with at least one turn is ubiquitous if and only if it has a (finite) odd number of turns.*

Theorem 5.1.3 solves Problem 5.1.2, except for the case of the oriented double ray without any turns. This raises the following problem:

Problem 5.1.4. *Is the oriented double ray without turns ubiquitous?*

More generally, one can investigate ubiquity of digraphs whose underlying undirected graphs are trees. However, even the question which undirected trees are subgraph-ubiquitous is unsolved and only known for ubiquity with respect to weaker relations such as the topological minor relation. Therefore it might be sensible also to discuss ubiquity of digraphs with respect to weaker relations such as butterfly minors. Moreover, since proving or disproving the ubiquity of consistently oriented double rays is not easy, we propose to initially consider *out-trees*, i.e. trees in which all arcs are oriented away from the root.

Problem 5.1.5. *Which out-trees are ubiquitous concerning a fitting notion of ubiquity?*

The proof of Theorem 5.1.3 is constructive, employing results from the previous chapter and investigating symmetry properties of oriented double rays. We begin by proving in Section 5.3 that R is non-ubiquitous if it has an even but non-zero or infinite number of turns: In Subsection 5.3.1, we address the case where R has an even but non-zero number of turns. If R has infinitely many turns, we distinguish whether R is non-periodic (Subsection 5.3.2) or periodic (Subsection 5.3.3) (periodicity is defined in Section 5.2). Finally, in Section 5.4 we show that any oriented double ray with an odd number of turns is ubiquitous.

5.2 Preliminaries

For digraphs D and D', we write $D \cong D'$ if D is isomorphic to D' and $D \leq D'$ if D is isomorphic to a subdigraph of D'.

In this chapter, rays and double rays are digraphs together with linear orders on their sets of vertices: A *ray* or *double ray* is a digraph R together with a linear order $\leq_R$ on $V(R)$ isomorphic to $\mathbb{N}$ or $\mathbb{Z}$, respectively, such that for all vertices v, w of R:

- if v and w are consecutive in the linear order, then either $vw \in A(R)$ or $wv \in A(R)$ but not both, and
- if v and w are not consecutive in the linear order, then $vw, wv \notin A(R)$.

Hence, rays and double rays are oriented counterparts of undirected rays and double rays, together with linear orders.

We say that a (double) ray is a *subdigraph* of another (double) ray or is *isomorphic* to another (double) ray if this is true for the digraphs in the usual sense, regardless of the linear orders.

In the following, we extend some definitions for rays from the previous chapter to rays and double rays as in the setting of this chapter, and add some new definitions. Let R be a ray or a double ray and let $v \in V(R)$. We write vR for the subdigraph of R induced by all vertices w of R with $w \geq_R v$ and Rv for the subdigraph of R induced by all vertices w of R with $w \leq_R v$. An infinite subdigraph of this form is called a *tail* of R. We call R *periodic* if R has a non-trivial $\leq_R$-preserving endomorphism. (Note that when R is a ray, any endomorphism must preserve $\leq_R$. However, the same is not true for double rays.) Let $v \in V(R)$ and let f be a non-trivial $\leq_R$-preserving endomorphism of R such that the distance d between v and $f(v)$ in R is minimal. Then we say that R has *periodicity* d (and d is independent of the choices of v and f). We call the unique $\leq_R$-minimal vertex of a ray R the *root* of R. A maximal (possibly infinite) directed path contained in a ray or double ray, i.e. a maximal connected subdigraph whose arcs are consistently oriented, is called a *phase*.

5.3 Negative results

In this section we prove the backwards implication of Theorem 5.1.3, which is divided into three different parts. First we show that double rays with a (finite) even, non-zero number of turns are non-ubiquitous in Theorem 5.3.1. Then we show that non-periodic double rays with infinitely many turns are non-ubiquitous in Theorem 5.3.5, and lastly that periodic double rays with infinitely many turns are non-ubiquitous in Theorem 5.3.6.

5.3.1 Double rays with an even, non-zero number of turns

Any double ray R with an even, non-zero number of turns contains an in-ray and an out-ray. By glueing together in- and out-rays of the members of a thick R-tribe[1] in a specific way we can show:

Theorem 5.3.1. *Any double ray with an even, non-zero number of turns is non-ubiquitous.*

Proof. Let R be a double ray with an even, non-zero number of turns. Let s be the first and t be the last turn of R. Since the number of turns is even, exactly one of Rs and tR is an in-ray and exactly one an out-ray. By possibly reversing the order $\leq_R$, we may assume that the former is an in-ray and the latter an out-ray. Let $p \in \mathbb{N}$ be the length of the longest finite phase of R. We will construct a digraph D containing arbitrarily many but not infinitely many disjoint copies of R.

For the construction of D, we use the auxiliary set

$$I := \{(n, m) \in \mathbb{N}^2 : n \leq m\}$$

ordered by the colexicographic order $\leq_{\text{col}}$. For $(n, m) \in I$, we write $(n, m)^+$ for the successor of (n, m) under $\leq_{\text{col}}$ and, if there exists a predecessor of (n, m) under $\leq_{\text{col}}$, we refer to this predecessor as $(n, m)^-$.

Let $(R(n, m))_{(n,m)\in I}$ be a family of pairwise disjoint copies of R. For $(n, m) \in I$, write $s(n, m)$ for the first turn of $R(n, m)$ and $t(n, m)$ for the last turn of $R(n, m)$.

Next, we define two families of vertices of $R(n, m)$ for every $(n, m) \in I$: Let $(v^{(i,j)}_{(n,m)})_{(i,j)\in I}$ be a family of vertices of $R(n, m)s(n, m)$ such that

- the order $\leq_{R(n,m)}$ and the order induced by $\leq_{\text{col}}$ on the superindices are reversed on $\{v^{(i,j)}_{(n,m)} : (i, j) \in I\}$,
- the vertices of $(v^{(i,j)}_{(n,m)})_{(i,j)\in I}$ have distance at least $p + 1$ to each other and to $s(n, m)$ in $R(n, m)$.

Let $(w^{(i,j)}_{(n,m)})_{(i,j)\in I}$ be a family of vertices of $t(n, m)R(n, m)$ such that

- the order $\leq_{R(n,m)}$ and the order induced by $\leq_{\text{col}}$ on the superindices coincide on $\{w^{(i,j)}_{(n,m)} : (i, j) \in I\}$,

[1]see Section 4.2 for the definition of (thick) R-tribes

- the vertices of $(w^{(i,j)}_{(n,m)})_{(i,j)\in I}$ have distance at least $p+1$ to each other and to $t(n,m)$ in $R(n,m)$.

Let D be the digraph constructed from the disjoint union $\bigsqcup_{(n,m)\in I} R(n,m)$ by identifying the two vertices $v^{(k,\ell)}_{(i,j)}$ and $w^{(i,j)}_{(k,\ell)}$ for any $(i,j),(k,\ell)\in I$ with $j\neq \ell$ (see Figure 5.1). We simply refer to the vertex of D that is obtained by identification of $v^{(k,\ell)}_{(i,j)}$ and $w^{(i,j)}_{(k,\ell)}$ as $v^{(k,\ell)}_{(i,j)}$.

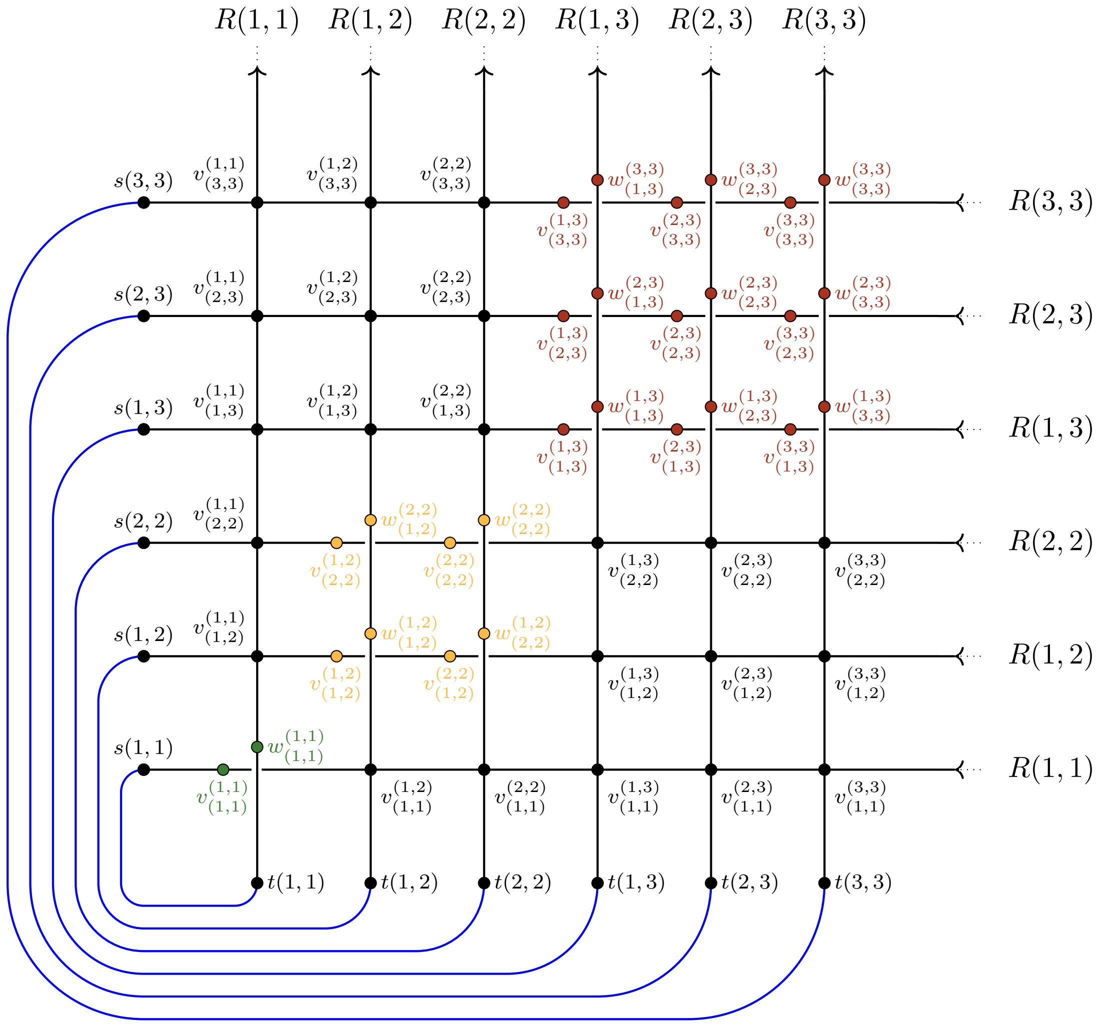

Figure 5.1: The digraph D constructed in the proof of Theorem 5.3.1. The horizontal paths represent parts of the in-rays $R(n,m)s(n,m)$ and the vertical paths represent parts of the out-rays $t(n,m)R(n,m)$. The blue paths connecting the vertices $s(n,m)$ and $t(n,m)$ represent the finite paths $s(n,m)R(n,m)t(n,m)$ that consist of the union of all finite phases of $R(n,m)$.

The digraph D contains m disjoint copies of R for every $m\in\mathbb{N}$, as the double rays $R(1,m),\dots,R(m,m)$ are disjoint. Thus it remains to prove that D does not contain infinitely many disjoint copies of R. In Claim 1, we investigate how an out-ray can lie in D, and in Claim 2, we investigate how an in-ray with a fixed root can lie in D. After that, we deduce from the two claims and from the fact that R contains both an out-ray and an in-ray, that D does not contain infinitely many copies of R.

Claim 1. Any out-ray in D has a tail that coincides with $xR(n,m)$ for some $(n,m) \in I$ and some $x \in V(R(n,m))$.

Proof. Let S be an arbitrary out-ray in D. If S is completely contained in some $R(n,m)$, then we are done.

Thus we can assume that S is not completely contained in some $R(n,m)$. This implies that S must contain an identification vertex, i.e. $v_{(i,j)}^{(k,\ell)} \in V(S)$ for some $(i,j),(k,\ell) \in I$. Let $(n,m) \in I$ be $\leq_{\text{col}}$-minimal with the property that S contains a vertex $v_{(i,j)}^{(n,m)}$ for some $(i,j) \in I$. We show that $v_{(i,j)}^{(n,m)}S$ is a tail of $t(n,m)R(n,m)$.

Suppose for a contradiction that $v_{(i,j)}^{(n,m)}S$ is not a tail of $t(n,m)R(n,m)$. Set $v := v_{(i,j)}^{(n,m)}$ if $v_{(i,j)}^{(n,m)} \notin V(t(n,m)R(n,m))$ (i.e. if $j = m$), and otherwise let v be the last vertex of $v_{(i,j)}^{(n,m)}S$ such that $v_{(i,j)}^{(n,m)}Sv$ is contained in the out-ray $t(n,m)R(n,m)$. In either case, there is $(k,\ell) \in I$ such that $v = v_{(k,\ell)}^{(n,m)}$ since in the latter case v had to be identified with another vertex in the construction. Then the first arc of vS is contained in $R(k,\ell)$. Since S is an out-ray, it also contains $v_{(k,\ell)}^{(n,m)^-}$ if $(n,m) \neq (1,1)$, or $s(k,\ell)$ if $(n,m) = (1,1)$ (see Figure 5.1). This contradicts either the minimality of (n,m) or the fact that $s(k,\ell)$ is a turn, respectively. ∎

Claim 2. For any $(n',m'),(n,m) \in I$, any in-ray in D with root $s(n',m')$ contains a vertex of $\{v_{(i,j)}^{(n,m)} : (i,j) \in I\}$.

Proof. Let $(n,m),(n',m') \in I$ and let S be an arbitrary in-ray in D with root $s(n',m')$. As no vertex of $s(n',m')S(n',m')t(n',m')$ has been identified with other vertices in the construction of D and $t(n',m')$ is a turn, S contains the vertex $v_{(n',m')}^{(1,1)}$.

We consider the set $\mathcal{X} := \{v_{(i,j)}^{(k,\ell)} \in V(S) : (i,j),(k,\ell) \in I, (k,\ell) \leq_{\text{col}} (n,m)\}$ and let $(k,\ell) \in I$ be $\leq_{\text{col}}$-maximal with the property that an element of $\mathcal{X}$ has superindex (k,ℓ). We show $(k,\ell) = (n,m)$, which implies the statement of Claim 2.

Suppose for a contradiction that $(k,\ell) <_{\text{col}} (n,m)$. Let $(i,j) \in I$ be $\leq_{\text{col}}$-minimal with the property that $v_{(i,j)}^{(k,\ell)}$ is an element of $\mathcal{X}$. The first arc of $v_{(i,j)}^{(k,\ell)}S$ lies either in $R(i,j)s(i,j)$ or in $t(k,\ell)R(k,\ell)$. In the former case, the in-ray S must also contain $v_{(i,j)}^{(k,\ell)^+}$ (see Figure 5.1), contradicting the maximality of (k,ℓ). In the latter case, S also contains $t(k,\ell)$ if $v_{(i,j)}^{(k,\ell)}$ is the first vertex of $t(k,\ell)R(k,\ell)$ that has been identified, or S contains a vertex $v_{(i',j')}^{(k,\ell)}$ for $(i',j') <_{\text{col}} (i,j)$ otherwise (see Figure 5.1). This contradicts either the fact that $t(k,\ell)$ is a turn or the minimality of (i,j), respectively. ∎

Let $\hat{R}$ be an arbitrary copy of R in D. By Claim 1, there is $(n,m) \in I$ and a vertex $x \in V(R(n,m))$ such that an out-ray in $\hat{R}$ coincides with $xR(n,m)$. To prove that D cannot contain infinitely many disjoint copies of R, it suffices to show that any copy of R in D has an in-ray that starts in some $s(n',m')$. Then by Claim 2, every copy of R in D contains a vertex of $\{v_{(i,j)}^{(n,m)} : (i,j) \in I\}$, contradicting that the set $\{v_{(i,j)}^{(n,m)} : (i,j) \in I\} \setminus V(xR(n,m))$ is finite.

Let $\tilde{R}$ be any copy of R in D and let $\tilde{R}y$ be the unique phase of $\tilde{R}$ that forms an in-ray. By construction of D, any phase of length at most p of $\tilde{R}$ is contained in $s(i,j)R(i,j)t(i,j)$ for some $(i,j) \in I$, which immediately yields $y \in \{s(n',m'), t(n',m')\}$ for some $(n',m') \in I$. Since $\tilde{R}y$ is an in-ray, we must have $y = s(n',m')$ as desired. This completes the proof. □

5.3.2 Non-periodic double rays with infinitely many turns

In Lemma 5.3.2, Lemma 5.3.3 and Theorem 5.3.4 we investigate symmetry properties of non-periodic double rays and show that any such double ray R has a tail which is isomorphic to only very specific other tails of R. In Theorem 5.3.5, we use this result to reduce the non-ubiquity of non-periodic double rays with infinitely many turns to the non-ubiquity of rays with infinitely many turns (Theorem 5.1.1).

Lemma 5.3.2. *For every non-periodic double ray R there exists $v^* \in V(R)$ such that Rv^* is non-periodic.*

Proof. Suppose for a contradiction that Rw is periodic for all $w \in V(R)$. Let $p \in \mathbb{N}$ be minimal among the periodicities of Rw for all $w \in V(R)$ and let $v \in V(R)$ such that Rv has periodicity p. We show that Rw has periodicity p for any $w \in V(R)$. This then implies that R is periodic with periodicity p, a contradiction.

For $w >_R v$, let f be any non-trivial endomorphism of Rw, which exists since Rw is periodic. By concatenating f with itself multiple times if necessary, we obtain an endomorphism of Rw whose image is contained in Rv. Thus it remains to prove that Ru has periodicity p for any $u <_R v$. Since Ru is a tail of Rv, Ru has periodicity at most p, and by minimality of p, Ru has periodicity exactly p. □

Lemma 5.3.3. *For every non-periodic double ray R there exists $v^* \in V(R)$ such that for all $v \geq_R v^*$ and all $w \in V(R)$:*

$$Rv \cong Rw \Rightarrow v = w.$$

Proof. We choose v^* such that Rv^* is non-periodic according to Lemma 5.3.2. Let $v \geq_R v^*$ and $w \in V(R)$ such that $Rv \cong Rw$. If $w <_R v$, we can restrict the isomorphism $Rv \to Rw$ to a non-trivial endomorphism of Rv^*, contradicting that Rv^* is non-periodic. If $w >_R v$, we can restrict the isomorphism $Rw \to Rv$ to a non-trivial endomorphism of Rv^*, which again is a contradiction. Thus $v = w$ holds. □

Theorem 5.3.4. *For every non-periodic double ray R there exists $\hat{v} \in V(R)$ such that for all $w \in V(R)$:*

(1) $R\hat{v} \not\leq \hat{v}R$,

(2) $R\hat{v} \cong Rw \Rightarrow \hat{v}R \cong wR$, and

(3) $R\hat{v} \cong wR \Rightarrow \hat{v}R \cong Rw$.

