End space theory for directed and undirected graphs

Dissertation

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

vorgelegt
im Fachbereich Mathematik
von
Ruben Melcher

Hamburg
2021

Prüfungskommission:

Vorsitz:	Professor Armin Iske	(Universität Hamburg)
Erstgutachter und Betreuer:	Professor Reinhard Diestel	(Universität Hamburg)
Zweitgutachter:	Professor Nathan Bowler	(Universität Hamburg)
Drittgutachter (extern):	Associate Professor (Reader) Agelos Georgakopoulos	(University of Warwick)
Mitglied und Schriftführer:	PD Dr. Ralf Holtkamp	(Universität Hamburg)

Datum der Disputation: 13. Juli 2021

To Mareike

Contents

1. Introduction

This dissertation aims to make a step towards a deeper understanding of ends of infinite graphs. This first chapter serves the purpose of giving a concise introduction and an overview of the results of this dissertation.

Ends of graphs are one of the most important concepts in modern infinite graph theory. They can be thought of as points at infinity to which the rays of an infinite graph converge. Formally, an *end* of an (undirected) graph G is an equivalence class of its rays where two rays are equivalent if for every finite vertex set $X \subseteq V(G)$ they have a tail in the same component of $G - X$. For example, infinite complete graphs or grids have one end while the binary tree has continuum many ends, one for every rooted ray [25].

One typical use of ends is to extend finite to infinite theorems of graphs, in that graphs together with their ends naturally form topological spaces; in these spaces, topological arcs and circles take the role of paths and cycles, respectively. However, their use goes far beyond this.

Only recently the infinite cycle space theory which ends make possible uncovered a deep connection between combinatorics and a long-standing problem in topology, see [36]. Within combinatorics, ends show up in problems which, at first sight, have nothing to do with ends. For example, the reconstruction conjecture for locally finite graphs essentially depends on the number of ends of the graph, see [5] for a detailed overview. Moreover, ends arise naturally in other branches of mathematics such as in group theory via Cayley graphs or in topology as ends of CW complexes.

Now, we have to be modest; in this thesis, we will not touch on every use of ends in mathematics, although the methods that we use vary from general topology to set theory.

This dissertation offers three parts; each part is divided into multiple chapters. In the first part, we consider those topological spaces which are formed by the ends of a graph. We will study two fundamental topological properties of these spaces and our results will have consequences of both combinatorial and topological nature. In the second part, we will focus on how the rays in an end can link up in the graph. Here we will exploit set theory in order to solve an old conjecture of Halin. Finally, in the third part, we propose a new notion of ends of digraphs and develop a corresponding theory of their end spaces. There have been a few attempts on that before but not with very encouraging results. In this last part, we extend to directed graphs a number of techniques that are fundamental in the study of ends of graphs. Furthermore, we introduce a topological space $|D|$ which is naturally formed by an infinite directed graph D and its ends. We show that a number of well-known theorems about finite directed graphs extend to this space while they do not generalise verbatim to infinite directed graphs.

Let us proceed with an overview of the chapters of this thesis. At this point we want to remark that each chapter starts with its own more comprehensive introduction.

1.1. Part I: End spaces

1.1.1. Chapter 2: Approximating infinite graphs by normal trees

A rooted tree T contained in a graph G is *normal* in G if the endvertices of every T-path in G are comparable in the tree-order of T. (In finite graphs, normal spanning trees are their depth-first search trees; see [25] for precise definitions.) Normal spanning trees are perhaps the most useful structural tool in infinite graph theory. Their importance arises from the fact that they capture the separation properties of the graph they span, and so in many situations it suffices to deal with the much simpler tree structure instead of the whole graph. For example, the end space of G coincides, even topologically, with the end space of any normal spanning tree of G. However, not every connected graph has a normal spanning tree, and the structure of graphs without normal spanning trees is still not completely understood [6, 30].

In order to harness and transfer the power of normal spanning trees to arbitrary connected graphs G, one might try to find an 'approximate normal spanning tree': a normal tree in G which spans the graph up to some arbitrarily small given error term. We show that every connected graph can be approximated by a normal tree, up to some arbitrarily small error phrased in terms of neighbourhoods around its ends, see Theorem 2.1. The existence of such approximate normal trees has consequences of both combinatorial and topological nature.

On the combinatorial side, we show that a graph has a normal spanning tree as soon as it has normal spanning trees locally at each end; i.e., the only obstruction for a graph to having a normal spanning tree is an end for which none of its neighbourhoods has a normal spanning tree, see Theorem 2.4.3.

On the topological side, we show that the end space $\Omega(G)$, as well as the spaces $|G| = G \cup \Omega(G)$ naturally associated with a graph G, are always paracompact. This gives unified and short proofs for a number of results by Diestel, Sprüssel and Polat, see Section 2.4, and answers an open question about metrizability of end spaces by Polat, see Theorem 2.4.1.

1.1.2. Chapter 3: Countably determined ends and graphs

Halin [45] defined the ends of an infinite graph 'from below' as equivalence classes of rays in the graph, where two rays are equivalent if no finite set of vertices separates them. As a complementary description of Halin's ends, Diestel and Kühn [28] introduced the notion of directions of infinite graphs. These are defined 'from above': A *direction* of a graph G is a map f, with domain the collection $\mathcal{X} = \mathcal{X}(G)$ of all finite vertex sets of G, that assigns to every $X \in \mathcal{X}$ a component $f(X)$ of $G - X$ such that $f(X) \supseteq f(X')$ whenever $X \subseteq X'$.

Every end ω of G defines a direction f_ω of G by letting $f_\omega(X)$ be the component of $G - X$ that contains a subray of every ray in ω. Diestel and Kühn showed that the natural map $\omega \mapsto f_\omega$ is in fact a bijection between the ends of G and its directions. This correspondence is now well-known and has become a standard tool in the study of infinite graphs. See [12, 21, 23, 54–56] for examples.

Although every direction is induced by a ray, there exist directions of graphs that are not uniquely determined by any countable subset of their choices. We characterise these directions and their countably determined counterparts in terms of star-like substructures or rays of the graph, see Theorem 3.1 and Theorem 3.2.

Curiously, there exist graphs whose directions are all countably determined but which cannot be distinguished all at once by countably many choices.

We structurally characterise the graphs whose directions can be distinguished all at once by countably many choices, and we structurally characterise the graphs which admit no such countably many choices. Our characterisations are phrased in terms of normal trees and tree-decompositions, see Theorem 3.3 and Theorem 3.4.

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

1.2. Part II: The degree of an end

1.2.1. Chapter 4: A strengthening of Halin's grid theorem

The *degree* of an end is the maximum cardinality of a collection of pairwise disjoint rays in that end, see Halin [44]. One of the cornerstones of infinite graph theory, *Halin's grid theorem* [44], says that every graph with an end of infinite degree contains a subdivision of the hexagonal half-grid, see Figure 4.1.1, whose rays belong to that end. We show that for every infinite collection $\mathcal{R}$ of disjoint equivalent rays in a graph G there is a subdivision of the hexagonal half-grid in G such that all its vertical rays belong to $\mathcal{R}$. This result strengthens Halin's grid theorem by giving control over which specific set of rays is used, while its proof is significantly shorter.

1.2.2. Chapter 5: Halin's end degree conjecture

Halin conjectured that the end degree can be characterised in terms of certain typical ray configurations, which would generalise his famous grid theorem. In particular, every end of regular uncountable degree κ would contain a *star of rays*, i.e. a configuration consisting of a central ray R and κ neighbouring rays $(R_i\colon i<\kappa)$ all disjoint from each other and each R_i sending a family of infinitely many disjoint paths to R so that paths from distinct families only meet in R.

We show that Halin's conjecture fails for end degree $\aleph_1$, holds for $\aleph_2, \aleph_3, \ldots, \aleph_\omega$, fails for $\aleph_{\omega+1}$, and is undecidable (in ZFC) for the next $\aleph_{\omega+n}$ with $n \in \mathbb{N}$, $n \geq 2$. Further results include a complete solution for all cardinals under GCH, complemented by a number of consistency results.

1.3. Part III: Ends of digraphs

Graphs together with their ends naturally form a topological space $|G|$. Many well-known theorems of finite graph theory extend to this space, while they do not generalise verbatim to infinite graphs. Examples include Nash-William's tree-packing theorem [22], Fleischner's Hamiltonicity theorem [37], and Whitney's planarity criterion [8]. In the formulation of these theorems, topological arcs and circles take the role of paths and cycles, respectively.

For directed graphs, a similarly useful notion and theory of ends has never been found. There have been a few attempts, most notably by Zuther [70], but not with very encouraging results. In this part, we propose a new notion of ends of digraphs and develop a corresponding theory of their end spaces. As for undirected graphs, the *ends* of a digraph are points at infinity to which its rays converge. Unlike for undirected graphs, some ends are joined by *limit edges* which will turn out to be crucial for what is to come, for example obtaining the end space of a directed graph as a natural (inverse) limit of its finite contraction minors.