Proof. Let v^* be as in Lemma 5.3.3.

Case 1: There is $v \geq_R v^*$ such that $Rv \ncong wR$ for all $w \in V(R)$.

In this case we set $\hat{v} := v$, which directly implies (1) and (3). For (2), since $\hat{v} \geq v^*$ and v^* was picked with the property of Lemma 5.3.3, $R\hat{v} \cong Rw$ implies $\hat{v} = w$ and thus $\hat{v}R = wR$.

Case 2: For all $v \geq_R v^*$ there is $\alpha(v) \in V(R)$ such that $Rv \cong \alpha(v)R$.

Claim 1. For every $v \geq_R v^*$, the vertex $\alpha(v)$ is unique.

Proof. Suppose for a contradiction that there are $\alpha(v) <_R \alpha(v)' \in V(R)$ such that $\alpha(v)R \cong Rv \cong \alpha(v)'R$. Since $\alpha(v)'R$ is a proper tail of $\alpha(v)R$, it follows that $Rv \cong \alpha(v)'R$ is isomorphic to a proper tail of $Rv \cong \alpha(v)R$. Thus there exists a non-trivial endomorphism of Rv, which contradicts that v^* has the property of Lemma 5.3.3. ∎

Claim 2. There is $\hat{v} \geq_R v^*$ such that $\alpha(\hat{v}) <_R \hat{v}$.

Proof. Let v and v' be any vertices of R with $v^* \leq_R v <_R v'$. Since Rv is a proper tail of Rv', it follows that $\alpha(v)R \cong Rv$ is isomorphic to a proper tail of $\alpha(v')R \cong Rv'$. Thus there is a vertex $w >_R \alpha(v')$ such that $wR \cong \alpha(v)R$. Then $w = \alpha(v)$ by Claim 1. In conclusion, we have established that $\alpha(v) >_R \alpha(v')$ whenever $v <_R v'$. Therefore, it is possible to pick a vertex $\hat{v}$ which is sufficiently large with respect to $\leq_R$, so that $\alpha(\hat{v}) <_R \hat{v}$. ∎

We show that any vertex $\hat{v}$ as in Claim 2 satisfies properties (1) to (3). For (1), assume that $R\hat{v} \leq \hat{v}R$. This means that there is $v' \geq_R \hat{v}$ with $R\hat{v} \cong v'R$. A contradiction since by Claim 1 $\alpha(\hat{v})$ is unique, but $\alpha(\hat{v}) <_R \hat{v} \leq_R v'$ by choice of $\hat{v}$. For (2), recall that $\hat{v} \geq_R v^*$. Thus, Lemma 5.3.3 implies that the only vertex $w \in V(R)$ with $R\hat{v} \cong Rw$ is $\hat{v}$. This proves (2) since clearly $\hat{v}R \cong \hat{v}R$. For (3), as $\alpha(\hat{v})$ is unique by Claim 1, it is enough to prove $\hat{v}R \cong R\alpha(\hat{v})$. Let $\psi : R\hat{v} \to \alpha(\hat{v})R$ be an isomorphism, which exists by choice of $\alpha(\hat{v})$. Then ψ maps $\alpha(\hat{v})$ to $\hat{v}$, since the distance between $\alpha(\hat{v}) \in V(R\hat{v})$ and the root of $R\hat{v}$ and the distance between $\hat{v} \in V(\alpha(\hat{v})R)$ and the root of $\alpha(\hat{v})R$ are the same. Thus ψ can be restricted to an isomorphism $R\alpha(\hat{v}) \to \hat{v}R$.

This completes the proof. □

Now we combine Theorem 5.1.1 and Theorem 5.3.4 to prove:

Theorem 5.3.5. *Any non-periodic double ray R with infinitely many turns is non-ubiquitous.*

Proof. Let R be any such double ray. We construct a digraph D that contains arbitrarily many but not infinitely many copies of R. Without loss of generality, for every $v \in V(R)$ the ray vR contains infinitely many turns (otherwise reverse the order $\leq_R$). Let $\hat{v} \in V(R)$ be as in Theorem 5.3.4. As $\hat{v}R$ contains infinitely many turns, there is a digraph D' containing arbitrarily many but not infinitely many disjoint copies of $\hat{v}R$ by Theorem 5.1.1. We construct D from D' and a family $(S_x)_{x \in V(D')}$ of disjoint copies of $R\hat{v}$ by identifying the root of S_x with x for each $x \in V(D')$.

By construction, D contains arbitrarily many disjoint copies of R. We have to show that D does not contain infinitely many disjoint copies of R, which implies the theorem. It suffices to prove that each copy of R in D has a tail isomorphic to $\hat{v}R$ that is contained in the subdigraph D' of D. Then D cannot contain infinitely many disjoint copies of R since D' does not contain infinitely many disjoint copies of $\hat{v}R$.

Let $\tilde{R}$ be any copy of R in D. If $\tilde{R}$ is completely contained in D', we are done. Thus we can suppose that there is $x \in V(D')$ and $w \in V(\tilde{R})$ such that either $S_x = \tilde{R}w$ or $S_x = w\tilde{R}$.

In the former case, we have $R\hat{v} \cong S_x = \tilde{R}w$ and thus $\hat{v}R \cong w\tilde{R}$ by Theorem 5.3.4 (2). It follows from (1) that $R\hat{v} \not\leq \hat{v}R \cong w\tilde{R}$. Hence $w\tilde{R}$ cannot have a tail in any S_y for $y \in V(D')$. Thus $w\tilde{R}$ is the desired tail of $\tilde{R}$ which is isomorphic to $\hat{v}R$ and contained in D'.

Similarly, in the latter case, we have $R\hat{v} \cong S_x = w\tilde{R}$ and thus $\hat{v}R \cong \tilde{R}w$ by (3). It follows from (1) that $R\hat{v} \not\leq \hat{v}R \cong \tilde{R}w$. Hence $\tilde{R}w$ cannot have a tail in any S_y for $y \in V(D')$ and $\tilde{R}w$ is the desired tail of $\tilde{R}$. □

5.3.3 Periodic double rays with infinitely many turns

Let R be a periodic double ray with infinitely many turns and let $\hat{R}, \tilde{R}$ be disjoint copies of R. By periodicity of R, one can show that identifying a turn of $\hat{R}$ of out-degree 2 and a turn of $\tilde{R}$ of in-degree 2 results in a digraph in which a copy of R has to be completely contained in either $\hat{R}$ or $\tilde{R}$. We use this fact to prove:

Theorem 5.3.6. *Any periodic double ray with infinitely many turns is non-ubiquitous.*

Proof. Let R be any periodic double ray with infinitely many turns and denote the periodicity of R by $p \in \mathbb{N}$. We will construct a digraph D containing arbitrarily many but not infinitely many copies of R.

We set

$$I := \{(n, m) \in \mathbb{N}^2 : n \leq m\}$$

and let $(R(n, m))_{(n,m)\in I}$ be a family of pairwise disjoint copies of R. Let D be the digraph constructed from the disjoint union $\bigsqcup_{(n,m)\in I} R(n, m)$ by identifying pairwise disjoint pairs of vertices such that for any $(n, m), (n', m') \in I$:

(i) no vertices of $R(n, m)$ and $R(n', m')$ have been identified with each other if $m = m'$,

(ii) exactly one vertex of $R(n, m)$ and exactly one vertex of $R(n', m')$ have been identified if $m \neq m'$,

(iii) if $v \in R(n, m)$ and $w \in R(n', m')$ have been identified with each other, then either the out-degree of v in $R(n, m)$ is 2 and the out-degree of v in $R(n', m')$ is 0 or vice versa, and

(iv) two vertices $v \neq w \in R(n, m)$ that have been identified with other vertices have distance at least p in $R(n, m)$.

A graph D satisfying (i) to (iv) can be constructed by enumerating all unordered pairs $\{R(n,m), R(n',m')\}$ of double rays with $m \neq m'$ and recursively identifying suitable turns of the two rays in each pair.

The digraph D contains arbitrarily many disjoint copies of R, as the double rays $R(1,m), \ldots, R(m,m)$ are disjoint for any $m \in \mathbb{N}$ by (i). To prove that D does not contain infinitely many disjoint copies of R, it suffices to show that any copy of R in D is of the form $R(n,m)$ for some $(n,m) \in I$: Then any infinite family of disjoint copies of R in D would contain two rays $R(n,m)$, $R(n',m')$ with $m \neq m'$ by definition of I. However, $R(n,m)$ and $R(n',m')$ are not disjoint in D by (ii).

Suppose for a contradiction that there is a copy $\hat{R}$ of R in D that is not contained in some $R(n,m)$ for $(n,m) \in I$. Then there are $(n,m) \neq (n',m') \in I$ and $v \in V(\hat{R})$ such that one arc of $\hat{R}$ incident with v is contained in $R(n,m)$ and the other arc of $\hat{R}$ incident with v is contained in $R(n',m')$. We assume without loss of generality that the first arc of $v\hat{R}$ is contained in $vR(n,m)$ (and not in $R(n,m)v$, $R(n',m')v$ or $vR(n',m')$).

Since $R(n,m)$ has periodicity p, the arc a of $R(n,m)$ preceding v has the same orientation as the p-th arc a' of $R(n,m)$ succeeding a. Similarly, the arc b of $\hat{R}$ preceding v has the same orientation as the p-th arc b' of $\hat{R}$ succeeding b. As no vertices of distance at most $p-1$ to v in D other than v were identified by (iv), the first p arcs of $v\hat{R}$ coincide with the first p arcs of $vR(n,m)$ and in particular we have $a' = b'$. Hence the arcs a, b either both point towards v or both point away from v, contradicting (iii) since $a \in A(R(n',m'))$ and $b \in A(R(n,m))$. □

5.4 Positive results

In this section we prove Theorem 5.4.2. We need the following lemma, which essentially states that pairs of disjoint out-rays are ubiquitous with an additional property for the roots of the out-rays. In preparation we define: Let U be the disjoint union of two out-rays and X a set of pairs of vertices. We say that U *is rooted in* X if there is $(x,y) \in X$, such that x and y are the roots of the two rays which are the components of U. If $U = \{(x,y)\}$, we also say that U *is rooted in* (x,y).

Lemma 5.4.1. *Let U be the disjoint union of two out-rays, let D be a digraph and let X be a set of pairs of vertices of D. If there exists a thick U-tribe $\mathcal{F}$ in D where all members of $\mathcal{F}$ are rooted in X, then D contains infinitely many disjoint copies of U that are rooted in X.*

We omit the proof of Lemma 5.4.1 since it is very similar to the proof of Theorem 4.3.2.

Theorem 5.4.2. *Any double ray with an odd number of turns is ubiquitous.*

Proof. Let R be a double ray with a (finite) odd number of turns, D a digraph and $\mathcal{E}$ a thick R-tribe in D. We show that D contains infinitely many disjoint copies of R. Let $\hat{R}$ be a finite connected subdigraph of R that contains all turns of R as internal vertices. Since R has an odd number of turns, by deleting the internal vertices of $\hat{R}$ the digraph R falls apart into a disjoint union U of either two in-rays or two out-rays. By Lemma 4.2.1, we may assume that the latter is the case.

Next, we apply Lemma 4.3.3 to D, R, $\hat{R}$ and $\mathcal{E}$, which yields a thick R-tribe $\mathcal{F}$ in D that is forked at $\hat{R}$. Let $\mathcal{F}'$ be the U-tribe resulting from $\mathcal{F}$ by deleting the internal vertices of the copy of $\hat{R}$ in R' from each member R' of $\mathcal{F}$. Let D' be the union of all members of $\mathcal{F}'$. Further, let X be the set of pairs $(x, y) \in V(D) \times V(D)$ for which there is a member R' of $\mathcal{F}$ such that $xR'y$ is the copy of $\hat{R}$ in R'. This means that any member of $\mathcal{F}'$ is rooted in X.

Now we apply Lemma 5.4.1 to D', $\mathcal{F}'$ and X, which yields an infinite set $\mathcal{U}$ of disjoint copies of U in D' that are rooted in X. We join any $U \in \mathcal{U}$ with $xR'y$, where U is rooted in $(x, y) \in X$ and R' is a member of $\mathcal{F}$ such that $xR'y$ is the copy of $\hat{R}$ in R'. Since $\mathcal{F}$ is forked at $\hat{R}$, this gives an infinite family of disjoint copies of R. □

Finally, the main result Theorem 5.1.3 is an easy consequence of Theorems 5.3.1, 5.3.5, 5.3.6 and 5.4.2.

Chapter 6

End spaces and tree-decompositions

6.1 Introduction

In this chapter we settle the question up to which complexity the topological spaces $|G|$ formed by an infinite graph G together with its ends can still be encoded by tree-decompositions of finite adhesion of the underlying graph G.

To state our results more precisely, recall that a *separation* of a graph G is an unordered pair $\{A, B\}$ of sets of vertices in G such that $A \cup B = V(G)$ and G has no edge between $A \setminus B$ and $B \setminus A$, which is equivalent to saying that its *separator* $A \cap B$ separates A from B. The cardinal $|A \cap B|$ is the *order* of the separation $\{A, B\}$ and the sets A, B are its *sides*.

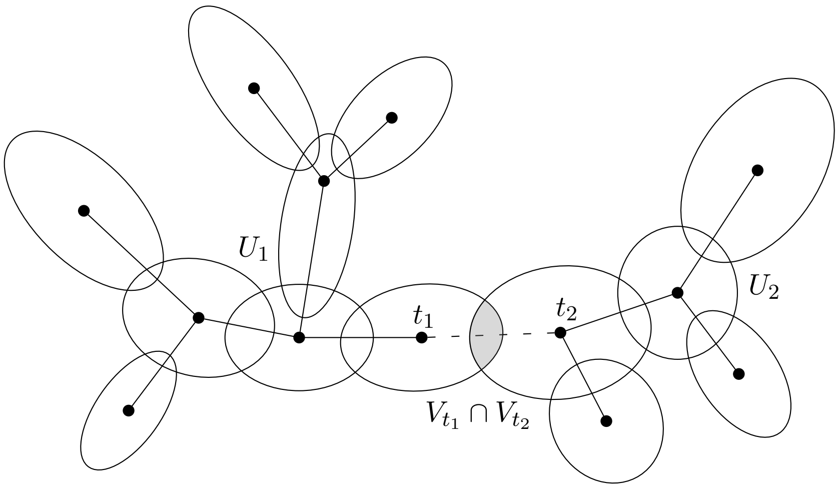

Figure 6.1: $V_{t_1} \cap V_{t_2}$ separates U_1 from U_2.

An *end* of a connected (infinite) graph G is an equivalence class of rays, where two rays are equivalent if for every finite order separation of the graph G, the rays eventually belong both to the same side of the separation. The set of ends is denoted by $\Omega(G)$. The space $|G|$ is the topological space on $G \cup \Omega(G)$ equipped with a natural topology called MTop, described in detail in Section 6.2. For locally finite connected graphs, this $|G|$ is precisely the Freudenthal compactification of G. A longstanding quest in graph theory is to understand end spaces of infinite graphs that are not necessarily locally finite, cf. [22, 24, 25, 53, 54, 56, 67, 68].

Note that the parts of a tree-decomposition $(T, \mathcal{V})$ of a graph G mirror the separation properties of the tree: just like removing any edge $e = t_1t_2$ from T gives rise to two components T_1 and T_2 of $T - e$, so does removing $X_e := V_{t_1} \cap V_{t_2}$ from G separate any part of T_1 from any part of T_2, see Figure 6.1. More formally, writing $U_1 = \bigcup\{V_t : t \in T_1\}$ and $U_2 = \bigcup\{V_t : t \in T_2\}$, we have that $\{U_1, U_2\}$ is a separation of G with separator X_e. If all such separations are of finite order, we say the tree-decomposition has *finite adhesion*.

Now consider how the ends of a graph G interact with a tree-decomposition $\mathcal{T}$ of finite adhesion. As every edge e of T induces a finite order separation $\{A_e, B_e\}$ of G, any end of G has to choose one side of $T - e$, and we may visualize this decision by orienting e accordingly. Then for a fixed end, all the edges point either towards a unique node or towards a unique end of T, see Figure 6.2. In this way, each end of G *lives* in a part of $\mathcal{T}$ or *corresponds* to an end of T, and we may encode this correspondence by a map $f_{\mathcal{T}} : \Omega(G) \to V(T) \cup \Omega(T)$.

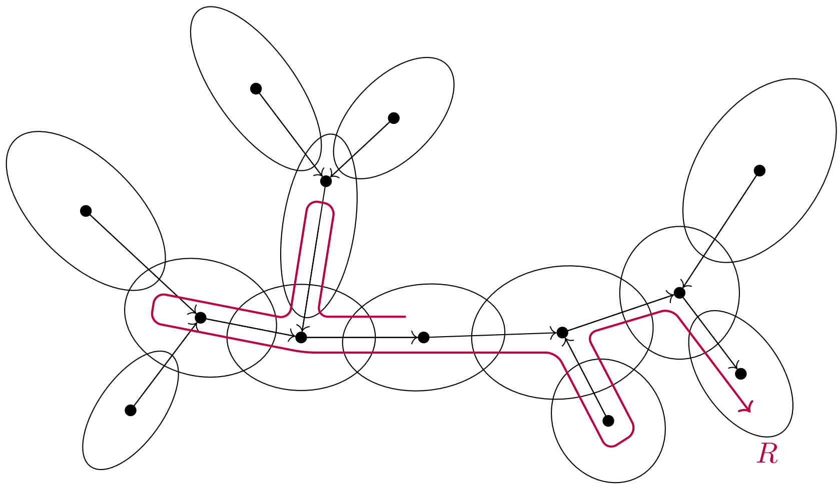

Figure 6.2: A ray R and its corresponding orientation of T

Tree-decompositions of finite adhesion have been used to study the structure of infinite graphs and their ends in e.g. [7, 16–19, 22, 62, 77]. Of course, some tree-decompositions of finite adhesion carry more information about the ends than others. For one, information content may be measured in terms of injectivity of $f_{\mathcal{T}}$. Indeed, a tree-decomposition consisting of a single part contains zero information, whereas a tree-decomposition $\mathcal{T}$ of finite adhesion that *distinguishes* all the ends, i.e. where $f_{\mathcal{T}}$ is injective, contains more information about the end space – although it may still give false hints, as for example ends of T may not represent real ends of G. So even better would be a bijective $f_{\mathcal{T}}$, in which case we say that $\mathcal{T}$ *represents* the ends of G. On the other hand, while the trivial tree-decomposition into a single part always exists, some graphs G, such as the binary tree with one dominating vertex added to every rooted ray (cf. Section 6.10), are too complex to be distinguished or represented by a tree-decomposition of finite adhesion.

Our first main result characterises precisely when these best-case scenarios occur; as a surprising by-product, we obtain that whenever a space $|G|$ can be distinguished by a tree-decomposition of finite adhesion, then it can also be represented. In fact, an even weaker condition suffices: As long as there is some tree-decomposition of finite adhesion into *≤1-ended parts*, i.e. a tree-decomposition such that at most one end is mapped to any given part under $f_{\mathcal{T}}$, we also get a tree-decomposition representing $|G|$.

Let's call a set of vertices $U \subseteq V(G)$ *slender* if its closure $\overline{U} \subseteq |G|$ is scattered of finite Cantor-Bendixson rank; in other words, if successively taking the Cantor-Bendixson derivative of its closure $\overline{U} \subseteq |G|$ yields the empty set after finitely many iterations, cf. Section 6.2.6.

With this notion, our first main result reads as follows.

Theorem 6.1.1. *The following are equivalent for any connected graph G with at least one end:*

(1) There is a tree-decomposition of finite adhesion that represents $\Omega(G)$.

(2) There is a tree-decomposition of finite adhesion that distinguishes $\Omega(G)$.

(3) There is a tree-decomposition of finite adhesion into $\leq$1-ended parts.

(4) $V(G)$ is a countable union of slender sets.

It is clear that any assertion from (1) to (3) implies the next. The idea for (3) ⇒ (4) is that for any fixed integer n, the union over all parts within distance n from the root is a slender set of vertices, and $V(G)$ clearly is a countable union of these sets. Thus, the main contribution behind Theorem 6.1.1 is the implication (4) ⇒ (1), which employs recently developed techniques of *envelopes* from [54, 62] and *rayless normal trees* from [56]. The proof of Theorem 6.1.1 is given in Section 6.8.

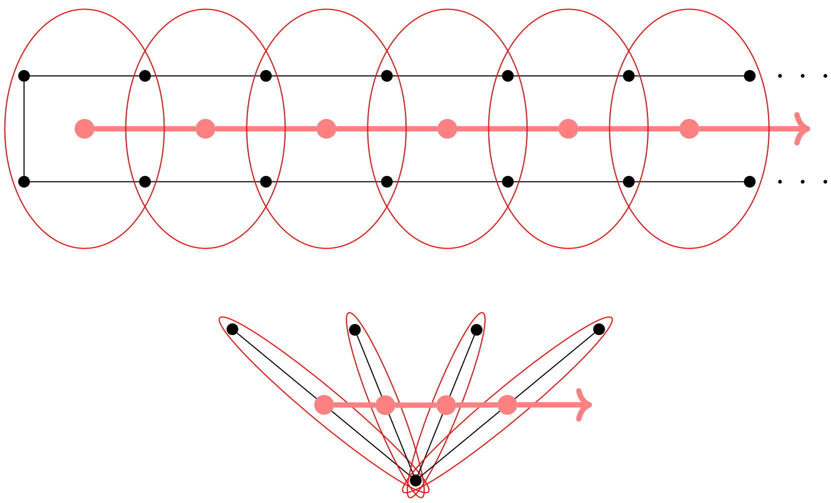

Figure 6.3: Examples of tree-decompositions (in red) of graphs (in black) failing to display their boundaries.

A slightly different way to measure information captured by some tree-decomposition of finite adhesion is motivated by the observation that end spaces of trees are well-understood: They are precisely the completely ultra-metrizable spaces. This suggests preferring tree-decompositions $\mathcal{T}$ where $f_{\mathcal{T}}$ sends as many ends to $\Omega(T)$ as possible. In this case, there is hope to understand the subset $\Psi = f_{\mathcal{T}}^{-1}[\Omega(T)] \subseteq \Omega(G)$ called the *boundary* of the tree-decomposition, with the best case being that $\mathcal{T}$ [*homeomorphically*] *displays* its boundary, meaning that $f_{\mathcal{T}}$ restricts to a bijection [homeomorphism] between Ψ and $\Omega(T)$, cf. Figures 6.3 and 6.4.

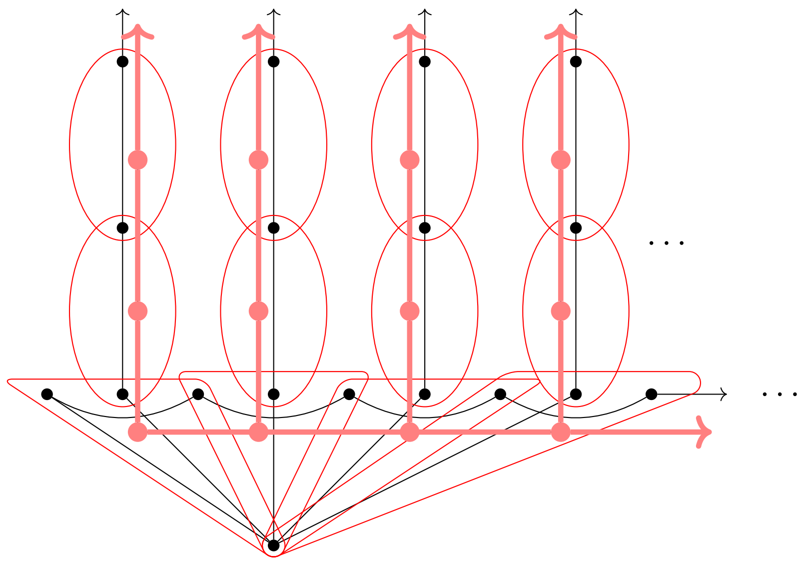

Figure 6.4: Example of a tree-decomposition (in red) that displays all ends of a countable star of rays (in black) but fails to display them homeomorphically.

At first glance, however, it does not seem useful at all when $f_{\mathcal{T}}$ maps all ends of G into $\Omega(T)$ but the function is very much non-injective. However, this information is enough to guarantee a normal spanning tree, from which the space $|G|$ is easily understood. Indeed, given previous work in the field due to Jung and Diestel [23, 24, 46], it is not hard to verify that the following assertions are equivalent, see Theorem 6.7.1 for details:

- There is a tree-decomposition of finite adhesion that (homeomorphically) displays $\Omega(G)$.
- There is a tree-decomposition of finite adhesion with boundary $\Omega(G)$.
- $|G|$ is (completely) metrizable.
- $V(G)$ is a countable union of closed sets in $|G|$.
- G has a normal spanning tree.

Now our second main result provides a local version of the above equivalences, characterising precisely which subsets Ψ of $\Omega(G)$ can be (homeomorphically) displayed. Indeed, a striking, recent result by Carmesin [18] says that it is always possible to display the set of undominated ends of a graph G. In [17], Bürger and Kurkofka partially localized Carmesin's result by constructing tree-decompositions of finite adhesion (with additional desirable properties) that display the boundary ∂U of prescribed infinite sets of vertices $U \subseteq V(G)$ where none of the ends in ∂U are dominated. Carmesin also asked for a characterisation of those pairs of a graph G and a subset $\Psi \subseteq \Omega(G)$ for which G has a tree-decomposition displaying Ψ [18, p. 549]. This problem has also been reiterated in [16, Problem 3.22]. Theorem 6.1.2 below answers this question.

Another set of questions in infinite topological graph theory concerns so-called Ψ-graphs, i.e. subspaces of $|G|$ of the form $|G|_\Psi = G \cup \Psi \subseteq |G|$ for a set of ends $\Psi \subseteq \Omega(G)$. Ψ-graphs have been studied in connection with infinite matroids [7,8,30]: For example, the topological circles (copies of the unit circle S^1) in $|G|_\Psi$ form the cycles of an infinite matroid whenever Ψ belongs to the Borel σ-algebra of $\Omega(G)$ [7].

It turns out that the correct generalisation of the 3rd bullet above about metrizability of $|G|$ involves precisely the property of complete metrizability of Ψ-graphs.

Theorem 6.1.2. *For any connected graph G and a set Ψ of ends of G the following are equivalent:*

(1) There is a tree-decomposition of finite adhesion homeomorphically displaying Ψ.

(2) There is a tree-decomposition of finite adhesion displaying Ψ.

(3) There is a tree-decomposition of finite adhesion with boundary Ψ.

(4) $|G|_\Psi$ is completely metrizable.

(5) Ψ is G_δ in $|G|$.

Note that from Theorem 6.1.2 one easily reobtains the above equivalences in the case $\Psi = \Omega$. Indeed, only item (5) needs to be commented on: For this, note that saying that $\Psi = \Omega$ is G_δ in $|G|$ means $\Psi = \Omega$ is a countable intersection of open sets, which turns out to be equivalent to $V(G)$ being a countable union of closed sets in $|G|$. Also note that Ψ being a G_δ means that Ψ is a fairly simple element of the Borel σ-algebra on $|G|$, and in fact, using Theorem 6.1.2 it is not hard to establish that $|G|_\Psi$ gives an infinite matroid in the special case from [7] where $\Psi \subseteq |G|$ is G_δ.

Carmesin's result that the undominated ends Ψ of any connected graph can always be displayed now follows easily from Theorem 6.1.2: Simply note that fixing any vertex v and considering the set $B_n(v)$ of all vertices in G within graph distance at most n from v, the set Ψ is the intersection of the countably many open sets $O_n = |G| \setminus \overline{B_n(v)}$ (for $n \in \mathbb{N}$) and hence G_δ, see Theorem 6.10.1.

Furthermore, Theorem 6.1.2 also provides tree-decompositions that (homeomorphically) display the undominated ends in the boundary ∂U of any fixed infinite set of vertices $U \subseteq V(G)$, strengthening the above mentioned result by Bürger and Kurkofka from [17]; see Theorem 6.10.2.

A number of natural questions remain on the topic which subsets of ends can be distinguished.

Problem 6.1.3. *Characterise which $\Psi \subseteq \Omega(G)$ can be distinguished.*

Given two distinct ends ω_1, ω_2 of a graph G write $n(\omega_1, \omega_2) \in \mathbb{N}$ for the minimal order of a separation in G that is oriented differently by ω_1 and ω_2. We say that a tree-decomposition $\mathcal{T}$ with decomposition tree T *efficiently distinguishes* a set of ends Ψ if $\mathcal{T}$ distinguishes Ψ with the additional property that for each $\psi_1 \neq \psi_2 \in \Psi$ there is an edge e on the path in T between $f_{\mathcal{T}}(\psi_1)$ and $f_{\mathcal{T}}(\psi_2)$ with $|X_e| = n(\omega_1, \omega_2)$.

Problem 6.1.4. *Characterise which $\Psi \subseteq \Omega(G)$ can be efficiently distinguished.*

An end ω of a graph is called *thin* if all families of disjoint ω-rays are finite, and *thick* otherwise. Our next problem extends a problem of Diestel [22], asking for which graphs there is a tree-decomposition of finite adhesion displaying precisely its thin ends. Carmesin [18] constructed a graph for which there is no such tree-decomposition, and we construct a different counterexample in Example 6.10.7 with help of our characterisation of displayable sets of ends from Theorem 6.1.2. We propose a different way in which a tree-decomposition of finite adhesion might distinguish the thin ends from the thick ends and ask which other bipartitions of $\Omega(G)$ can be distinguished in the same way:

Problem 6.1.5. *Characterise for which bipartitions $\Omega(G) = \Omega_1 \sqcup \Omega_2$ there is a tree-decompositions $\mathcal{T}$ of finite adhesion with $f_{\mathcal{T}}(\Omega_1) \cap f_{\mathcal{T}}(\Omega_2) = \varnothing$.*

We conclude with two problems concerning metrizability in end spaces.

Problem 6.1.6. *Characterise which subspaces $\Psi \subseteq \Omega(G)$ are metrizable or completely metrizable.*

Problem 6.1.7. *Characterise which spaces $|G|_\Psi$ are metrizable.*

For $\Psi = \Omega(G)$, an answer to Problem 6.1.6 is given in [56].

6.2 Preliminaries

6.2.1 Ends

A 1-way infinite path is called a *ray* and the subrays of a ray are its *tails*. Two rays in a graph $G = (V, E)$ are *equivalent* if no finite set of vertices separates them; the corresponding equivalence classes of rays are the *ends* of G. If ω is an end of G and $R \in \omega$, we call R an ω-ray. The set of ends of a graph G is denoted by $\Omega = \Omega(G)$.

The *degree* $\deg(\omega)$ of an end ω is the supremum of the sizes of collections of pairwise disjoint rays in ω; Halin showed that this supremum is always attained, see [26, Theorem 8.2.5]. Ends are called *thin* if they have finite degree, and *thick* otherwise.

We say that a vertex v *dominates* a ray R if there is a subdivided star with centre v and leaves in R. Whenever two rays R and R' are equivalent, a vertex dominates R if and only if it dominates R'. Thus we can say that a vertex *dominates* an end ω if and only if it dominates one ray (and thus all rays) from ω. In this case, we say ω is *dominated*.

6.2.2 Ends and directions

If $X \subseteq V$ is finite and $\omega \in \Omega$, there is a unique component of $G - X$ that contains a tail of every ω-ray. We denote this component by $C(X, \omega) = C_G(X, \omega)$ and say that ω *lives in* $C(X, \omega)$.

A *direction* on G is a function d that assigns to every finite $X \subseteq V$ one of the components of $G - X$ so that $d(X) \supseteq d(X')$ whenever $X \subseteq X'$. For every end ω, the map $X \mapsto C(X, \omega)$ is easily seen to be a direction. Conversely, every direction is defined by an end in this way:

Theorem 6.2.1 (Diestel & Kühn [28]). *For every direction d on a graph G there is an end ω such that $d(X) = C(X, \omega)$ for every finite $X \subseteq V(G)$.*

6.2.3 Star-Comb Lemma

Given a set of vertices U, a *comb attached to U* consists of a ray R together with infinitely many disjoint R–U paths (possibly trivial). A *star attached to U* is a subdivided infinite star with all leaves in U.

Lemma 6.2.2 (Star-Comb Lemma [26, Lemma 8.2.2]). *Let U be an infinite set of vertices in a connected graph G. Then G contains a star or a comb attached to U.*

6.2.4 End spaces

If $X \subseteq V$ is a finite set of vertices and $\omega \in \Omega$, then

$$\Omega(X, \omega) = \Omega_G(X, \omega) = \{\varphi \in \Omega \colon C(X, \varphi) = C(X, \omega)\}$$

denotes the set of all ends that live in $C(X, \omega)$. We put $\hat{C}(X, \omega) = C(X, \omega) \cup \Omega(X, C)$.

The collection of singletons $\{v\}$ for $v \in V(G)$ together with all sets of the form $\hat{C}(X, \omega)$ for finite $X \subseteq V$ and $\omega \in \Omega(G)$ forms a basis for a topology on $V(G) \cup \Omega(G)$. This topology is Hausdorff, and it is *zero-dimensional* in the sense that it has a basis consisting of closed-and-open sets.

Given a set of vertices $U \subseteq V(G)$, we write ∂U for its boundary, i.e. the set of ends in $\overline{U}$. It is well-known that $\omega \in \partial U$ if and only if there is a comb attached to U with spine in ω.

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 $\iota \colon \Omega(H) \to \Omega(G)$ is well-defined. A subgraph $H \subseteq G$ is *end-faithful* if this inclusion map ι is a bijection from $\Omega(H)$ onto $\partial H \subseteq \Omega(G)$.[1]

We now describe one common way to extend this topology on $V(G) \cup \Omega(G)$ to a topology on $|G| = G \cup \Omega(G)$, the graph G together with its ends. This topology, called MTop, has a basis formed by all open sets of G considered as a metric length-space (i.e. every edge together with its endvertices forms a unit interval of length 1, and the distance between

[1] In the literature, the term end-faithful subgraph is sometimes used only for subgraphs $H \subseteq G$ with $\partial H = \Omega(G)$.

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

$$\hat{C}_\varepsilon(X,\omega) := \hat{C}(X,\omega) \cup \mathring{E}_\varepsilon(X, C(X,\omega)),$$

where $\mathring{E}_\varepsilon(X, C(X,\omega))$ denotes the open ball around $C(X,\omega)$ in G of radius $\varepsilon < 1$.

6.2.5 Tree orders and normal trees

Recall that the *tree order* of a tree T with root r is a partial order on $V(T)$ which is defined by setting $u \le v$ if u lies on the unique path rTv from r to v in T. Given $n \in \mathbb{N}$, the *nth level* T^n of T is the set of vertices at distance n from r in T, and by $T^{\le n}$ we denote the union over the first n levels. The *down-closure* of a vertex v is the set $\lceil v \rceil := \{\, u \colon u \le v \,\}$. 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 spanning tree T of a graph G is *normal* in G if the endvertices of every edge of G are comparable in the tree order of T. Normal spanning trees are always end-faithful [26, Lemma 8.2.3].

A rooted, not necessarily spanning, 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 subgraph $H \subseteq G$, a path P in G is said to be an *H-path* if P is non-trivial and meets H exactly in its endvertices. Clearly, if T is spanning, this reduces to the earlier condition, as in this case all T-paths are chords. 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. The following result can be found in [56].

Theorem 6.2.3. *Let G be a connected graph. For every open cover $\mathcal{O}$ of $\Omega(G)$, there is a rayless normal tree T in G such that for every component C of $G - T$ there is a set $O \in \mathcal{O}$ such that $\partial C \subseteq O$.*

Additionally, we need the following easy lemma: We say a set of vertices U in a graph G has *finite adhesion*, if and only if every component of $G - U$ has a finite neighbourhood in U.

Lemma 6.2.4. *Let G be a connected graph and T a rayless normal tree in G. Then T has finite adhesion in G. Moreover, for every finite set $U \subseteq V(G)$ there is a rayless normal tree $T^* \supseteq T$ in G such that $U \subseteq V(T^*)$.*

Proof. For the proof that T has finite adhesion in G, let C be any component of $G - T$. Since T is normal, the neighbourhood of C is a chain in the tree order of T, and this chain is finite because T^* is rayless.

Next, let T^* be a rayless normal tree in G extending the tree T which contains maximally many vertices from U. We show that T^* contains all vertices from U. Suppose for a contradiction that there is a vertex $u \in U$ with $u \notin V(T^*)$ and let C be the component of $G - T^*$ containing u. We showed in the first paragraph of this proof that the neighbourhood

$N(C)$ of C is a finite chain in the tree order of T^*. Let v be its maximal element and v' a neighbour of v in C. Then the union of T^* with the edge vv' and a v'–u path in C is again a rayless normal tree with T as a subgraph, contradicting the maximality of T. □

6.2.6 Topological notions

A subspace Y of a topological space X is *discrete* if every singleton of Y is open in the subspace topology.

A G_δ*-set* of a topological space X is a countable intersection of open sets. An F_σ*-set* is a countable union of closed sets. Note that the complement of a G_δ-set is always a F_σ-set and vice versa.

Lemma 6.2.5. *Let G be a graph and $\Psi \subseteq \Omega(G)$. Then $V(G) \cup \Psi$ is F_σ in $|G|$ if and only if $G \cup \Psi$ is F_σ in $|G|$.*

Proof. The backwards direction follows from that fact that $V \cup \Omega$ is closed in $|G|$, so $V \cup \Psi$ is closed in $G \cup \Psi$, and closed subsets of F_σ-sets are themselves F_σ.