Let us now introduce a minimum of definitions in order to look at an example to convey some basic intuition. A *directed ray* is an infinite directed path that has a first vertex (but no last vertex). The directed subrays of a directed ray are its *tails*. We call a directed ray in a digraph D *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. We call two solid rays in a digraph D *equivalent* if for every finite vertex set $X \subseteq V(D)$ they have a tail in the same strong component of $D - X$. The classes of this equivalence relation are the *ends* of D. Note that two solid rays R and R' in D represent the same end if and only if D contains infinitely many disjoint paths from R to R' and infinitely many disjoint paths from R' to R. For a finite vertex set $X \subseteq V(D)$ and an end ω we write $C(X, \omega)$ for the unique strong component of $D - X$ that contains a tail of every ray that represents ω; the end ω is then said to *live* in that strong component.

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

For example, the digraph D in Figure 1.3.1 has two ends joined by a limit edge. Both the upper ray R and the lower ray R' are solid in D because the vertex set of any tail of R or R' is strongly connected in D. Deleting finitely many vertices of D always results in precisely two infinite strong components (and finitely many finite strong components) spanned by the vertex sets of tails of R or R'. Here the component spanned by the vertex set of a tail of R' sends an edge to the one spanned by the vertex set of a tail of R.

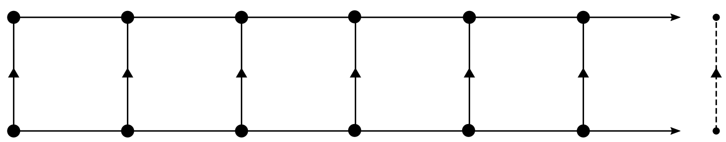

Figure 1.3.1.: A digraph with two ends (depicted as small dots) linked by a limit edge (depicted as a dashed line). Every undirected edge in the figure represents a pair of inversely directed edges.

1.3.1. Chapter 6: Ends of digraphs I: basic theory

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

As our main result in the first chapter of this part, we show that the notion of *directions* of an undirected graph, a tangle-like description of its ends, extends to digraphs: there is a one-to-one correspondence between the 'directions' of a digraph and its ends and limit edges.

In the course of this we extend to digraphs a number of fundamental tools and techniques for the study of ends of graphs, such as the star-comb lemma and Schmidt's ranking of rayless graphs.

1.3.2. Chapter 7: Ends of digraphs II: the topological point of view

In the second chapter of this part, we introduce the topological space $|D|$ formed by a digraph D together with its ends and limit edges. We then characterise those digraphs that are compactified by this space. Furthermore, we show that if $|D|$ is compact, it is the inverse limit of finite contraction minors of D.

To illustrate the use of this, we extend to the space $|D|$ two statements about finite digraphs that do not generalise verbatim to infinite digraphs. The first statement is the characterisation of finite Eulerian digraphs by the condition that the in-degree of every vertex equals its out-degree. The second statement is the characterisation of strongly connected finite digraphs by the existence of a closed Hamilton walk.

1.3.3. Chapter 8: Ends of digraphs III: normal arborescences

In the third chapter of this part, we consider normal spanning trees, one of the most important structural tools in infinite graph theory. In finite graphs, normal spanning trees are precisely the depth-first search trees [25].

As a directed analogue of normal spanning trees, we introduce and study *normal spanning arborescences* of digraphs. These are generalisations of depth-first search trees to infinite digraphs, which promise to be as powerful for a structural analysis of digraphs as normal spanning trees are for graphs. We show that normal spanning arborescences are end-faithful: every end of the digraph is represented by exactly one ray in the normal spanning arborescence that starts from the root. We further show that this bijection extends to a homeomorphism between the end space of a digraph D, which may include limit edges between ends, and the end space of any normal arborescence with limit edges induced from D. Finally, we prove a Jung-type criterion for the existence of normal spanning arborescences.

1.3.4. Chapter 9: Hamiltonicity in infinite tournaments

In the first three chapters of this part, we introduced a notion of ends in digraphs for which the fundamental techniques of undirected end space theory naturally generalise to digraphs. In the second chapter, we showed that a digraph D together with its ends and limit edges naturally forms a topological space $|D|$. So the scene is set now to attempt, also for digraphs D, to extend finite to infinite theorems by letting the naturally oriented topological paths and circles in $|D|$ take the role of directed paths and cycles in D. The purpose of this chapter is to make a start on this programme, with two well-known Hamiltonicity theorems for digraphs.

Two folklore theorems in finite graph theory, due to Rédei [62] and Camion [16], respectively, say that every finite tournament has a Hamilton path, and every finite strongly connected tournament has a Hamilton cycle. In this chapter, we show that these results have natural analogues in the space $|D|$. We prove that for all countable tournaments D its compactification $|D|$ by its ends and limit edges contains a topological Hamilton path: a topological arc that contains every vertex. If D is strongly connected, then $|D|$ contains a topological Hamilton circle. We shall see that ends and limit edges are both crucial for such extensions to exist: there exists a countable tournament D whose compactification by just the ends of the underlying undirected graph contains no topological Hamilton path. Similarly, D has no topological Hamilton path in $|D|$ that avoids all its limit edges, and D has no spanning ray or double ray.