Conversely, assume $V \cup \Psi = \bigcup_{n\in\mathbb{N}} X_n$ is a countable union of closed sets X_n of vertices and ends in $|G|$. Without loss of generality, we have $X_n \subseteq X_{n+1}$. Let $V_n = X_n \cap V(G)$. Then $\bigcup_{n\in\mathbb{N}} V_n = V(G)$ and $G = \bigcup_{n\in\mathbb{N}} G[V_n]$. But then the induced subsets $G[X_n] := G[V_n] \cup X_n$ are also closed in $|G|$, and so $G \cup \Psi = \bigcup_{n\in\mathbb{N}} G[X_n]$ is F_σ in $|G|$, too. □

A set of vertices U in a graph G is *dispersed* if it can be separated from any ray in G by a finite set of vertices. This is equivalent to the property of U being closed in $|G|$.

Given a topological space X, the *derived space* X' is the subspace $X' \subseteq X$ obtained by removing all isolated points from X. By transfinite induction, one defines a decreasing sequence of subsets of X by setting $X^{(0)} = X$, $X^{(\alpha+1)} = (X^{(\alpha)})'$ in the successor case, and $X^{(\lambda)} = \bigcap_{\alpha<\lambda} X^{(\alpha)}$ for all limit ordinals λ. For cardinality reasons, this transfinite sequence must eventually be constant. The smallest ordinal α such that $X^{(\alpha+1)} = X^{(\alpha)}$ is called the *Cantor–Bendixson rank* of X. If $X^{(\alpha)} = \varnothing$ for some ordinal α, then X is *scattered.*

A set of vertices $U \subseteq V(G)$ is *slender* if $X = \overline{U} \subseteq |G|$ satisfies $X^{(n)} = \varnothing$ for some $n \in \mathbb{N}$.

6.3 Tree-decompositions

A [rooted] *tree-decomposition* of a graph G is a pair $\mathcal{T} = (T, \mathcal{V})$ where T is a [rooted] tree and $\mathcal{V} = (V_t : t \in T)$ is a family of vertex sets of G called *parts* such that the following holds:

(T1) for every vertex v of G there exists $t \in T$ such that $v \in V_t$;

(T2) for every edge e of G there exists $t \in T$ such that $e \in G[V_t]$; and

(T3) $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever t_2 lies on the t_1–t_3 path in T.

Let $e = xy$ be any edge of T and let T_x and T_y be the two components of $T - e$ with $x \in T_x$ and $y \in T_y$. Each edge $e = xy$ of T in a tree-decomposition gives rise to a separator $X_e := V_x \cap V_y$ called the separator *induced by* the edge e, which separates $A_x = \bigcup_{t \in T_x} V_t$ from $A_y = \bigcup_{t \in T_y} V_t$. The tree-decomposition has *finite adhesion* if all separators of G induced by the edges of T are finite.

Given a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion of G, any end ω of G orients each edge $e = xy$ of T according to whether ω lives in a component of $G[A_x] - X_e$ or $G[A_y] - X_e$. This orientation of T points towards a node of T or to an end of T, and ω *lives* in that part for that node or *corresponds* to that end, respectively.

Let $f_{\mathcal{T}} \colon \Omega(G) \to V(T) \cup \Omega(T)$ be the function mapping every end of G to the node or end of T that it lives in or corresponds to, respectively. We say that $\mathcal{T}$ *distinguishes* the ends of G if $f_{\mathcal{T}}$ is injective, and it *represents* the ends of G if $f_{\mathcal{T}}$ is bijective.

We call $f_{\mathcal{T}}^{-1}[\Omega(T)]$ the *boundary* of $\mathcal{T}$, and $f_{\mathcal{T}}^{-1}[V(T)]$ the *interior* of $\mathcal{T}$. We say that $\mathcal{T}$ *displays* a subset $\Psi \subseteq \Omega(G)$ if Ψ is the boundary of $\mathcal{T}$ and $f_{\mathcal{T}} \restriction \Psi \to \Omega(T)$ is bijective, and it *homeomorphically displays* Ψ if $f_{\mathcal{T}} \restriction \Psi \to \Omega(T)$ is a homeomorphism. We say that $\mathcal{T}$ [*bijectively*] *distributes* a subset $\Xi \subseteq \Omega(G)$ if Ξ is the interior of $\mathcal{T}$ and $f_{\mathcal{T}} \restriction \Xi \to V(T)$ is injective [bijective]. Finally, we say that $\mathcal{T}$ *realises* [*represents*] a partition $\Omega(G) = \Xi \sqcup \Psi$ of the end space of G if $\mathcal{T}$ [bijectively] distributes Ξ and displays Ψ.

We conclude this section with a sufficient condition for tree-decompositions to (homeomorphically) display their boundary. We say a rooted tree-decomposition $(T, \mathcal{V})$ is *upwards connected* if for every edge $e = xy \in E(T)$ with $x < y$, the induced subgraph $H_e := G[A_y \setminus A_x] = G[A_y] - X_e$ (with A_x, A_y and X_e as above) is non-empty and connected (or equivalently, H_e is a component of $G - X_e$).

Lemma 6.3.1. *Every upwards connected rooted tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion of a graph G homeomorphically displays its boundary.*

Proof. Let Ψ be the boundary of $\mathcal{T}$. We show that $f := f_{\mathcal{T}} \restriction \Psi : \Psi \to \Omega(T)$ is a homeomorphism.

For the proof that f is injective, let $\psi_1 \neq \psi_2 \in \Psi$ and let R_i be the $f(\psi_i)$-ray in T starting in the root of T for $i = 1, 2$. There is a finite vertex set $S \subseteq V(G)$ such that ψ_1 and ψ_2 live in different components of $G - S$. By (T1) there is a finite subtree T' of T containing the root of T such that $S \subseteq \bigcup_{t \in T'} V_t =: G'$. We denote the unique T'–$(T \setminus T')$ edge in R_i by e_i for $i = 1, 2$. Then ψ_i lives in H_{e_i} (as defined above), which is a component of $G - G'$ since $\mathcal{T}$ is upwards connected. Since ψ_1 and ψ_2 live in different components of $G - S$, they also live in different components of $G - G'$. It follows that $H_{e_1} \neq H_{e_2}$. Therefore $e_1 \neq e_2$, $R_1 \neq R_2$, and thus $f(\psi_1) \neq f(\psi_2)$.

Next, for the proof that f is onto, for each end ω of T we find an end $\psi \in \Psi$ such that $f_{\mathcal{T}}(\psi) = \omega$. Let $R = re_0v_1e_1v_2e_2\ldots$ be the ω-ray in T starting in the root of T. We have $\bigcap_{i \in \mathbb{N}} H_{e_i} = \varnothing$ because each H_{e_i} contains only vertices from parts V_t such the distance of t to the root of T is greater than i. In particular, for every finite subset S of $V(G)$ there is a minimal integer i such that $S \cap H_{e_i} = \varnothing$. Since H_{e_i} is connected and non-empty, there is a unique component $d(S)$ of $G - S$ with $H_{e_i} \subseteq d(S)$. The function d defines a direction on G

because the components $H_{e_0} \supseteq H_{e_1} \supseteq \dots$ are nested and non-empty. By Theorem 6.2.1, there is an end ψ of G such that $C_G(S, \psi) = d(S)$ for every finite subset $S \subseteq V(G)$. In particular, we have $d(X_{e_i}) = H_{e_i}$ for the separator X_{e_i} corresponding to e_i, and hence ψ lives in H_{e_i} for all $i \in \mathbb{N}$. Consequently, ψ lies in the boundary of $\mathcal{T}$ and $f(\psi) = \omega$.

We now argue that f is continuous (this part of the argument works for any tree-decomposition displaying Ψ and doesn't yet require upwards connectedness). Indeed, let $\psi \in \Psi$ and $f(\psi) = \omega \in \Omega(T)$. For continuity, consider an arbitrary basic open neighbourhood $\Omega_T(T', \omega)$ of $\omega \in \Omega(T)$. Since T is a tree, there is a unique $C_T(T', \omega)$–T' edge $e = tt'$. Then $X_e = V_t \cap V_{t'}$ is finite since $\mathcal{T}$ had finite adhesion. Now $C_G(X_e, \psi)$ lies completely on one side of the separation (A_e, B_e), and so all ends in $C_G(X_e, \psi)$ orient e towards ω, showing that $f[\Omega_G(X_e, \psi)] \cap \Psi \subseteq \Omega_T(T', \omega)$ as desired.

Finally, we show that f^{-1} is continuous (this part fails without upwards connectedness, cf. Figure 6.4). Let $f^{-1}(\omega) = \psi \in \Psi$ as before and consider a basic open neighbourhood $\Omega_G(S, \psi) \cap \Psi$ of $\psi \in \Psi$. Let T' be a finite subtree of T which contains the root of T and such that $S \subseteq \bigcup_{t \in V(T')} V_t$. Since ψ orients the unique $C_T(T', \omega)$–T' edge e towards ω and H_e is connected, it follows that $H_e = C_G(X_e, \psi) \subseteq C_G(S, \psi)$. Thus, all ends that orient e towards ω live in $C_G(S, \psi)$, giving $f^{-1}[\Omega_T(T', \omega)] \subseteq \Omega_G(S, \psi) \cap \Psi$ as desired. □

6.4 Tree-decompositions displaying all ends

In this section we answer the question which graphs have a tree-decomposition displaying all ends. It turns out that those are exactly the graphs with a normal spanning tree. A characterisation of those graphs by forbidden minors can be found in [64].

Theorem 6.4.1. *The following are equivalent for any connected graph G:*

(1) There is an upwards connected tree-decomposition of finite adhesion with connected parts that homeomorphically displays $\Omega(G)$.

(2) There is a tree-decomposition of finite adhesion displaying $\Omega(G)$.

(3) There is a tree-decomposition of finite adhesion with boundary $\Omega(G)$.

(4) $|G|$ is (completely) metrizable.

(5) $\Omega(G)$ is G_δ in $|G|$.

(6) G has a normal spanning tree.

Proof. The equivalence (5) ⇔ (6) is a well-known result by Jung characterising the existence of normal spanning trees [46]. In Jung's language, a connected graph has a normal spanning tree if and only if $V(G)$ is a countable union of dispersed sets; since dispersed sets are precisely the sets of vertices which are closed in $|G|$, this is equivalent to $V = V(G)$ being F_σ in $|G|$. By Lemma 6.2.5, this is equivalent to G being F_σ in $|G|$, which by taking complements is the same as $\Omega(G)$ being G_δ in $|G|$.

The equivalence (4) ⇔ (6) is due to Diestel [24].[2] The implications (1) ⇒ (2) and (2) ⇒ (3) are trivial.

For (3) ⇒ (5) suppose we have a tree-decomposition $(T, \mathcal{V})$ with root r of finite adhesion with boundary $\Omega(G)$. We claim that $G[D_n]$ is closed, where

$$D_n := \bigcup_{t \in T^{\leq n}} V_t.$$

Indeed, for any end ω of G there is a unique ray $R = t_0 t_1 t_2 \ldots$ starting at the root $t_0 = r$ corresponding to this end. Then $V_{t_n} \cap V_{t_{n+1}}$ is a finite separator that separates D_n from the tails of all ω-rays. Hence no end lives in the closure of D_n, so $G[D_n]$ is closed. It follows from property (T1) and (T2) of a tree-decomposition that $G = \bigcup_{n \in \mathbb{N}} G[D_n]$ is F_σ, so by taking complements in $|G|$, we see that $\Omega(G)$ is G_δ in $|G|$.

Lastly, we show (6) ⇒ (1). Something similar has been done in [23]. Assume that G has a normal spanning tree T with root r. For every vertex t of T, we define $V_t := \lceil t \rceil$ and show that $\mathcal{T} := (T, (V_t)_{t \in T})$ is a tree-decomposition of G of finite adhesion that homeomorphically displays all its ends. Since T is normal, the end vertices of any edge vw of G are comparable in the tree order. If say $v < w$, then e belongs to the part V_w per definition, giving (T2). Further, if a vertex lies in two parts V_v and V_w, it lies in $\lceil v \rceil \cap \lceil w \rceil$ and hence in all V_t for vertices t on the unique v–w path in T. Thus we get property (T3), so we have a tree-decomposition. It is clear that all parts are connected and since all parts are finite, also all adhesion sets are finite. Finally, $\mathcal{T}$ is clearly upwards connected. Therefore it follows from Lemma 6.3.1 that $\mathcal{T}$ homeomorphically displays its boundary, which contains all ends of G since all parts are finite. □

6.5 Envelopes

Let G be a connected graph. An *envelope* for a set of vertices $U \subseteq V(G)$ is a set of vertices $U^* \supseteq U$ of finite adhesion (i.e. such that every component of $G - U^*$ has only finitely many neighbours in U^*) with $\partial U^* = \partial U$. In [54, Theorem 3.2] it is proven that every set of vertices in a connected graph admits a connected envelope.

In the following, however, we need a stronger notion of an envelope that works for a set $X \subseteq V(G) \cup \Omega(G)$ of vertices and ends (and in particular, for a set X consisting of ends only): An *envelope* for such a set $X \subseteq V(G) \cup \Omega(G)$ is a set of vertices $X^* \supseteq X \cap V(G)$ of finite adhesion such that $\partial X^* = \overline{X} \cap \Omega(G)$, where the closure $\overline{X}$ of X is taken in $|G|$.

Theorem 6.5.1. *Any set consisting of vertices and ends in a graph G admits an envelope.*

Proof. Let $X \subseteq V(G) \cup \Omega(G)$ be a given set of vertices and ends in a graph G, and write $V(X) := X \cap V(G)$. Let $\mathcal{R}$ be an inclusionwise maximal set of pairwise disjoint rays of ends in $\overline{X}$. Put

$$X' := V(X) \cup \bigcup_{R \in \mathcal{R}} V(R)$$

[2]For (6) ⇒ (4), Diestel only verifies that his metric is topologically compatible; but it not hard to see that his metric is in fact complete. See also Theorem 6.7.1.

and let $\mathcal{S}$ be the set of all centres of (infinite) stars attached to X'. We will show that

$$X^* := X' \cup \mathcal{S}$$

is an envelope for X. The verification relies on the following two claims:

Claim 6.5.2. *If S is a finite set of vertices and C is a component of $G - S$ such that $X' \cap C$ is finite, then $X^* \cap C = X' \cap C$.*

Only $X^* \cap C \subseteq X' \cap C$ requires proof. For this consider some $v \in \mathcal{S}$. By definition, v is the centre of an infinite star attached to X'. Since S is finite and X' meets C finitely, it follows that $v \notin C$. Hence, $C \cap \mathcal{S} = \varnothing$ and so $X^* \cap C = X' \cap C$ as claimed.

Claim 6.5.3. *If S is a finite set of vertices and C is a component of $G - S$ such that $\overline{C} \cap X = \varnothing$, then $X^* \cap C$ is finite.*

To see the claim, consider some finite set of vertices S, and assume that C is a component of $G - S$ such that $\overline{C}$ avoids X. First, we show that $\overline{C} \cap \overline{X} = \varnothing$. For this, observe that the set $\overline{C} \cup \mathring{E}_{1/2}(S, C)$ is open and disjoint from X and so it is disjoint from $\overline{X}$. In particular, $\overline{C}$ is disjoint from $\overline{X}$. Hence every ray $R' \in \mathcal{R}$ meets C finitely. Furthermore, every ray from $\mathcal{R}$ which meets C also meets S, and since the rays in $\mathcal{R}$ are pairwise disjoint, at most $|S|$ rays from $\mathcal{R}$ meet C. So $\bigcup_{R \in \mathcal{R}} V(R)$ meets C finitely, and hence so does X'. By Claim 6.5.2, also $X^* \cap C = X' \cap C$ is finite. This establishes the claim.

To see $\partial X^* = \overline{X} \cap \Omega(G)$, we show both inclusions separately. For $\supseteq$ consider any end $\varepsilon \notin \partial X^*$. Then $C(S, \varepsilon) \cap X^* = \varnothing$ for some finite set of vertices S. Consider a ray R in ε that is completely contained in $C(S, \varepsilon)$. Then R is disjoint from any ray in $\mathcal{R}$. By maximality of $\mathcal{R}$, this means that $\varepsilon \notin \overline{X}$.

For $\subseteq$ consider any end $\varepsilon \notin \overline{X}$. Then there is a finite set of vertices S such that $\hat{C}(S, \varepsilon)$ avoids X. By Claim 6.5.3, also X^* intersects $C(S, \varepsilon)$ finitely, witnessing $\varepsilon \notin \partial X^*$.

To see that X^* has finite adhesion, suppose for a contradiction that there is a component C of $G - X^*$ with infinite neighbourhood. Then by a routine application of the Star-Comb Lemma 6.2.2, we either find a star or a comb attached to X^* whose centre v or spine R is contained in C. The ray case results in an immediate contradiction as follows: If ε denotes the end with $R \in \varepsilon$, then the comb attached to X^* with spine R witnesses that $\varepsilon \in \partial X^*$. Since $\partial X^* = \overline{X} \cap \Omega(G)$ by the earlier observation, we get $R \in \varepsilon \in \overline{X}$. But then the existence of R contradicts the maximality of $\mathcal{R}$.

In the star case, note that for all finite sets of vertices S disjoint from v, the component C of $G - S$ containing v meets X^* infinitely. Then C also meets X' infinitely by Claim 6.5.2. But then it is straightforward to inductively construct a star with centre v attached to X', violating the maximality of $\mathcal{S}$. The completes the proof that X^* is an envelope for X. □

Note that the envelopes constructed in Theorem 6.5.1 are in general neither connected nor end-faithful. But we can easily obtain both properties with the following construction.

For a given subgraph $H \subseteq G$ of finite adhesion, we define a *torso-extension* $H' \supseteq H$ as follows: First, we make H induced. Then for each component C of $G - H$, we let

$T_C \subseteq G[C \cup N(C)]$ be a finite tree such that all vertices from $N(C)$ are leaves of T_C. We add all these T_C to H to obtain H'.

Lemma 6.5.4. *Let G be connected. Whenever $H \subseteq G$ is a subgraph of finite adhesion, then every torso-extension H' is an end-faithful connected subgraph of G of finite adhesion with $\partial H' = \partial H$.*

Proof. Since inside of each component of $G - H$ we only add a finite subgraph to H, also H' has finite adhesion.

By construction, every vertex of $H' \setminus H$ is connected via a finite path in H' to a vertex of H. Hence for connectivity of H' it remains to show that there is a path in H' between every two vertices $v, w \in H$.

Since G is connected, there is a v–w path P in G. We consider P as a sequence of edges between vertices of H and segments inside of components C of $G - H$ together with their end-vertices in $N(C)$. After replacing each of those segments in a component C by a path in T_C between the same end-vertices, we obtain a finite v–w walk P' contained in H'. So H' is connected.

To see that $\partial H' = \partial H$, only $\subseteq$ requires proof. If $\omega \notin \partial H$, then ω lives in a unique component C of $G - H$. Since $H' \cap C$ is finite it follows that ω also lives in a unique component C' of $G - H'$ with $C' \subseteq C$ and hence $\omega \notin \partial H'$ by finite adhesion of H'.

We now argue that H' contains an ω-ray for every end ω in $\partial H' = \partial H$. Suppose without loss of generality that $H \neq \varnothing$ and fix any ω-ray $R = r_0 r_1 r_2 \ldots$ in G with $r_0 \in V(H)$. By finite adhesion of H, the ray R contains infinitely many vertices of H. We will construct a ray $R' \subseteq H'$ that meets R infinitely as follows: If $R \subseteq H'$, there is nothing to do. Otherwise, let r_{n_0} be the first vertex on R outside of H', and consider the component $C_0 \ni r_{n_0}$ of $G - H$. Let r_{k_0} be the last vertex of R in C_0. Replace $r_{n_0-1} R r_{k_0+1}$ by an r_{n_0-1}–r_{k_0+1} path P_0 in $T_{C_0} \subseteq H'$ and call the resulting ray R_1. Note that $R_1 \cap H \subseteq R \cap H$. Now we iterate the same step for R_1 to find a new ray R_2 and so on. This yields a sequence of rays $R_1, R_2, R_3, \ldots$ with $R_n \cap H \subseteq R \cap H$ and which agree on larger and larger initial segments contained in H'. The union of these segments is a ray $R' \subseteq H'$ with $R' \cap H \subseteq R \cap H$, so R' is an ω-ray in H' as desired.

To see that H' is end-faithful, it remains to show that any two rays R_1 and R_2 in H' that are equivalent in G are also equivalent in H'. By assumption there is a collection $\mathcal{P}$ of infinitely many disjoint R_1–R_2 paths in G. We will find infinitely many such paths in H'. Let P be an R_1–R_2 path in G with endvertices $r_1 \in R_1$ and $r_2 \in R_2$. As in the second paragraph, we find a r_1–r_2 walk P' in H'.

Consider the finitely many components of $G - H$ that meet P' and delete from $\mathcal{P}$ all paths that meet one of these components – by finite adhesion, $\mathcal{P}$ remains infinite. So we can find another R_1–R_2 path in H' disjoint to the first one. Iterating this construction, we find infinitely many disjoint R_1–R_2 paths in H', showing that R_1 and R_2 in are also equivalent in H'. □

A result similar to Lemma 6.5.4 for so-called tree sets is proved in [33, Section 2.6].

Corollary 6.5.5. *Any set consisting of vertices and ends in a connected graph G has a connected, end-faithful envelope.*

Whenever we refer to *the envelope* of X inside a connected graph G, we assume that we fixed one possible end-faithful connected choice and call it $\mathcal{E}_G(X)$.

6.6 From topology to tree-decompositions

In this section we employ the envelope technique in order to construct a tree-decomposition of finite adhesion adapted to some prescribed topological information. Roughly, given an infinite graph G and an increasing sequence of closed subsets $X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$ in $|G|$ such that $V(G) \subseteq \bigcup_{n\in\mathbb{N}} X_n$, we construct a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion such that precisely the ends of X_n live in parts indexed by the first n levels of T, and all other ends get displayed.[3]

However, we also want a device that ensures that all ends of some prescribed subcollection Δ of ends in $\bigcup_{n\in\mathbb{N}} X_n$ live in pairwise distinct parts of $\mathcal{T}$. It turns out that this can be achieved provided that each $\Delta_n := \Delta \cap (X_n \setminus X_{n-1})$ is a discrete set.

Lemma 6.6.1. *Let G be a graph and $\Xi \subseteq \Omega(G)$. Suppose that there is a sequence $X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$ of subsets of $V(G) \cup \Xi$ that are closed in $|G|$ with $V(G) \cup \Xi = \bigcup_{n\in\mathbb{N}} X_n$. Denote $\Xi_n := X_n \cap \Omega(G)$ and let Δ_n be a discrete subset of $\Xi_n \setminus \bigcup_{i<n} \Xi_i$ for all $n \in \mathbb{N}$. Then there exists a sequence of induced subgraphs of finite adhesion $G_0 \subseteq G_1 \subseteq G_2 \subseteq \dots$ of G such that the following holds for all $n \in \mathbb{N}$:*

(i) For every component C of $G - G_n$, the set $(C \cap G_{n+1}) \cup N(C)$ is connected in G;

(ii) $\partial G_{3n+2} \subseteq \Xi_n \setminus \Delta_n$;

(iii) for every component C of $G - G_{3n+2}$, there is at most one end from Δ_n contained in ∂C;

(iv) $X_n \cap V(G) \subseteq V(G_{3n+3})$;

(v) $\partial G_{3n+3} = \Xi_n$.

Proof. Set $G_0 := \varnothing$. We will inductively define subgraphs $G_0, G_1, \dots$ of G all of finite adhesion so that (i) – (v) are satisfied.

Every step of the construction follows the same general pattern: To construct G_{n+1} from G_n consider the current set $\mathcal{C}_n$ of components of $G - G_n$. For every $D \in \mathcal{C}_n$ we consider the subgraph $\tilde{D} := G[D \cup N(D)]$ of G. Each time we will define a set of vertices $V_D \subseteq V(\tilde{D})$ of finite adhesion in $\tilde{D}$ containing $N(D)$. Then also $G_{n+1} := G_n \cup \bigcup_{D\in\mathcal{C}_n} V_D$ has finite adhesion in G since any component C of $G - G_{n+1}$ is also a component of $\tilde{D} - V_D$ for some $D \in \mathcal{C}_n$. Furthermore, we will make sure that V_D is connected so that (i) is satisfied.

[3] In the actual proof, we arrange for technical reasons that the ends of X_n live precisely in parts indexed by the first $3n + 3$ levels of T.

Next, we make two observations concerning the end space of $\tilde{D}$, which both follow from the fact that $N(D)$ is finite: Firstly, we have $\partial D = \partial \tilde{D}$ in G, and secondly, the inclusion map ι as mentioned in Section 6.2.4 is a homeomorphism from $\Omega(\tilde{D})$ to $\partial D \subseteq |G|$. Via this homeomorphism, we will in the following identify the spaces $\Omega(\tilde{D})$ and $\partial D \subseteq |G|$.

Now for the actual construction of the sequence $G_0, G_1, \ldots,$ we proceed in steps of three. Suppose that G_{3n} has already been defined. We demonstrate how to recursively construct

$$G_{3n} \rightsquigarrow G_{3n+1} \rightsquigarrow G_{3n+2} \rightsquigarrow G_{3n+3} = G_{3(n+1)}$$

in order to satisfy $(i)-(v)$ for the three indices $3n+1$, $3n+2$ and $3n+3$.

1. Step $3n \rightsquigarrow 3n+1$**.**
Let D be any component from $\mathcal{C}_{3n}$. Since Δ_n is discrete in $|G|$, also $\Delta_n \cap \partial D$ is discrete in ∂D. Thus there is a set $\mathcal{O}_D = \{O_\omega : \omega \in \Delta_n \cap \partial D\}$ of open subsets of ∂D with $O_\omega \cap \Delta_n = \{\omega\}$ for all $\omega \in \Delta_n \cap \partial D$. Applying Corollary 6.5.5 we consider the envelope

$$V_D := \mathcal{E}_{\tilde{D}}(((\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D) \cup N(D)),$$

which is a connected vertex set of finite adhesion in $\tilde{D}$ (cf. Figure 6.5).

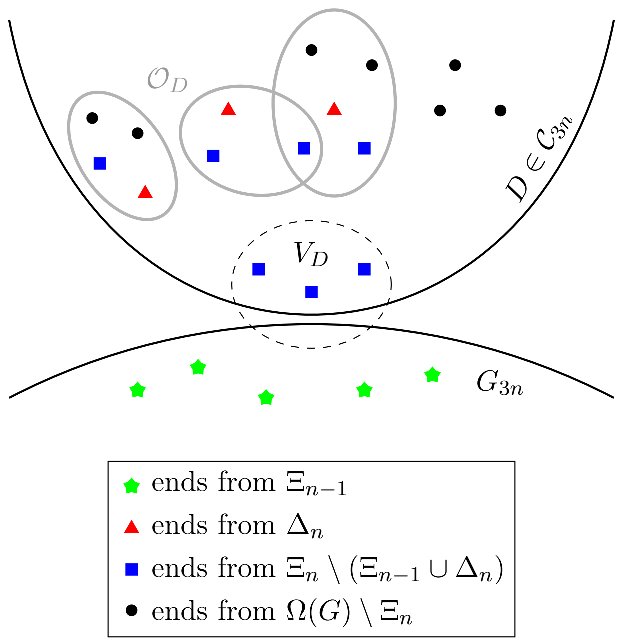

Figure 6.5: Construction step $3n \rightsquigarrow 3n+1$.

We now determine which ends are contained in ∂V_D. Since both X_n and $\Omega(G)$ are closed in $|G|$, also $\Xi_n = X_n \cap \Omega(G)$ is closed in $|G|$. Hence $(\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D$ is closed in the subspace ∂D of $|G|$. Since $N(D)$ is finite and therefore does not have any ends in its closure, it follows from the definition of an envelope and our identification of $\Omega(\tilde{D})$ with ∂D that

$$\partial V_D = \overline{((\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D) \cup N(D))} \cap \partial D = \overline{(\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D} = (\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D.$$

Hence by (v) for G_{3n}, the graph $G_{3n+1} = G_{3n} \cup \bigcup_{D \in \mathcal{C}_{3n}} V_D$ satisfies

(vi) $\partial G_{3n+1} \subseteq \Xi_n \setminus \Delta_n$.

Next, we show that

(vii) for every component C of $G - G_{3n+1}$, there is an open cover $\mathcal{O}$ of ∂C such that each set from $\mathcal{O}$ contains at most one end from Δ_n.

Let C be a component of $G - G_{3n+1}$ and D' the component of $G - G_{3n}$ with $C \subseteq D'$. We show that *(vii)* is fulfilled with

$$\mathcal{O} := \{O \cap \partial C : O \in \mathcal{O}_{D'}\} \cup \{\partial C \setminus \Xi_n\}.$$

Clearly, all sets in $\mathcal{O}$ are open in ∂C and contain at most one end from Δ_n. For the proof that $\partial C \subseteq \bigcup \mathcal{O}$, we observe that C and $V_{D'}$ are disjoint and the neighbourhood of C is finite. Therefore, ∂C and $\partial V_{D'}$ are disjoint. Since $\partial V_{D'} = (\Xi_n \cap \partial D') \setminus \bigcup \mathcal{O}_{D'}$, we have $\partial C \cap \Xi_n \subseteq \bigcup \mathcal{O}_{D'}$ and therefore $\partial C \subseteq \bigcup \mathcal{O}$.

2. Step $3n + 1 \rightsquigarrow 3n + 2$**.**
Let D be any component from $\mathcal{C}_{3n+1}$. By *(vii)* there exists an open cover $\mathcal{O}$ of ∂D such that each set from $\mathcal{O}$ contains at most one end from Δ_n (cf. Figure 6.6). Then by Theorem 6.2.3, there is a rayless normal tree T in $\tilde{D}$ such that for every component C' of $\tilde{D} - T$ there is a set $O \in \mathcal{O}$ with $\partial C' \subseteq O$. By Lemma 6.2.4 there exists a rayless normal tree T^* in $\tilde{D}$ such that $V(T) \cup N(D) \subseteq V(T^*)$. We define $V_D := V(T^*)$. Then every component C of $G - V_D$ is contained in a component C' of $G - T$ and thus there is a set $O \in \mathcal{O}$ with $\partial C \subseteq O$. Then by *(vii)*, C contains at most one end from Δ_n. Hence $G_{3n+2} = G_{3n+1} \cup \bigcup_{D \in \mathcal{C}_{3n+1}} V_D$ satisfies *(iii)*. Furthermore, T^* has finite adhesion in $\tilde{D}$ by Lemma 6.2.4. Finally, normal trees are end-faithful by [26, Lemma 8.2.3], so from the fact that T^* is rayless it follows that $\partial T^* = \varnothing$. Therefore $\partial G_{3n+2} = \partial G_{3n+1}$ and *(ii)* is a consequence of *(vi)*.

3. Step $3n + 2 \rightsquigarrow 3n + 3$**.**
Again let D be any component from $\mathcal{C}_{3n+2}$ (the components of $G - G_{3n+2}$). We define

$$V_D := \mathcal{E}_{\tilde{D}}((X_n \cap \overline{D}) \cup N(D)).$$

Then it follows from the definition of an envelope that

$$X_n \cap V(\tilde{D}) \subseteq ((X_n \cap \overline{D}) \cup N(D)) \cap V(\tilde{D}) \subseteq V_D.$$

Therefore $G_{3n+3} = G_{3n+2} \cup \bigcup_{D \in \mathcal{C}_{3n+2}} V_D$ satisfies *(iv)*. Furthermore, since $N(D)$ is finite and X_n is closed, we have

$$\partial V_D = \overline{((X_n \cap \overline{D}) \cup N(D))} \cap \partial D = X_n \cap \partial D = \Xi_n \cap \partial D.$$

Then together with *(ii)* we obtain $\partial G_{3n+3} = \Xi_n$ which proves *(v)*. □

Theorem 6.6.2. *Let G be a connected graph and $\Xi \subseteq \Omega(G)$. Suppose that there is a sequence $X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$ of subsets of $V(G) \cup \Xi$ that are closed in $|G|$ with $V(G) \cup \Xi = \bigcup_{n \in \mathbb{N}} X_n$. Denote $\Xi_n := X_n \cap \Omega(G)$ and let Δ_n be a discrete subset of $\Xi_n \setminus \bigcup_{i<n} \Xi_i$ for all $n \in \mathbb{N}$. Then there is an upwards connected tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion with connected parts which homeomorphically displays $\Omega(G) \setminus \Xi$ such that the boundary of every part contains at most one end from $\bigcup_{n \in \mathbb{N}} \Delta_n$.*

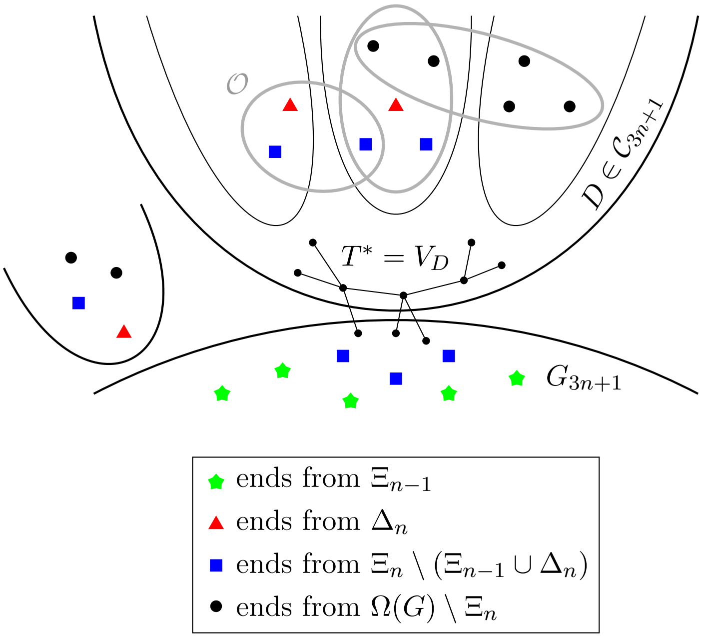

Figure 6.6: Construction step $3n+1 \rightsquigarrow 3n+2$.

Proof. Let $G_0 \subseteq G_1 \subseteq G_2 \subseteq \dots$ be the sequence from Lemma 6.6.1 with properties *(i)* – *(v)* and suppose without loss of generality that $G_0 = \varnothing$. This sequence gives rise to a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion and into connected parts as follows: Write $\mathcal{C}_n$ for the set of components of $G - G_n$. We define a tree order $\leq_T$ on $T := \bigsqcup_{n \in \mathbb{N}} \mathcal{C}_n$ as follows: For all $C_n \in \mathcal{C}_n$ and $C_m \in \mathcal{C}_m$, let $C_n \leq_T C_m$ if and only if $C_n \supseteq C_m$ and $n \leq m$; this will be our decomposition tree. Note that $G_0 = \varnothing$ ensures T has a root whose associated part is G. The part corresponding to a node $C \in \mathcal{C}_n$ of T will be $N(C) \cup (C \cap G_{n+1})$ (which is precisely the set V_C from the proof of Lemma 6.6.1). Then it is readily checked that all properties (T1) – (T3) of a tree-decomposition are satisfied, in particular (T1) holds by *(iv)*. All parts of $\mathcal{T}$ are connected by *(i)*.

It is clear from the construction that $\mathcal{T}$ is upwards connected. Furthermore, by *(v)* the interior of $\mathcal{T}$ is Ξ and hence its boundary is $\Omega(G) \setminus \Xi$. Therefore $\mathcal{T}$ homeomorphically displays $\Omega(G) \setminus \Xi$ by Lemma 6.3.1.

It is left to show that in every part of $\mathcal{T}$ there lives at most one end from $\bigcup_{n \in \mathbb{N}} \Delta_n$. For any $n \in \mathbb{N}$, we have $\Delta_n \subseteq \partial G_{3n+3} \setminus \partial G_{3n+2}$ by *(ii)* and *(v)*.

Since this inclusion holds for all $n \in \mathbb{N}$, it follows that $\partial G_{3n+3} \setminus \partial G_{3n+2}$ does not contain ends from $\Delta_{n'}$ for any $n' \neq n$. Furthermore, by *(iii)* every component in $\mathcal{C}_{3n+2}$ contains at most one end from Δ_n in its boundary. Hence all ends from Δ_n are contained in the boundaries of parts of the form $N(C) \cup (C \cap G_{3n+3})$ for $C \in \mathcal{C}_{3n+2}$, and in the boundary of every such part there is no end from $\Delta_{n'}$ for any $n' \neq n$ and at most one end from Δ_n. This finishes the proof. □

6.7 Tree-decompositions displaying sets of ends

In this section we prove our characterisation announced in Theorem 6.1.2 of *displayable* subsets of $\Omega(G)$, i.e. subsets which can be (homeomorphically) displayed by a tree-decomposition of finite adhesion.

Theorem 6.7.1. *For any connected graph G and any set Ψ of ends of G the following are equivalent:*

(1) There is an upwards connected tree-decomposition of finite adhesion with connected parts that homeomorphically displays Ψ.

(2) There is a tree-decomposition of finite adhesion displaying Ψ.

(3) There is a tree-decomposition of finite adhesion with boundary Ψ.

(4) $|G|_\Psi$ is completely metrizable.

(5) Ψ is G_δ in $|G|$.

Proof. We demonstrate the following sequence of implications:

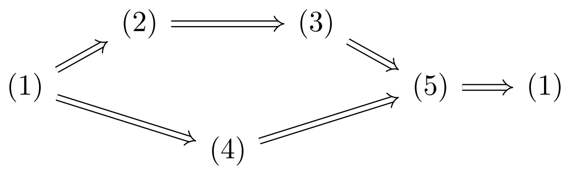

The implications (1) $\Rightarrow$ (2) $\Rightarrow$ (3) are trivial.