1.4. Preliminaries

For graph theoretic terms we follow the terminology of Diestel [25], and in particular [25, Chapter 8] for ends of graphs. For topological notions we follow the terminology of Engelking [33] and for set theory we generally follow the textbook by Jech [49].

In Part I and Part II, we usually consider simple undirected graphs without multiple edges or loops, although our proofs (with obvious changes) also work for multigraphs with all parallel edges of finite multiplicity. For a graph G we write $V(G)$ and $E(G)$ for its vertex and edge set, respectively. We write $\Omega(G)$ for its set of ends. If G is understood, we just write V, E and Ω for $V(G)$, $E(G)$ and $\Omega(G)$, respectively. Usually, we write ω for an end of a graph. However, this will cause a clash of notation in Chapter 5 with the smallest infinite ordinal; there we usually write ε for ends of graphs and there ω denotes the smallest infinite ordinal.

In Part III, we usually consider directed graphs without multiple edges or loops but possibly with inversely directed edges between distinct vertices. Also for directed graphs, our proofs and concepts (with obvious changes) work for directed multigraphs with all parallel edges of finite multiplicity. We will call directed graphs digraphs and if there is no risk of confusion, we shall omit the word 'directed', as in 'directed path' or 'directed ray'. We refer to 'edges' of a digraph and not to 'arcs' of a digraph. We write edges as ordered pairs (v, w) of vertices v, w of D, and we usually write (v, w) simply as vw. For a digraph D, we write $V(D)$ and $E(D)$ for its set of vertices and set of edges, respectively. Also, we write $\Omega(D)$ and $\Lambda(D)$ for its set of ends and set of limit edges, respectively.

Finally, we remark that each chapter will have its own preliminary section and its own introduction. This makes every chapter self-contained and the interested reader can read each chapter independently. However, this might cause some redundancy and I apologise to those readers who are bothered by it.

Part I.

End spaces

2. Approximating infinite graphs by normal trees

2.1. Introduction

A rooted tree T contained in a graph G is *normal* in G if the endvertices of every T-path in G are comparable in the tree-order of T. (In finite graphs, normal spanning trees are their depth-first search trees; see [25] for precise definitions.) Normal spanning trees are perhaps the most useful structural tool in infinite graph theory. Their importance arises from the fact that they capture the separation properties of the graph they span, and so in many situations it suffices to deal with the much simpler tree structure instead of the whole graph. For example, the end space of G coincides, even topologically, with the end space of any normal spanning tree of G. However, not every connected graph has a normal spanning tree, and the structure of graphs without normal spanning trees is still not completely understood [6, 30].

In order to harness and transfer the power of normal spanning trees to arbitrary connected graphs G, one might try to find an 'approximate normal spanning tree': a normal tree in G which spans the graph up to some arbitrarily small given error term. To formalize this idea, recall that, as usual, a *neighbourhood* of an end is the component of $G - X$ which contains a tail of every ray of that end, for some (arbitrarily large) finite set of vertices $X \subseteq V(G)$. We say that a graph G can be *approximated by normal trees* if for every selection of arbitrarily small neighbourhoods around its ends there is a normal tree $T \subseteq G$ such that every component of $G - T$ is included in one of the selected neighbourhoods and every end of G has some neighbourhood in G that avoids T.

Our approximation result for normal trees in infinite graphs then reads as follows:

Theorem 2.1. *Every connected graph can be approximated by normal trees.*

Note that the normal trees provided by our theorem will always be rayless, see also Section 2.3.

We indicate the potential of Theorem 2.1 by a number of applications. Our first two applications are of combinatorial nature: we exhibit in Section 2.4 two new existence results for normal spanning trees that Theorem 2.1 implies. One of these, Theorem 2.4.3, says that if every end of a connected graph G has a neighbourhood which has a normal spanning tree then G itself has a normal spanning tree.

Interestingly, Theorem 2.1 may not only be read as a structural result for connected graphs: it also implies and extends a number of previously hard results about topological properties of end spaces [21, 27, 28, 60, 61, 67]. Denote by $\Omega(G)$ the end space of a graph G, and by $|G|$ the space on $G \cup \Omega(G)$ naturally associated with the graph G and its ends; see the next section for precise definitions. When G is locally finite and connected, then $\Omega(G)$ is compact, and $|G|$ is the well-known Freudenthal compactification of G. For arbitrary G, the spaces $\Omega(G)$ and $|G|$ are usually non-compact and far from being completely understood.