(1) $\Rightarrow$ (4): Let $(T, \mathcal{V})$ be a tree-decomposition of finite adhesion of G homeomorphically displaying Ψ with a fixed root r of T. We begin by defining a complete metric d_T on $V(T) \cup \Omega(T)$. Assign to every $e \in E(T)$ a number $\ell(e)$: If $e \in E(T)$ is a T^n–T^{n+1} edge (i.e. an edge between level n and level $n+1$ of T), we set $\ell(e) = 1/2^n$. If P is a (possibly infinite) path in T, we say that the finite number $\sum_{e \in E(P)} \ell(e)$ is the *length* of P. Now we define $d_T(x, y)$ for all $x, y \in V(T) \cup \Omega(T)$: If x and y are both vertices, let $d_T(x, y)$ be the length of the unique x–y path in T. If x is a vertex and y is an end, then let $d_T(x, y)$ be the length of the unique ray from y which starts in x. Similarly, if both x and y are ends, let $d_T(x, y)$ be the length of the unique double ray in T between x and y. It is straight-forward to check that d_T defines a complete metric on $V(T) \cup \Omega(T)$.

We now use d_T to define a metric d on $G \cup \Psi$. For every vertex $v \in V(G)$, let v_T be the least vertex of T with respect to the tree order such that v is contained in the part V_{v_T} (this is well-defined according to (T3)). Additionally, for every end $\omega \in \Psi$, let ω_T be the end of T which ω corresponds to. For all $x, y \in V \cup \Psi$, we define

$$d(x,y) = \begin{cases} 0 & \text{if } x = y, \\ 1/2^n & \text{if } x \neq y \in V(G) \text{ and } x_T = y_T \text{ lies in the } n\text{th level of } T, \\ d_T(x_T, y_T) & \text{if } x_T \neq y_T. \end{cases}$$

Next, we prove that d is a metric on $V(G) \cup \Psi$. It is clear that $d(x,x) = 0$ and $d(x,y) > 0$ for all $x \neq y$ and that d is symmetric. We show that triangle inequality holds: Let x, y, z be pairwise distinct elements of $V(G) \cup \Psi$. We need to show that

$$d(x,z) \leq d(x,y) + d(y,z). \qquad (*)$$

Clearly, $(*)$ holds if $x_T = y_T = z_T$. If $x_T = y_T \neq z_T$, then $d(x,z) = d(y,z)$ and hence $(*)$ follows. A similar argument works if $y_T = z_T$. Next, suppose that $x_T = z_T \neq y_T$ and let n be the level of x_T in T. Then $d(x,z) = 1/2^n$ and since $\ell(e) \geq 1/2^n$ for every edge e of T with endvertex x_T also $d(x,y) \geq 1/2^n$, which proves $(*)$. Finally, if x_T, y_T and z_T are pairwise distinct, then $(*)$ follows from the triangle inequality for d_T. This finishes the proof of $(*)$.

For the proof that d is complete, let $(x_n)_{n\in\mathbb{N}}$ be a Cauchy-sequence in $V(G) \cup \Psi$. Hence $((x_n)_T)_{n\in\mathbb{N}}$ is a Cauchy-sequence in T because $d_T(v_T, w_T) \leq d(v,w)$ for all $v, w \in V(G) \cup \Psi$. If $((x_n)_T)_{n\in\mathbb{N}}$ is eventually constant, then $(x_n)_{n\in\mathbb{N}}$ is eventually contained in V_t for some $t \in V(T)$. If t lies in the nth level of T, then $d(v,w) \geq 1/2^n$ for all $v \neq w \in V_t$. Hence also $(x_n)_{n\in\mathbb{N}}$ is eventually constant. Otherwise, if $((x_n)_T)_{n\in\mathbb{N}}$ is not eventually constant, then $((x_n)_T)_{n\in\mathbb{N}}$ converges to an end ω of T and thus $(x_n)_{n\in\mathbb{N}}$ converges to the end of G which corresponds to ω. Finally, we extend d to a complete metric on $G \cup \Psi$ by relating every edge vw of G linearly to a real closed interval of length $d(v,w)$. We omit the details.

It is left to show that the metric d induces the subspace topology on $G \cup \Psi$ inherited from $|G|$. We need to show for any given $x \in G \cup \Psi$ that

(†) every MTop-basic open neighbourhood of x in $G \cup \Psi$ contains an open ε-ball around x with respect to d, and vice versa.

This is clear if x is an inner point of an edge. Next, let $x \in V(G)$ be a vertex and n the level of x_T in T. Then (†) is true because every edge of G which has x as an endvertex has length at least $1/2^n$ and at most 1.

Now suppose that $x \in \Psi$ and let $\hat{C}_\varepsilon(S,x)$ be a basic open neighbourhood of x in $|G|$ for some $\varepsilon \leq 1$. Let n be the maximum level of T containing a vertex s_T for some $s \in S$. We show that the open ball B in $|G|$ with respect to d with radius $\varepsilon/2^n$ and centre x is a subset of $\hat{C}_\varepsilon(S,x)$. First, consider the open ball B' in T with respect to the metric d_T with radius $\varepsilon/2^n$ and centre x', where x' is the end of T which x corresponds to. Let e be the edge of T which is contained in the normal x'-ray in T and connects a node u_n form the nth level of T to a node u_{n+1} from the $n+1$st level. Then B' is completely contained in the closure of the component D of $T - e$ with $u_{n+1} \in V(D)$ since

$$d_T(u_{n+1}, x') = \sum_{i \geq n+1} 1/2^i = 1/2^n \geq \varepsilon/2^n.$$

In particular, every vertex in B' lies in the $n+1$st level of T or above. Next, it follows from the definition of the metric d that every vertex in B is contained in a part V_t with $t \in B' \subseteq \overline{D}$, but no vertex of B can be contained in a part V_t such that the level of t in T is at most n. Therefore all vertices in B and similarly also all ends in B are contained in $\overline{H_e}$, where H_e is the subgraph of G from the definition of upwards connectedness. Since H_e is disjoint from S, connected by upwards connectedness of $\mathcal{T}$, and x orients e towards x', we have $\overline{H_e} \subseteq \hat{C}(S, \omega)$. Hence all vertices and ends in B and all edges with both endvertices in B are contained in $\hat{C}_\varepsilon(S, x)$; it is left to show the same for points of edges in B with only one endvertex in B. Every such edge f, however, has its other endvertex in V_{u_n} by (T3), and as u_n lies in the nth level of T, the length of f with respect to d is at least $1/2^n$. Recall that any point p on f in B has distance less than $\varepsilon/2^n$ to x and therefore also to the end vertex of f in B. Thus p is contained in $\hat{C}_\varepsilon(S, x)$, as desired.

Conversely, let B be an open ε-ball around x with respect to d of radius $0 < \varepsilon \leq 1$. Let $\omega \in \Omega(T)$ be the end of T corresponding to x and R the rooted ω-ray in T. Choose $n \in \mathbb{N}$ such that $1/2^n < \varepsilon$ and let $t^i \in V(T)$ be the node in $R \cap T^i$ for $i \in \{n+2, n+3\}$. Then define S as the separator induced by the edge $t^{n+2}t^{n+3}$ of T in G. Now $C := \hat{C}_{1/2^{n+1}}(S, x)$ is a subset of B: Let y be any point in C; we have to show that $d(y, x) < \varepsilon$. First suppose that $y \in C(S, x)$ and let w be a vertex from the part $V_{t^{n+3}}$. For any point $z \in C(S, x)$ we have

$$d(w, z) \leq \sum_{i \geq n+3} 1/2^i = 1/2^{n+2}.$$

Hence

$$d(y, x) \leq d(y, w) + d(w, x) \leq 1/2^{n+2} + 1/2^{n+2} = 1/2^{n+1} < \varepsilon.$$

Next, suppose that y is an inner point of an S–$C(S, x)$ edge with endvertex v in $C(S, x)$. We have seen above that $d(v, x) \leq 1/2^{n+1}$. Hence it follows from the choice of C that

$$d(y, x) \leq d(y, v) + d(v, x) \leq 1/2^{n+1} + 1/2^{n+1} < \varepsilon$$

which proves $C \subseteq B$.

(4) $\Rightarrow$ (5): Assume that $|G|_\Psi$ is completely metrizable. We claim that

- Ψ is G_δ in $|G|_\Psi$, and
- $|G|_\Psi$ is G_δ in $|G|$.

This implies (5) as being G_δ is transitive.

Since closed subsets of metrizable spaces are always G_δ [34, Corollary 4.1.12], we get that Ψ is G_δ in $|G|_\Psi$. Next, by a well-known result of Čech [34, Theorem 4.3.26] all completely metrizable spaces, and so in particular $|G|_\Psi$, are Čech-complete, and by [34, Exercise 3.9.A], all Čech-complete spaces are G_δ in their closures. Thus we conclude that $|G|_\Psi$ is G_δ in its closure $|G|$.

(3) $\Rightarrow$ (5): Let $(T, \mathcal{V})$ be a tree-decomposition of finite adhesion of G with boundary Ψ. Fix a root r of T and denote by E_n the set of all edges between the nth and $n+1$st level of T. For every edge $e \in E_n$, let (A_e, B_e) be the respective separation of G such that $V_r \subseteq A_e$ and let $S_e = A_e \cap B_e$ be the corresponding finite adhesion set. Note that A_e contains every part V_t with $t \in T^{\leq n}$. We denote

$$\mathcal{C}_e := \bigcup \{\hat{C}_{1/2}(S_e, \omega) : \omega \in \partial B_e\}.$$

Then $O_n := \bigcup_{e \in E_n} \mathcal{C}_e$ is an open set in $|G|$ because it is a union of open sets. We show that $\Psi = \bigcap_{n \in \mathbb{N}} O_n$. Clearly, $\Psi \subseteq \bigcap_{n \in \mathbb{N}} O_n$. For the converse inclusion, let $\omega \in \bigcap_{n \in \mathbb{N}} O_n$. We show that ω does not live in any part of $(T, \mathcal{V})$ and therefore lies in the boundary of $(T, \mathcal{V})$. Indeed, if $\omega \in \partial V_t$ for $t \in T^n$, then ω is not contained in O_{n+1}, a contradiction.

(5) $\Rightarrow$ (1): Let $\Psi \subseteq \Omega(G)$ be a G_δ set in $|G|$. Hence $G \cup \Xi$ where $\Xi := \Omega(G) \setminus \Psi$ is an F_σ set in $|G|$ and by Lemma 6.2.5, also $V(G) \cup \Xi$ is an F_σ set in $|G|$. This means that $V(G) \cup \Xi = \bigcup_{n \in \mathbb{N}} X_n$ is a countable union of sets X_n which are closed in $|G|$, we may assume that $X_0 \subseteq X_1 \subseteq \cdots$. By applying Theorem 6.6.2 (with $\Delta_n = \varnothing$) there is an upwards connected tree-decomposition of finite adhesion into connected parts that homeomorphically displays $\Psi = \Omega(G) \setminus \bigcup_{n \in \mathbb{N}} X_n$. □

Corollary 6.7.2. *Displayable sets of ends are completely metrizable.*

Proof. The implication (2) $\Rightarrow$ (4) in Theorem 6.7.1 says that for every displayable set of ends $\Psi \subseteq \Omega(G)$ in a graph G we have that $|G|_\Psi$ is completely metrizable. Since $\Psi \subseteq |G|_\Psi$ is closed, and closed subspaces of completely metrizable spaces are again completely metrizable, it follows that Ψ is completely metrizable. □

Corollary 6.7.3. *Let G be a graph with a displayable set of ends $\Psi \subseteq \Omega(G)$ and let Φ be a subset of Ψ. Then Φ is (homeomorphically) displayable if and only if Φ is a G_δ set in Ψ.*

Proof. Immediate from (2) $\Leftrightarrow$ (5) in Theorem 6.7.1 and transitivity of the G_δ-property. □

Corollary 6.7.4. *Let G be a graph with a normal spanning tree. Then a subset $\Phi \subseteq \Omega(G)$ is (homeomorphically) displayable if and only if Φ is a G_δ set in $\Omega(G)$.*

Proof. Follows from (6) $\Rightarrow$ (5) in Theorem 6.4.1 together with the previous corollary for $\Psi = \Omega(G)$. □

6.8 Tree-decompositions distributing sets of ends

In this section we characterise which subsets of ends of a graph can be distributed by a tree-decomposition of finite adhesion. Recall that a topological space $X \subseteq Z$ has a *σ-discrete expansion in Z* if it can be written as a disjoint union $X = \bigsqcup_{n \in \mathbb{N}} X_n$ such that all X_n are discrete and all $Y_n := \bigcup_{i \leq n} X_i$ are closed in Z.

Theorem 6.8.1. *Let G be a connected graph and $\Xi \subseteq \Omega(G)$ a subset of ends of G. Then the following are equivalent:*

(i) There is a tree-decomposition of finite adhesion distributing Ξ.

(ii) $V(G)$ is a countable union of slender vertex sets U_n such that $\bigcup_{n \in \mathbb{N}} \partial U_n = \Xi$.

(iii) $V(G) \cup \Xi$ has a σ-discrete expansion in $|G|$.

(iv) There is an upwards connected tree-decomposition of finite adhesion with connected parts realising $(\Xi, \Xi^\complement)$.

Proof. We will show a cyclic chain of implications. For $(i) \Rightarrow (ii)$, suppose we have a tree-decomposition $(T, \mathcal{V})$ with root r of finite adhesion that distributes Ξ.

We define

$$U_n = \bigcup_{t \in T^{\leq n}} V_t.$$

By property (T1) of a tree-decomposition, it is clear that $V(G) \subseteq \bigcup_{n\in\mathbb{N}} U_n$. Since Ξ is the interior of $(T, \mathcal{V})$, we also have $\Xi = \bigcup_{n\in\mathbb{N}} \partial U_n$ as desired.

Furthermore, each U_n is slender: Clearly, all vertices are isolated in $|G|$. Additionally, $\partial U_n \setminus \partial U_{n-1}$ consists of at most one end for each part V_t for $t \in T^n$ and hence all ends in $\partial U_n \setminus \partial U_{n-1}$ are isolated points of U_n. Therefore, each $\overline{U_n}$ has Cantor-Bendixson rank at most $n+1$ by induction.

For $(ii) \Rightarrow (iii)$, suppose $V(G)$ is a countable union of slender vertex sets U_n such that $\bigcup_{n\in\mathbb{N}} \partial U_n = \Xi$. Without loss of generality, the sequence of the U_n is increasing. Write $X_n = \overline{U_n}$ and let $Y_0 = X_0$ and $Y_{n+1} = X_{n+1} \setminus X_n$. By assumption, each Y_n has finite Cantor-Bendixson rank say k_n. Recall that $Y_n^{(0)} := Y_n$ and $Y_n^{(i+1)}$ denotes the derived space of $Y_n^{(i)}$ for all $i \in \mathbb{N}$. Since Y_n has rank k_n, we have $Y_n^{(k_n)} = \varnothing$. Let $Z_{n,i} := Y_n^{(i)} \setminus Y_0^{(i+1)}$ be the subset of Y_n consisting of all elements that get deleted when forming $Y_n^{(i+1)}$ for $0 \leq i \leq k_n - 1$. We claim that

$$Z_{0,k_0-1}, Z_{0,k_0-2}, \ldots, Z_{0,0}, Z_{1,k_1-1}, Z_{1,k_1-2}, \ldots, Z_{1,0}, Z_{2,k_2-1}, Z_{2,k_2-2}, \ldots$$

is the desired σ-discrete expansion of $V(G) \cup \Xi$.

First of all, since $V(G) \cup \Xi = \bigcup_{n\in\mathbb{N}} Y_n$ and this union is disjoint, the above sequence has union $V(G) \cup \Xi$. By the definition of rank, it is also clear that all sets in the sequence are discrete. It remains to show that the union over finite initial segments is closed. Clearly, each such union is of the form

$$Y = X_n \cup Z_{n+1,k_{n+1}-1} \cup \cdots \cup Z_{n+1,i} \subseteq X_{n+1}$$

for some $i < k_{n+1}$, and this set is closed in $|G|$ as X_{n+1} is closed in $|G|$ and Y is closed in X_{n+1} by the definition of the Cantor-Bendixson rank.

For $(iii) \Rightarrow (iv)$, let $(X'_n)_{n\in\mathbb{N}}$ be a σ-discrete expansion for $V(G) \cup \Xi$. Then we apply Theorem 6.6.2 for the closed sets $X_n := \bigcup_{i \leq n} X'_i$ and the discrete sets $\Delta_n := X'_n \cap \Omega(G)$ to obtain an upwards connected tree-decomposition of G of finite adhesion into connected parts displaying $\Xi^{\complement}$ such that all ends from $\Xi = \bigcup_{n\in\mathbb{N}} \Delta_n$, and hence all ends from the interior of $\mathcal{T}$ live in pairwise distinct parts. In other words, this tree-decomposition realises $(\Xi^{\complement}, \Xi)$.

Next, it is clear that (iv) implies (i), which completes the proof. □

We have now all results in place to prove our main result Theorem 6.1.1 from this chapter; the following theorem contains even more equivalent properties:

Theorem 6.8.2. *The following are equivalent for any connected graph G with at least one end:*

(1) There is an upwards connected tree-decomposition of finite adhesion that represents $\Omega(G)$ such that all parts induce connected subgraphs.

(2) There is a tree-decomposition of finite adhesion that represents all ends in $\Omega(G)$.

(3) There is a tree-decomposition of finite adhesion that distinguishes all ends in $\Omega(G)$.

(4) There is a tree-decomposition of finite adhesion into $\leq$1-ended parts.

(5) Some subset $\Xi \subseteq \Omega(G)$ of ends can be distributed.

(6) $V(G)$ is a countable union of slender sets.

Proof. The implications (1) ⇒ (2) ⇒ (3) ⇒ (4) ⇒ (5) are trivial. The implication (5) ⇒ (6) follows from $(i) \Rightarrow (ii)$ in Theorem 6.8.1. Finally, for (6) ⇒ (1) note that due to $(ii) \Rightarrow (iv)$ in Theorem 6.8.1, we immediately get from (6) that there is an upwards connected tree-decomposition of finite adhesion into connected parts that realises $(\Xi, \Xi^{\complement})$. But then it follows from the subsequent Lemma 6.8.3 that there also is such a tree-decomposition $\mathcal{T}'$ that represents some partition (Ξ', Ψ') of $\Omega(G)$ with $\Xi \subseteq \Xi'$, and so $\mathcal{T}'$ represents all ends in $\Omega(G)$ as desired. □

Lemma 6.8.3. *If a connected graph G with at least one end admits a tree-decomposition $\mathcal{T}$ of finite adhesion that realises some partition (Ξ, Ψ) of $\Omega(G)$, then there also is such a tree-decomposition $\mathcal{T}'$ that represents some partition (Ξ', Ψ') of $\Omega(G)$ with $\Xi \subseteq \Xi'$.*

Moreover, whenever $\mathcal{T}$ has connected parts or is upwards connected, we can obtain the same for $\mathcal{T}'$.

Proof. Suppose we are given a tree-decomposition $(T, \mathcal{V})$ of finite adhesion realising some partition (Ξ, Ψ) of $\Omega(G)$. We will perform two rounds of contractions on T to make sure that we represent some partition (Ξ', Ψ') of $\Omega(G)$ with $\Xi \subseteq \Xi'$.

First, pick a maximal family $\mathcal{R}$ of disjoint rays in T such that no end of G lives in a part corresponding to one of the nodes of a ray in $\mathcal{R}$. Then consider a new tree-decomposition $(\dot{T}, \dot{\mathcal{V}})$ where $\dot{T}$ is obtained from T by contracting each ray in $\mathcal{R}$. For every $R \in \mathcal{R}$ we define a corresponding part $\dot{V}_R = \bigcup_{t \in R} V_t$. Since the set of separators of $(\dot{T}, \dot{\mathcal{V}})$ is a subset of the set of separators of $(T, \mathcal{V})$, it follows that also $(\dot{T}, \dot{\mathcal{V}})$ has finite adhesion. And since by assumption on $(T, \mathcal{V})$ there corresponds precisely one end of G to any ray $R \in \mathcal{R}$, it follows that $(\dot{T}, \dot{\mathcal{V}})$ realises (Ξ', Ψ') where Ξ' is the union of Ξ together with all ends of G that correspond to a ray in $\mathcal{R}$, and Ψ' is its complement.

Next, note that by maximality of $\mathcal{R}$, every ray of $\dot{T}$ contains infinitely many nodes whose corresponding parts in $\dot{\mathcal{V}}$ contain an end of G. Therefore, if we pick any partition $\mathcal{P}$ of $V(\dot{T})$ into subtrees such that each subtree P contains a unique node for which there is an end ω_P of G living in the corresponding part of $\dot{\mathcal{V}}$, then all $P \in \mathcal{P}$ are necessarily rayless.

Now consider a new tree-decomposition $(T', \mathcal{V}')$ where T' is obtained from $\dot{T}$ by contracting each subtree in $\mathcal{P}$. Naturally, $V(T') = \mathcal{P}$, and for each $P \in V(T')$ we define $V'_P = \bigcup_{t \in P} \dot{V}_t$. Since $\mathcal{T}'$ arises from $\mathcal{T}$ by contracting subtrees, it is clear that $\mathcal{T}'$ has finite adhesion, connected parts, or is upwards connected if the same is true for $\mathcal{T}$. Lastly, $(T', \mathcal{V}')$ now represents the partition (Ξ', Ψ'), as in each part V'_P there lives precisely the single end ω_P from Ξ', and since all P were rayless and $(\dot{T}, \dot{\mathcal{V}})$ displays Ψ', also $(T', \mathcal{V}')$ displays Ψ'. □

Corollary 6.8.4. *If a connected graph G with at least one end admits a rayless tree-decomposition $\mathcal{T}$ of finite adhesion that distributes $\Omega(G)$, then there also is such a tree-decomposition that bijectively distributes $\Omega(G)$. Moreover, whenever $\mathcal{T}$ has connected parts or is upwards connected, we can obtain the same for $\mathcal{T}'$.*

6.9 Tree-decompositions distributing all ends

In the previous section we stated a topological characterisation for the sets of ends that can be distributed. If we are interested in distributing all ends of G, we can obtain a combinatorial characterisation in terms of the underlying graph.

The following is a convenient description of the Cantor-Bendixson rank of the space $V \cup \Omega(G) \subseteq |G|$ due to Jung [46, §3]: The *rank* $r(x)$ of a vertex or an end x in a graph $G = (V, E)$ is defined as follows: all vertices have rank 0. An end ω has rank 1, if there is a finite set $S \subseteq V$, such that $\hat{C}(S, \omega)$ contains no other end. For an ordinal α, we say an end ω has rank α, if it has not already been assigned a smaller rank and if there is a finite set $S \subseteq V$ such that all ends in $\hat{C}(S, \omega)$ have been assigned a rank, and all these ranks are strictly smaller than α.

For a graph G in which every end has a rank (i.e. for graphs where $V \cup \Omega(G)$ is scattered), we define the *end-rank* $r(G)$ as the supremum of the ranks of all points in $V \cup \Omega(G)$.[4]

Theorem 6.9.1. *The following are equivalent for any connected graph G:*

(i) There is an upwards connected rayless tree-decomposition of finite adhesion with connected parts distributing $\Omega(G)$.

(ii) There is a tree-decomposition of finite adhesion distributing $\Omega(G)$.

(iii) $V \cup \Omega(G)$ has a σ-discrete expansion.

(iv) G contains no end-faithful subdivision of the full binary tree T_2.

(v) Every end of G has a rank, i.e. $\Omega(G)$ is scattered.

Moreover, if $\Omega(G) \neq \varnothing$, we may add

(vi) There is an upwards connected rayless tree-decomposition of finite adhesion with connected parts bijectively distributing $\Omega(G)$.

[4]We remark that in this formulation, $r(G)$ and the Cantor-Bendixson rank of $V \cup \Omega(G)$ may differ by ± 1.

Proof. $(i) \Leftrightarrow (ii) \Leftrightarrow (iii)$ is a special case of Theorem 6.8.1.

For the implication $(iii) \Rightarrow (iv)$ note that any subspace of $V \cup \Omega(G)$ inherits the property of having a σ-discrete expansion. However, the end space of a binary tree does not have a σ-discrete expansion: Indeed, any discrete set in a compact metric space is just countable; but the end space of a binary tree is uncountable, so not a countable union of countable sets.

The equivalence $(iv) \Leftrightarrow (v)$ is the content of Jung's [46, Satz 4].

We prove $(v) \Rightarrow (iii)$ by transfinite induction on the end-rank α of G. In the base case $r(G) = 0$, i.e. when $\Omega(G) = \varnothing$, we may take the trivial expansion consisting just of the vertex set.

Now let $\alpha > 0$, and suppose that all graphs of rank $< \alpha$ admit a σ-discrete expansion. First, let $\Phi \subseteq \Omega(G)$ consist of all ends of rank α. Clearly, Φ is a closed discrete subset of $\Omega(G)$. By Corollary 6.5.5, there is a connected envelope U for Φ, i.e. U is a connected set of vertices in G of finite adhesion such that $\partial U = \Phi$. Write $\mathcal{P}$ for the collection of components of $G - U$, and note that for each $P \in \mathcal{P}$, all ends living in P have rank $< \alpha$.

Now for each component $P \in \mathcal{P}$ individually, consider a collection $\mathcal{C}_P = \{C_P(S_\omega, \omega) : \omega \in \Omega(P)\}$ such that each set $C_P(S_\omega, \omega)$ witnesses the rank of ω inside the graph P. By Theorem 6.2.3, there is a rayless normal tree N_P in P such that every component D of $P - N_P$ is included in an element of $\mathcal{C}_P$ and hence satisfies $r(D) < \alpha$. Note that $U' = U \cup \bigcup_{P \in \mathcal{P}} N_P$ also is an envelope for Φ, but now, writing $\mathcal{P}'$ for the collection of components of $G - U'$, we have $r(D) < \alpha$ for every $D \in \mathcal{P}'$. By induction assumption, each $D \in \mathcal{P}'$ admits a σ-discrete expansion

$$V(D) \cup \Omega(D) = \bigcup_{n \geq 1} X_{D,n}.$$

Then $X_0 := \overline{U'} = U' \cup \Phi$ together with

$$X_n := \bigcup_{D \in \mathcal{P}'} X_{D,n}$$

for $n \geq 1$ gives the desired σ-discrete expansion of $V \cup \Omega(G)$. Indeed, to see that $X_0 \cup X_1 \cup \cdots \cup X_n$ is closed for every $n \in \mathbb{N}$, note that every end ω of G outside of this set lives in some component D for $D \in \mathcal{P}'$. Let $S \subseteq V(D)$ be finite such that $\hat{C}_D(S, \omega)$ is a basic open set inside $V(D) \cup \Omega(D)$ separating ω from the closed set $X_{D,1} \cup \cdots \cup X_{D,n}$. But then $\hat{C}_G(S \cup N(D), \omega)$ is a basic open neighbourhood of ω in $V \cup \Omega(G)$ witnessing that ω does not belong to the closure of $X_0 \cup X_1 \cup \cdots \cup X_n$. This completes the induction step and the proof of $(v) \Rightarrow (iii)$.

Finally, the moreover part $(i) \Leftrightarrow (vi)$ is immediate from Corollary 6.8.4. □

Using different methods, Polat showed that $\Omega(G)$ has a σ-discrete expansion if and only if every end of G has a rank [68, Theorem 8.11].

6.10 Applications

6.10.1 Tree-decompositions displaying special subsets of ends

Through our main characterisation, we can now give a short proof of the main result from Carmesin's [18].

Theorem 6.10.1. *Every connected graph G has a tree-decomposition of finite adhesion with connected parts that displays precisely the undominated ends of G.*

Proof. Let Ξ be the set of all ends of G which are dominated. By Theorem 6.7.1 it suffices to show that $\Omega(G) \setminus \Xi$ is a G_δ set in $|G|$, and by Lemma 6.2.5 it is equivalent to show that $V(G) \cup \Xi$ is an F_σ set in $|G|$. Choose an arbitrary vertex $u \in V(G)$ and for all $n \in \mathbb{N}$ write X_n for the set of all vertices of G with distance at most n to u. We show that $V(G) \cup \Xi = \bigcup_{n\in\mathbb{N}} \overline{X_n}$. We have $V(G) = \bigcup_{n\in\mathbb{N}} X_n$ because G is connected. It is left to show that the ends in $\bigcup_{n\in\mathbb{N}} \overline{X_n}$ are precisely the dominated ends of G.

Consider any end $\omega \in \Omega(G)$ and let R be an ω-ray in G. First, suppose that ω is dominated and let v be the centre of an infinite subdivided star S with leaves in R. Furthermore, suppose that $v \in X_n$. Then $S - v$ is a comb attached to $N(v) \subseteq X_{n+1}$ and therefore ω is contained in $\overline{X_{n+1}}$.

Now assume for a contradiction that some $\overline{X_n}$ contains an undominated end ω, and choose n minimal with that property. Then there is a comb C attached to X_n with spine $R \in \omega$. By minimality of n, there is an infinite set $\mathcal{T}$ of teeth of C which lie in $X_n \setminus X_{n-1}$. The neighbourhood of $\mathcal{T}$ in X_{n-1} is finite, again by minimality of n. Since every vertex in $X_n \setminus X_{n-1}$ has a neighbour in X_{n-1}, there is vertex $v \in X_{n-1}$ with infinitely many neighbours in $\mathcal{T}$. Hence ω is dominated by v, a contradiction. □

The following generalises a corresponding result from [17, Theorem 2].

Theorem 6.10.2. *For every infinite set of vertices U in a connected graph G, there is a tree-decomposition of G of finite adhesion that displays precisely the undominated ends of ∂U.*

Proof. Without loss of generality, we may assume that U has finite adhesion (Theorem 6.5.1). Consider the contraction minor $H \preceq G$ obtained from G by contracting each component C of $G - U$ to a single vertex v_C (of finite degree).

Claim 6.10.3. *The inclusion $U \hookrightarrow H$ induces a bijection $\partial U \to \Omega(H)$ that preserves the property of being dominated.*

This claim is proven just like Lemma 6.5.4.

Claim 6.10.4. *The contractions resulting in H induce a natural continuous surjection $f\colon |G| \to |H|$.*

To see that f is continuous, consider some end $\omega \in |G|$. If $\omega \notin \partial U$, then $f(\omega) = v_C$ for some component C, and f is continuous at ω. If $\omega \in \partial U$, then $f(\omega) = \omega' \in \Omega(H)$ by

Claim 6.10.3. Let $C_H(X', \omega')$ be an arbitrary basic open neighbourhood around ω' in H. Let $X \subseteq U$ be the finite set of vertices where we replace every vertex of the form v_C in X' by $N(C)$. It remains to verify that

$$f[C_G(X, \omega)] \subseteq C_H(X', \omega').$$

But this is clear: for every $v \in C_G(X, \omega)$, any $v - \omega$-ray R avoiding X is mapped to a locally finite connected subgraph in H avoiding X' which includes an $f(v) - \omega'$-ray R'.

Now we apply Theorem 6.10.1 inside H to see that there is a tree-decomposition of finite adhesion displaying the undominated ends Ψ of H. Hence Ψ is G_δ in $|H|$ by Theorem 6.7.1, say $\Psi = \bigcap_{n \in \mathbb{N}} O_n$ with O_n open in $|H|$. But then by Claim 6.10.4,

$$f^{-1}(\Psi) = f^{-1}(\bigcap_{n \in \mathbb{N}} O_n) = \bigcap_{n \in \mathbb{N}} f^{-1}(O_n)$$

is G_δ in $|G|$. Thus $f^{-1}(\Psi)$ can be displayed by a tree-decomposition of finite adhesion of G, again by Theorem 6.7.1. This completes the proof as $f^{-1}(\Psi)$ is the set of all undominated ends in ∂U by Claim 6.10.3. □

6.10.2 Counterexamples

Consider the full infinite binary tree T_2, and let $X \subseteq \Omega(T_2)$ be any set of ends. A *binary tree with tops* X is the graph with vertex set $T_2 \sqcup X$, all edges of T_2, and such that the neighbourhood of $x \in X$ consists of infinitely many nodes on its corresponding normal ray in T_2.

We reobtain Carmesin's observation that a T_2 with uncountably many tops does not admit a tree-decomposition of finite adhesion displaying all its ends, but now with significantly shorter proof.

Example 6.10.5. No binary tree with uncountably many tops admits a tree-decomposition of finite adhesion displaying all its ends.

Proof. These graphs do not have normal spanning trees by [29, Proposition 3.3], and so the result follows from Theorem 6.4.1. □

We can prove the following stronger result by Carmesin [18, p.7] with only a little more work.

Example 6.10.6. No binary tree with uncountably many tops admits a tree-decomposition of finite adhesion distinguishing all its ends.

Proof. Let G be a binary tree with uncountably many tops. Suppose for a contradiction that $V(G)$ is a countable union of slender sets. Then one of the slender sets U contains uncountably many of the tops. Write $\mathcal{R}$ for the set of all normal rays of T_2 which have a corresponding top in U. We call a vertex v of T_2 *good*, if it lies in uncountably many rays from $\mathcal{R}$. It is clear that the root of T_2 is good. We now show that for each good vertex v, there are two incomparable good vertices above v in the tree-order:

Suppose not for a contradiction. It is clear that at least one upper neighbour in T_2 of each good vertex is good. This implies that there is a ray R of good vertices above v. Since per assumption all good vertices above v are comparable, no other vertex above v outside the ray R is good. But this ray has only countable many neighbours in T_2. As no such neighbour above v is good, every neighbour of R above v lies on only countably many rays from $\mathcal{R}$. But then also v lies on only countably many rays from $\mathcal{R}$, which is a contradiction since v is good.

From this claim follows that there is a subdivided binary tree inside G such that each branch vertex is good.

It follows that ∂U itself contains the end space of a subdivided binary tree. But the end space of a binary tree is not scattered, a contradiction. It follows from Theorem 6.8.2 that $\Omega(G)$ cannot be distinguished. □

We conclude this section with a new example of a graph G witnessing that the thin ends of G cannot always be displayed, that is based on topological considerations only (a different example is given by Carmesin in [18, Example 3.3]). More precisely, since displayable subsets of ends are always completely metrizable by Corollary 6.7.2, it suffices to construct a graph where the thin ends are not completely metrizable.

As a warm-up, consider the binary tree T, and call a normal ray of T *rational* if its corresponding $0-1$-sequence becomes eventually constant, and *irrational* otherwise. Let $\Sigma \subseteq \Omega(T_2)$ be the subspace of rational ends. By Sierpinski's characterisation [75], every countable metric space without isolated points – so in particular Σ – is homeomorphic to the rational numbers $\mathbb{Q}$. Thus, Σ is not completely metrizable, and hence not displayable.

We now modify T such that all irrational ends become thick, and all rational ends remain thin. A binary tree with *fat* tops Z is a graph with vertex set $T \sqcup Z$, all edges of T_2, and such that the neighbourhood of $z \in Z$ consists of infinitely many nodes on some normal ray R_z of T_2. Thus, the difference between a tree with tops and tree with fat tops is that a normal ray may now have more than one top vertex.

Example 6.10.7. There is a binary tree with uncountably many fat tops such that its thin ends cannot be displayed.

Construction. Starting from the binary tree T, let $\{R_i : i \in \mathbb{N}\}$ be an enumeration of the rational rays in T. We now add infinitely many top-vertices above each irrational ray, and connect them to their rays such that

(1) each top-vertex z dominates its corresponding irrational ray R_z, and

(2) for each rational ray R_i, at most i vertices on R_i have top-vertices as neighbours.

Once the construction has finished, it is clear that the resulting graph G is as desired. The end space $\Omega(G) = \Omega(T)$ remains unchanged. From (2) it is easy to see that every rational end $\omega_i \ni R_i$ has end-degree 1 in G (and hence is thin), since the corresponding ray R_i has a tail of vertices of degree 3 whose edges are cut edges. All irrational ends are dominated by their infinitely many top-vertices and thus become thick.

By Sierpinski's characterisation [75], the set of rational / thin ends is homeomorphic to the rational numbers $\mathbb{Q}$, so not completely metrizable, and hence not displayable by Corollary 6.7.2.

It remains to describe how to connect a top-vertex z to its irrational ray R_z. For each top-vertex z and every $j \in \mathbb{N}$, let r_z^j be the $\leq_T$-minimal vertex in $R_z \setminus (R_0 \cup \ldots \cup R_j)$. Now let the neighbours of z be exactly the vertices in $\{r_z^0, r_z^1, r_z^2, \ldots\}$. Since $r_z^0 \leq r_z^1 \leq r_z^2 \leq \cdots$ is cofinal in R_z, the top-vertex z dominates R_z, establishing property (1).

Next, consider the ith rational ray R_i. Again, for $j \leq i$ let r_i^j be the $\leq_T$-minimal vertex in $R_i \setminus (R_0 \cup \ldots \cup R_j)$ (if it exists). Then it is clear that $r_i^0, r_i^1, \ldots, r_i^{i-1}$ are the only vertices on R_i adjacent to top-vertices, giving (2). □

Appendix

Chapter 7

English summary

Chapter 2

Two graphs are of the same *topological type* if they are topological minors of each other. We show that for any infinite cardinal κ there are exactly κ^+ distinct topological types of trees of size κ.

Chapter 3

A class $\mathcal{G}$ of graphs has the *Erdős-Pósa property (EPP)* if there is a function $f : \mathbb{N} \to \mathbb{N}$ such that for every graph Γ and every $k \in \mathbb{N}$ there are either k disjoint graphs from $\mathcal{G}$ in Γ or there is a set $X \subseteq V(\Gamma)$ of size at most $f(k)$ such that $\Gamma - X$ contains no graph from $\mathcal{G}$. Moreover, $\mathcal{G}$ has the *κ-EPP*, where κ is an infinite cardinal, if for any graph Γ there are either κ disjoint graphs from $\mathcal{G}$ in Γ or there is a set X of vertices of Γ of size less than κ such that $\Gamma - X$ contains no graph from $\mathcal{G}$.