Polat has shown that $\Omega(G)$ is ultrametrizable if and only if G contains a topologically end-faithful normal tree [60, Theorem 5.13], and has proved as a crucial auxiliary step that end spaces are always collectionwise normal [60, Lemma 4.14]. Changing focus from $\Omega(G)$ to $|G|$, Sprüssel has shown that $|G|$ is normal [67], and Diestel has characterised when $|G|$ is metrizable or compact [21] in terms of certain normal spanning trees in G. Our combinatorial Theorem 2.1 provides, in just a few lines, new and unified proofs for all these results. Additionally, Theorem 2.1 shows that metrizable end spaces are always ultrametrizable (Theorem 2.4.1), answering an open question by Polat.

Finally, Theorem 2.1 brings new progress to an old problem of Diestel, which asks for a topological charactersation of all end spaces [27, Problem 5.1]. Indeed, note that Theorem 2.1 translates to the topological assertion that every open cover of an end space can be refined to an open partition cover, Corollary 2.3.1. This last property is known in the literature as ultra-paracompactness. It implies that all spaces $|G|$ are paracompact (Corollary 2.3.2), and that all end spaces $\Omega(G)$ are even hereditarily ultra-paracompact (Corollary 2.5.3).

This chapter is organised as follows: The next section contains a recap on end spaces and other technical terms. Section 2.3 contains the proof of our main result, and Section 2.4 derives the consequences outlined above. Section 2.5 indicates a simple argument showing that subspaces of end spaces inherit their property of being ultra-paracompact.

2.2. End spaces of graphs: a reminder

For graph theoretic terms we follow the terminology in [25], and in particular [25, Chapter 8] for ends in graphs and the spaces $\Omega(G)$ and $|G|$. 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. The set of ends of a graph G is denoted by $\Omega = \Omega(G)$. If $X \subseteq V$ is finite and $\omega \in \Omega$, there is a unique component of $G - X$ that contains a tail of every ray in ω, which we denote by $C(X, \omega)$. If C is any component of $G - X$, we write $\Omega(X, C)$ for the set of ends ω of G with $C(X, \omega) = C$, and abbreviate $\Omega(X, \omega) := \Omega(X, C(X, \omega))$. Finally, if $\mathscr{C}$ is any collection of components of $G - X$, we write $\Omega(X, \mathscr{C}) := \bigcup \{ \Omega(X, C) \colon C \in \mathscr{C} \}$.

The collection of all sets $\Omega(X, C)$ with $X \subseteq V$ finite and C a component of $G - X$ form a basis for a topology on Ω. This topology is Hausdorff, and it is *zero-dimensional* in that it has a basis consisting of closed-and-open sets. Note that when considering end spaces $\Omega(G)$, we may always assume that G is connected; adding one new vertex and choosing a neighbour for it in each component does not affect the end space.

We now describe two common ways to extend this topology on $\Omega(G)$ to a topology on $|G| = G \cup \Omega(G)$, the graph G together with its ends. The first topology, called Top, has a basis formed by all open sets of G considered as a 1-complex, together with basic open neighbourhoods for ends of the form

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

where $\mathring{E}_*(X, C(X, \omega))$ denotes any union of half-open intervals of all the edges from the edge cut $E(X, C(X, \omega))$ with endpoint in $C(X, \omega)$.

As the 1-complex topology on G is not first-countable at vertices of infinite degree, it is sometimes useful to consider a metric topology on G instead: The second topology commonly considered, called MTop, has a basis formed by all open sets of G considered as a metric length-space (i.e. every edge together with its endvertices forms a unit interval of length 1, and the distance between two points of the graph is the length of a shortest arc in G between them), together with basic open neighbourhoods for ends of the form

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

where $\mathring{E}_\varepsilon(X, C(X,\omega))$ denotes the open ball around $C(X,\omega)$ in G of radius ε. Note that both topologies Top and MTop induce the same subspace topology on $\hat{V}(G) := V(G) \cup \Omega(G)$ and $\Omega(G)$, the last of which coincides with the topology on $\Omega(G)$ described above. Polat observed that $\hat{V}(G)$ is homeomorphic with $\Omega(G^+)$, where G^+ denotes the graph obtained from G by gluing a new ray R_v onto each vertex v of G so that R_v meets G precisely in its first vertex v and R_v is distinct from all other $R_{v'}$, cf. [60, §4.16].