We show that if $\mathcal{G}$ consists of a single infinite graph that does not contain a path of length n for some $n \in \mathbb{N}$, then $\mathcal{G}$ has the EPP and the κ-EPP for all infinite cardinals κ. Furthermore, we show that the class of all subdivisions of any tree T has the κ-EPP for every uncountable cardinal κ, and if T is rayless, also the $\aleph_0$-EPP and the EPP. On the other hand, we also find for every infinite cardinal κ a graph that does not have the κ-EPP.

Chapters 4 and 5

We call a digraph H *ubiquitous* if every digraph D containing k disjoint copies of H for every $k \in \mathbb{N}$ also contains infinitely many disjoint copies of H. A *turn* of an oriented (double) ray is a vertex of in-degree 2 or out-degree 2. We prove that an oriented ray is ubiquitous if and only if it has a finite number of turns, and that an oriented double ray with at least one turn is ubiquitous if and only if it has an odd number of turns.

Chapter 6

The ends of a graph G interact in a natural way with a tree-decomposition $(T, \mathcal{V})$ of G of finite adhesion. As every edge e of T induces a finite order separation $\{A_e, B_e\}$ of G, each end of G has to choose one side of the separation and thus one component of $T - e$. By orienting the edges of T accordingly, we obtain an orientation of T for each end of G. Consider any set Ψ of ends of G. If for every end in Ψ the respective orientation of T points towards a (unique) end of T, and if this correspondence between Ψ and the ends of T is bijective, then we say that the tree-decomposition $(T, \mathcal{V})$ *displays* Ψ.

We find connections between sets of ends of G that can be displayed and their topological properties in the topological space $|G|$ formed by G together with its ends. In particular, we show that the following are equivalent:

- There is a tree-decomposition of finite adhesion displaying Ψ.
- The subspace of $|G|$ consisting of Ψ together with all vertices and edges of G is completely metrizable.
- Ψ is G_δ (i.e. a countable intersection of open sets) in $|G|$.

Chapter 8

Deutsche Zusammenfassung

Kapitel 2

Zwei Graphen sind vom selben topologischen Typ, wenn sie sich gegenseitig als topologische Minoren enthalten. Wir zeigen, dass es für jede unendliche Kardinalzahl κ genau κ^+ verschiedene topologische Typen von Bäumen der Größe κ gibt.

Kapitel 3

Eine Klasse $\mathcal{G}$ von Graphen hat die *Erdős-Pósa-Eigenschaft (EPE)*, wenn es eine Funktion $f : \mathbb{N} \to \mathbb{N}$ gibt, sodass es für jeden Graphen Γ und jedes $k \in \mathbb{N}$ entweder k disjunkte Graphen aus $\mathcal{G}$ in Γ gibt oder es eine Menge $X \subseteq V(\Gamma)$ der Größe höchstens $f(k)$ gibt, sodass $\Gamma - X$ keinen Graphen aus $\mathcal{G}$ enthält. Außerdem hat $\mathcal{G}$ die *κ-EPE*, wobei κ eine unendliche Kardinalzahl ist, wenn es für jeden Graphen Γ entweder κ disjunkte Graphen aus $\mathcal{G}$ in Γ gibt oder es eine Menge X von Ecken von Γ gibt, die kleiner als κ ist, sodass $\Gamma - X$ keinen Graphen aus $\mathcal{G}$ enthält.

Wir zeigen, dass $\mathcal{G}$ die EPE und die κ-EPE für alle unendlichen Kardinalzahlen κ hat, wenn $\mathcal{G}$ aus einem einzigen unendlichen Graphen besteht, der keinen Weg der Länge n für irgendein $n \in \mathbb{N}$ enthält. Zusätzlich zeigen wir, dass die Klasse aller Unterteilungen eines beliebigen Baums T die κ-EPE für jede überabzählbare Kardinalzahl κ hat, und wenn T strahlenlos ist, auch die $\aleph_0$-EPE und die EPE. Andererseits finden wir auch für jede unendliche Kardinalzahl κ einen Graphen, der nicht die κ-EPE hat.

Kapitel 4 und 5

Wir nennen einen gerichteten Graphen H *allgegenwärtig*, wenn jeder gerichtete Graph D, der k disjunkte Kopien von H für jedes $k \in \mathbb{N}$ enthält, auch unendlich viele disjunkte Kopien von H enthält. Ein *Richtungswechsel* eines orientierten Strahls oder Doppelstrahls ist eine Ecke von Eingangsgrad 2 oder Ausgangsgrad 2. Wir beweisen, dass ein orientierter Strahl genau dann allgegenwärtig ist, wenn er eine endliche Anzahl von Richtungswechseln hat, und dass ein orientierter Doppelstrahl mit mindestens einem Richtungswechsel genau dann allgegenwärtig ist, wenn er eine ungerade Anzahl von Richtungswechseln hat.

Kapitel 6

Die Enden von einem Graphen G interagieren auf natürliche Weise mit einer Baumzerlegung $(T, \mathcal{V})$ von G von endlicher Adhäsion. Da jede Kante e von T eine Teilung $\{A_e, B_e\}$ von G endlicher Ordnung induziert, muss jedes Ende von G eine Seite der Teilung und damit eine Komponente von $T - e$ wählen. Indem wir die Kanten von T entsprechend orientieren, erhalten wir eine Orientierung von T für jedes Ende von G. Betrachten wir eine beliebige Menge Ψ von Enden von G. Wenn für jedes Ende in Ψ die entsprechende Orientierung von T auf ein (eindeutiges) Ende von T zeigt, und wenn diese Beziehung zwischen Ψ und den Enden von T bijektiv ist, dann sagen wir, dass die Baumzerlegung $(T, \mathcal{V})$ die Menge Ψ *darstellt*.

Wir finden Zusammenhänge zwischen Mengen von darstellbaren Enden von G und ihren topologischen Eigenschaften im topologischen Raum $|G|$, der von G zusammen mit seinen Enden gebildet wird. Wir zeigen, dass die folgenden Aussagen äquivalent sind:

- Es gibt eine Baumzerlegung von endlicher Adhäsion, die Ψ darstellt.
- Der Unterraum von $|G|$, der aus Ψ zusammen mit allen Ecken und Kanten von G besteht, ist vollständig metrisierbar.
- Ψ ist G_δ (d.h. ein abzählbarer Schnitt von offenen Mengen) in $|G|$.

Chapter 9

Publications related to this dissertation

- Chapter 2 is based on [51].
- Chapter 3 is based on [50].
- Chapter 4 is based on [37].
- Chapter 5 is based on [36] and parts of [37].
- Chapter 6 is based on [47].
- Chapters 1, 7 and 8 are based on all of the papers cited above.

Chapter 10

Declaration of my contributions

- Chapter 2 is joint work with Max Pitz [51]. We conducted the research together and I drafted Sections 2.3 and 2.4.
- I created Chapter 3 on my own [50].
- Chapters 4 and 5 are joint work with Florian Gut and Florian Reich [36, 37]. We conducted the research together and wrote both chapters together.
- Chapter 6 is joint work with Marcel Koloschin and Max Pitz [47]. We conducted the research together and I drafted Section 6.6 and parts of Sections 6.2, 6.3, 6.7, and 6.10.

Acknowledgement

I would like to thank my supervisor Max Pitz for his help and support with all my research projects, for introducing me to infinite graph theory, and for teaching me mathematical writing. I also want to thank my co-workers and office mates Florian Gut, Florian Reich and Marcel Koloschin. We had a lot of fun working together or just chatting around the office. Special thanks to Florian Reich for countless inspiring mathematical discussions. Furthermore, I am grateful to Nathan Bowler for always listening to my ideas and giving me helpful feedback and impulses. Thanks also to everyone else in the Discrete Mathematics Research Group. We had a great time together and I hope to see some of you in the bouldering gym again! Finally, I would like to thank Nicola Lorenz, my sister Pia Krill and my parents Sibylle Schrenner-Krill and Jürgen Krill for all their love, care and support.

Bibliography

[1] S. Albrechtsen, R. W. Jacobs, P. Knappe, and M. Pitz, *Linked tree-decompositions into finite parts*, Preprint, arXiv:2405.06753 (2024).

[2] T. Andreae, *On disjoint configurations in infinite graphs*, J. Graph Theory **39** (2002), no. 4, 222–229.

[3] ———, *Classes of locally finite ubiquitous graphs*, J. Combin. Theory Ser. B **103** (2013), no. 2, 274–290.

[4] A. Atminas and V. Lozin, *Labelled induced subgraphs and well-quasi-ordering*, Order **32** (2015), no. 3, 313–328.

[5] J. Bang-Jensen and G. Gutin, *Digraphs: Theory, algorithms and applications*, Springer Science & Business Media, 2008.

[6] H. L. Bodlaender, *Treewidth: Algorithmic techniques and results*, Proc. 22nd MFCS, 1997, pp. 19–36.

[7] N. Bowler and J. Carmesin, *Infinite matroids and determinacy of games*, Preprint, arXiv:1301.5980 (2013).

[8] ———, *The ubiquity of psi-matroids*, Preprint, arXiv:1304.6973 (2013).

[9] N. Bowler, J. Carmesin, and J. Pott, *Edge-disjoint double rays in infinite graphs: A Halin type result*, J. Combin. Theory Ser. B **111** (2013), 1–16.

[10] N. Bowler, C. Elbracht, J. Erde, J. P. Gollin, K. Heuer, M. Pitz, and M. Teegen, *Topological ubiquity of trees*, J. Combin. Theory Ser. B **157** (2022), 70–95.

[11] ———, *Ubiquity of graphs with nowhere-linear end structure*, J. Graph Theory **103** (2023), no. 3, 564–598.

[12] ———, *Ubiquity of locally finite graphs with extensive tree-decompositions*, Combin. Theory **4** (2024), no. 2.

[13] R. Brignall, M. Engen, and V. Vatter, *A counterexample regarding labelled well-quasi-ordering*, Graphs Combin. **34** (2018), 1395–1409.

[14] J. Bruno, *A family of ω_1 topological types of locally finite trees*, Discrete Math. **340** (2017), no. 4, 794–795.

[15] J. Bruno and P. Szeptycki, *There are exactly ω_1 topological types of locally finite trees with countably many rays*, Fund. Math. **256** (2022), 243–259.

[16] C. Bürger and J. Kurkofka, *Duality theorems for stars and combs I: Arbitrary stars and combs*, J. Graph Theory **99** (2022), no. 4, 525–554.

[17] ———, *Duality theorems for stars and combs II: Dominating stars and dominated combs*, J. Graph Theory **99** (2022), no. 4, 555–572.

[18] J. Carmesin, *All graphs have tree-decompositions displaying their topological ends*, Combinatorica **39** (2019), 545–596.

[19] J. Carmesin, *Topological cycle matroids of infinite graphs*, European J. Combin. **60** (2017), 135–150.

[20] ———, *On Andreae's ubiquity conjecture*, J. Combin. Theory Ser. B **162** (2023), 68–70.

[21] J. Daligault, M. Rao, and S. Thomassé, *Well-quasi-order of relabel functions*, Order **14** (1990), 427–435.

[22] R. Diestel, *The end structure of a graph: recent results and open problems*, Discrete Math. **100** (1992), no. 1–3, 313–327.

[23] ———, *The depth-first search tree structure of $TK_{\aleph_0}$-free graphs*, J. Combin. Theory Ser. B **61** (1994), no. 2, 260–262.

[24] ———, *End spaces and spanning trees*, J. Combin. Theory Ser. B **96** (2006), no. 6, 846–854.

[25] ———, *Locally finite graphs with ends: a topological approach, I–III.*, Discrete Math. **311–312** (2010/11).

[26] ———, *Graph Theory*, 5th ed., Springer, 2015.

[27] ———, *A simple existence criterion for normal spanning trees*, Electron. J. Combin. **23** (2016), no. 2, P2.33.

[28] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, J. Combin. Theory Ser. B **87** (2003), no. 1, 197–206.

[29] R. Diestel and I. Leader, *Normal spanning trees, Aronszajn trees and excluded minors*, J. London Math. Soc. **63** (2001), no. 1, 16–32.

[30] R. Diestel and J. Pott, *Dual trees must share their ends*, J. Combin. Theory Ser. B **123** (2017), 32–53.

[31] G. Ding, *Subgraphs and well-quasi-ordering*, J. Graph Theory **16** (1992), no. 5, 489–502.

[32] B. Dushnik and E. W. Miller, *Partially ordered sets*, Amer. J. Math. **63** (1941), no. 3, 600–610.

[33] A. Elm and J. Kurkofka, *A tree-of-tangles theorem for infinite tangles*, Abh. Math. Sem. Univ. Hamburg **92** (2022), 139–178.

[34] R. Engelking, *General toplogy – revised and completed ed.*, Vol. 6, Heldermann Verlag Berlin, 1989.

[35] P. Erdős and L. Pósa, *On independent circuits contained in a graph*, Canad. J. Math. **17** (1965), 347–352.

[36] F. Gut, T. Krill, and F. Reich, *On the ubiquity of oriented double rays*, Preprint, arXiv:2310.09857 (2023).

[37] ———, *Ubiquity of oriented rays*, J. Graph Theory **107** (2024), no. 1, 200–211.

[38] R. Halin, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Math. Nachr. **30** (1965), 63–85.

[39] ———, *Graphen ohne unendliche Wege*, Math. Nachr. **31** (1966), 111–123.

[40] ———, *Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen*, Math. Nachr. **44** (1970), 119–127.

[41] ———, *Some path problems in graph theory*, Abh. Math. Sem. Univ. Hamburg **44** (1975), 175–186.

[42] ———, *Simplicial decompositions of infinite graphs*, Ann. Discrete Math. **3** (1978), 93–109.

[43] R. Halin, *The structure of rayless graphs*, Abh. Math. Sem. Univ. Hamburg **68** (1998), no. 1, 225–253.

[44] T. Jech, *Set theory, the third millennium edition*, Springer Monogr. Math., 2003.

[45] B. Jia, *Excluding long paths*, Preprint, arXiv:1503.08258 (2015).

[46] H. A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, Math. Nachr. **41** (1969), 1–22.

[47] M. Koloschin, T. Krill, and M. Pitz, *End spaces and tree-decompositions*, J. Combin. Theory Ser. B **161** (2023), 147–179.

[48] P. Komjáth, *A note on minors of uncountable graphs*, Math. Proc. Cambridge Philos. Soc. **117** (1995), no. 1, 7–9.

[49] N. Korpelainen and V. Lozin, *Two forbidden induced subgraphs and well-quasi-ordering*, Discrete Math. **311** (2011), 1813–1822.

[50] T. Krill, *The Erdős-Pósa property for infinite graphs*, Preprint, arXiv:2411.02561 (2024).

[51] T. Krill and M. Pitz, *The number of topological types of trees*, Combinatorica **44** (2024), 651–657.

[52] J. B. Kruskal, *Well-quasi-ordering, the tree theorem, and Vazsonyi's conjecture*, Trans. Amer. Math. Soc. **95** (1960), 210–225.

[53] J. Kurkofka and R. Melcher, *Countably determined ends and graphs*, J. Combin. Theory Ser. B **156** (2022), 31–56.

[54] J. Kurkofka and M. Pitz, *A representation theorem for end spaces of infinite graphs*, Preprint, arXiv:2111.12670 (2021).

[55] J. Kurkofka, *Ubiquity and the Farey graph*, European J. Combin. **95** (2021), 103326.

[56] J. Kurkofka, R. Melcher, and M. Pitz, *Approximating infinite graphs by normal trees*, J. Combin. Theory Ser. B **148** (2021), 173–183.

[57] D. Kühn, *On well-quasi-ordering infinite trees – Nash–Williams's theorem revisited*, Math. Proc. Cambridge Philos. Soc. **130** (2001), no. 3, 401–408.

[58] R. Laver, *Better-quasi-orderings and a class of trees*, Studies in Foundations and Combinatorics, Adv. in Math. Supplementary Studies **1** (1978), 31–48.

[59] L. Matthiesen, *There are uncountably many topological types of locally finite trees*, J. Combin. Theory Ser. B **96** (2006), no. 5, 758–760.

[60] C. St. J. A. Nash-Williams, *On well-quasi-ordering infinite trees*, Math. Proc. Cambridge Philos. Soc. **61** (1965), no. 3, 697–720.

[61] B. Oporowski, *A counterexample to Seymour's self-minor conjecture*, J. Graph Theory **14** (1990), no. 5, 521–524.

[62] M. Pitz, *Constructing tree-decompositions that display all topological ends*, Combinatorica **42** (2022), 763–769.

[63] M. Pitz, *A unified existence theorem for normal spanning trees*, J. Combin. Theory Ser. B **145** (2020), 466–469.

[64] ———, *Proof of Halin's normal spanning tree conjecture*, Isr. J. Math. **246** (2021), 353–370.

[65] ———, *Quickly proving Diestel's normal spanning tree criterion*, Electron. J. Combin. **28** (2021), no. 3, P3.59.

[66] ———, *A note on minor antichains of uncountable graphs*, J. Graph Theory **102** (2023), no. 3, 552–555.

[67] N. Polat, *Ends and multi-endings, I*, J. Combin. Theory Ser. B **67** (1996), 86–110.

[68] ———, *Ends and multi-endings, II*, J. Combin. Theory Ser. B **68** (1996), 56–86.

[69] M. Pouzet, *Un bel ordre d'abritement et ses rapports avec les bornes d'une multirelation*, C. R. Acad. Sci., Paris Sér. A–B **274** (1972), 1677–1680.

[70] R. Rado, *Partial well-ordering of sets of vectors*, Mathematika **1** (1954), no. 2, 89–95.

[71] J.-F. Raymond and D. M. Thilikos, *Recent techniques and results on the Erdős–Pósa property*, Discrete Appl. Math. **231** (2017), 25–43.

[72] N. Robertson and P. D Seymour, *Graph minors. IV. Tree-width and well-quasi-ordering*, J. Combin. Theory Ser. B **48** (1990), no. 2, 227–254.

[73] N. Robertson and P. D. Seymour, *Graph Minors. XX. Wagner's conjecture*, J. Combin. Theory Ser. B **92** (2004), no. 2, 325–357.

[74] R. Schmidt, *Ein Ordnungsbegriff für Graphen ohne unendliche Wege mit einer Anwendung auf n-fach zusammenhängende Graphen*, Arch. Math. (Basel) **40** (1983), no. 1, 283–288.

[75] W. Sierpiński, *Sur une propriété topologique des ensembles dénombrables denses en soi*, Fund. Math. **1** (1920), no. 1, 11–16.

[76] R. Thomas, *A counter-example to 'Wagner's conjecture' for infinite graphs*, Math. Proc. Cambridge Philos. Soc. **103** (1988), no. 1, 55–57.

[77] ———, *Well-quasi-ordering infinite graphs with forbidden finite planar minor*, Trans. Amer. Math. Soc. **312** (1989), no. 1, 279–313.

[78] C. Thomassen, *On the presence of disjoint subgraphs of a specified type*, J. Graph Theory **12** (1988), no. 1, 101–111.

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.