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 2.2.1 (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 $X \subseteq V(G)$.*

The *tree-order* of a rooted tree (T,r) 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* of T is the set of vertices at distance n from r in T. The *down-closure* of a vertex v is the set $\lceil v \rceil := \{\, u \colon u \le v \,\}$; its *up-closure* is the set $\lfloor v \rfloor := \{\, w \colon v \le w \,\}$. The down-closure of v is always a finite chain, the vertex set of the path rTv. A ray $R \subseteq T$ starting at the root is called a *normal ray* of T.

A rooted tree T contained in a graph G is *normal* in G if the endvertices of every T-path in G are comparable in the tree-order of T. Here, for a given graph H, a path P is said to be an *H-path* if P is non-trivial and meets H exactly in its endvertices. We remark that for a normal tree $T \subseteq G$ the neighbourhood $N(D)$ of every component D of $G - T$ forms a chain in T. A set U of vertices is *dispersed* in G if for every end ω there is a finite $X \subseteq V$ with $C(X,\omega) \cap U = \emptyset$, or equivalently, if U is a closed subset of $|G|$ (in either Top or MTop).

Theorem 2.2.2 (Jung [50]). *A vertex set in a connected graph is dispersed if and only if there is a rayless normal tree including it. Moreover, every rayless normal tree in a connected graph can be extended to a rayless normal tree that includes an arbitrary pre-specified dispersed vertex set of the graph. As a consequence, a connected graph has a normal spanning tree if and only if its vertex set is a countable union of dispersed sets.*

If H is a subgraph of G, then rays equivalent in H remain equivalent in G; in other words, every end of H can be interpreted as a subset of an end of G, so the natural inclusion map $\iota \colon \Omega(H) \to \Omega(G)$ is well-defined. A subgraph $H \subseteq G$ is *end-faithful* if this inclusion map ι is a bijection. The terms *end-injective* and *end-surjective* are defined accordingly. Normal trees are always end-injective; hence, normal trees are end-faithful as soon as they are end-surjective. Given a subgraph $H \subseteq G$, write $\partial_\Omega H \subseteq \Omega(G)$ for the set of ends ω of G which satisfy $C(X,\omega) \cap H \neq \emptyset$ for all finite $X \subseteq V(G)$.

For topological notions we follow the terminology in [33]. All spaces considered in this chapter are Hausdorff, i.e. every two distinct points have disjoint open neighbourhoods. An *ultrametric* space (X, d) is a metric space in which the triangle inequality is strengthened to $d(x,z) \leq \max\{d(x,y), d(y,z)\}$. A topological space X is *ultrametrizable* if there is an ultrametric d on X which induces the topology of X. A topological space is *normal* if for any two disjoint closed sets A_1, A_2 there are disjoint open sets U_1, U_2 with $A_i \subseteq U_i$. A space is *collectionwise normal* if for every *discrete* family $\{ A_s \colon s \in S \}$ of disjoint closed sets, i.e. a family such that $\bigcup \{ A_s \colon s \in S' \}$ is closed for any $S' \subseteq S$, there is a collection $\{ U_s \colon s \in S \}$ of disjoint open sets with $A_s \subseteq U_s$.

A collection $\mathcal{A}$ of sets is said to *refine* another collection $\mathcal{B}$ of sets if for every $A \in \mathcal{A}$ there is $B \in \mathcal{B}$ with $A \subseteq B$. A cover $\mathcal{V}$ of a topological space X is *locally finite* if every point of X has an open neighbourhood which meets only finitely many elements of $\mathcal{V}$. A topological space X is *paracompact* if for every open cover $\mathcal{U}$ of X there is a locally finite open cover $\mathcal{V}$ refining $\mathcal{U}$. All compact Hausdorff spaces and also all metric spaces are paracompact, which in turn are always normal and even collectionwise normal [33, Chapter 5.1]. A space is *ultra-paracompact* if every open cover has a refinement by an open partition.

Lastly, ordinal numbers are identified with the set of all smaller ordinals, i.e. $\alpha = \{ \beta \colon \beta < \alpha \}$ for all ordinals α.

2.3. Proof of the main result

This section is devoted to the proof of our main theorem, which we restate more formally:

Theorem 2.1. *For every collection $\mathscr{C} = \{ C(X_\omega, \omega) \colon \omega \in \Omega(G) \}$ in a connected graph G, there is a rayless normal tree T in G such that every component of $G - T$ is included in an element of $\mathscr{C}$.*

As every rayless normal tree $T \subseteq G$ is dispersed in G by Jung's Theorem 2.2.2, this technical variant of our main result is clearly equivalent to the formulation presented in the introduction.

Let us briefly discuss two other possible notions of 'approximating graphs by normal trees': First, Theorem 2.1 is significantly stronger than just requiring that (every component of) $G - T$ is included in the union $\bigcup \mathscr{C}$ of the selected neighbourhoods; the latter assertion is easily seen to be equivalent to Jung's Theorem 2.2.2. In the other direction, could one strengthen our notion of 'approximating by normal trees' and demand a normal rayless tree T such that for every end ω of G, the component of $G - T$ in which every ray of ω has a tail is included in $C(X_\omega, \omega)$? This notion, however, is too strong and such a T may not exist: Consider the graph $G = K^+$ (see Section 2.2) for an uncountable clique K, and let $\mathscr{C}$ be the collection of all the ray-components of $G - K$ (together with an arbitrary neighbourhood of the end of the clique K). Any normal tree for G satisfying our stronger requirements would restrict to a normal spanning tree of K, an impossibility.

Proof of Theorem 2.1. Given a collection $\mathscr{C} = \{ C(X_\omega, \omega) \colon \omega \in \Omega(G) \}$ in a connected graph G, call a subgraph $H \subseteq G$ *bounded* if there is an end $\omega \in \Omega(G)$ with $H \subseteq C(X_\omega, \omega)$, and *unbounded* otherwise.

We construct a sequence of rayless normal trees $T_1 \subseteq T_2 \subseteq \ldots$ extending each other all with the same root r as follows: Let T_1 be the tree on a single vertex r (for some arbitrarily chosen vertex $r \in G$) and suppose that T_n has already been constructed.

We claim that for every unbounded component D of $G - T_n$ there exists a finite separator $S_D \subseteq V(D)$ such that $D - S_D$ has either zero or at least two unbounded components. To see this, suppose the contrary; then the map d sending each finite vertex set in D to its unique unbounded component is a direction on D and hence defines an end ω of D by Theorem 2.2.1. But $d(X_\omega \cap D) \subseteq C(X_\omega, \omega)$ is bounded, a contradiction.

Now for every such unbounded D let S_D be a finite separator of the first kind in D if possible, and otherwise of the second kind. Note that the union of all finite vertex sets S_D is dispersed in G because T_n is a rayless normal tree. Since G is connected, we may use Jung's Theorem 2.2.2 to extend T_n in an inclusion minimal way to a rayless normal tree $T_{n+1} \supseteq T_n$ with root r that includes all the finite vertex sets S_D. This completes the construction.

Now consider the normal tree $T = \bigcup_{n \in \mathbb{N}} T_n$. We claim that T is rayless. Indeed, suppose otherwise, that there is a normal ray R in T belonging to the end $\omega \in \Omega(G)$ say.

Then, for every $n \in \mathbb{N}$, the ray R has a tail in an unbounded component D_n of $G - T_n$, and all finite separators S_{D_n} chosen for these components were of the second kind, since we never extended T_n into a component that was already bounded. In particular R meets each S_{D_n} in at least one vertex, s_n say. Now, fix for every S_{D_n} an unbounded component C_{n+1} of $D_n - S_{D_n}$ different from D_{n+1}. Every C_{n+1} has a neighbour, say u_n, in S_{D_n}. Moreover, the paths $P_n = s_n T u_n$ connecting s_n to u_n in T are pairwise disjoint, as each of them was constructed in the nth step.

From this, we obtain a contradiction as follows. Choose $n \in \mathbb{N}$ large enough so that X_ω avoids all of C_{n+1}, the path P_n and the tail of R that starts at the vertex s_n (this is indeed possible since any two components C_i and C_j are disjoint). Now, C_{n+1} is contained in $C(X_\omega, \omega)$, contradicting the fact that C_{n+1} is unbounded. This shows that ω cannot exist, and hence that T is rayless.

Finally, we claim that every component D of $G - T = G - \bigcup_{n \in \mathbb{N}} T_n$ is bounded. Since T is a normal tree, $N(D)$ is a chain in T, and since T is rayless, $N(D)$ is finite. Hence, there is $m \in \mathbb{N}$ such that $N(D) \subseteq T_m$, i.e. D is already a component of $G - T_m$. The fact that we have not extended T_m into D means that D is bounded. □

Corollary 2.3.1. *For every connected graph G and every open cover $\mathcal{U}$ of its end space $\Omega(G)$ there is a rayless normal tree T in G such that the collection of components of $G - T$ induces an open partition of $\Omega(G)$ refining $\mathcal{U}$. In particular, all end spaces $\Omega(G)$ are ultra-paracompact.*

Proof. Without loss of generality, the open cover $\mathcal{U}$ is of the form $\mathcal{U} = \{ \Omega(X_\omega, \omega) \colon \omega \in \Omega(G) \}$. Theorem 2.1 applied to $\mathscr{C} = \{ C(X_\omega, \omega) \colon \omega \in \Omega(G) \}$ yields a rayless normal tree T in G such that every component of $G - T$ is included in an element of $\mathscr{C}$. As every component of $G - T$ has a finite neighbourhood, it induces an open set of the end space. This gives the desired open partition of $\Omega(G)$ refining $\mathcal{U}$. □

Corollary 2.3.2. *All spaces $|G|$ are paracompact in both* TOP *and* MTOP.

Proof. First, we consider $|G|$ with MTOP. To show that $|G|$ is paracompact, suppose that any open cover $\mathcal{U}$ of $|G|$ consisting of basic open sets is given. The cover elements come in two types: basic open sets of G, and basic open neighbourhoods of ends. We write $\mathcal{U}_\Omega = \{ \hat{C}_{\varepsilon_i}(X_i, \omega_i) : i \in I \}$ for the collection consisting of the latter. As $\mathcal{U}_\Omega$ covers the end space of G, applying Theorem 2.1 to the collection $\mathscr{C} := \{ C(X_i, \omega_i) : i \in I \}$ yields a rayless normal tree T in G such that $\{ C(Y_j, \omega_j) : j \in J \}$, the collection of components of $G - T$ containing a ray, refines $\mathscr{C}$. For every $j \in J$ we choose $\varepsilon_j := \varepsilon_i$ for some $i \in I$ with $C(Y_j, \omega_j) \subseteq C(X_i, \omega_i)$, ensuring that the disjoint collection $\mathcal{V}_\Omega := \{ \hat{C}_{\varepsilon_j}(Y_j, \omega_j) : j \in J \}$ refines $\mathcal{U}_\Omega$.

Next, consider the quotient space H that is obtained from $|G|$ by collapsing every closed subset $C(Y_j, \omega_j) \cup \Omega(Y_j, \omega_j)$ with $j \in J$ to a single point. As the open sets in $\mathcal{V}_\Omega$ are disjoint, the quotient is well-defined and we may view H as a rayless multi-graph endowed with MTOP. Now consider the open cover $\mathcal{U}_H$ of H that consists of the quotients of the elements of $\mathcal{V}_\Omega$ on the one hand, and on the other hand, for every non-contraction point of H a choice of one basic open neighbourhood in G that is contained in some element of $\mathcal{U}$. Since metric spaces are paracompact, H admits a locally finite refinement $\mathcal{V}_H$ of $\mathcal{U}_H$ consisting of basic open sets of (H, MTOP). Then the open cover $\mathcal{V}$ of $|G|$ induced by $\mathcal{V}_H$ gives the desired locally finite refinement of $\mathcal{U}$.

A similar argument shows that $|G|$ with TOP is paracompact. Here, (H, TOP) is paracompact because all CW-complexes are. □

Note in particular that paracompactness implies normality and collectionwise normality, and hence we reobtain the previously mentioned results by Polat [60, Lemma 4.14] and Sprüssel [67, Theorems 4.1 & 4.2] as a straightforward consequence of our Corollary 2.3.2.

2.4. Consequences of the approximation result

In [60, Theorem 5.13] Polat characterised the graphs that admit an end-faithful normal tree as the graphs with ultrametrizable end space, and raised the question [61, §10] whether metrizability of the end space is enough to ensure the existence of an end-faithful normal tree. As our first application we show how using Theorem 2.1 provides a much simplified proof for Polat's result that simultaneously answers his question about the metrizable case in the affirmative:

Theorem 2.4.1. *For every connected graph G, the following are equivalent:*

(i) *The end space of G is metrizable,*
(ii) *the end space of G is ultrametrizable,*
(iii) *G contains an end-faithful normal tree.*

Proof. The implication (iii) ⇒ (ii) is routine, as the end space of any tree is ultrametrizable (see e.g. [48] for a detailed account), and $\Omega(T)$ and $\Omega(G)$ are homeomorphic for every end-faithful normal tree T of G (see e.g. [27, Proposition 5.5]). The implication (ii) ⇒ (i) is trivial.

Hence, it remains to prove (i) $\Rightarrow$ (iii). For this, consider the covers $\mathcal{U}_n$ for $n \in \mathbb{N}$ of $\Omega(G)$ given by the open balls with radius $1/n$ around every end; with respect to some fixed metric d inducing the topology of $\Omega(G)$. By applying Corollary 2.3.1 to the covers $\mathcal{U}_1, \mathcal{U}_2, \ldots$, and combining it with Jung's Theorem 2.2.2, it is straightforward to construct a sequence of rayless normal trees $T_1 \subseteq T_2 \subseteq \ldots$ all rooted at the same vertex such that the partition of $\Omega(G)$ given by the components of $G - T_n$ refines $\mathcal{U}_n$.

Any two ends $\omega \neq \eta$ of G are separated by any T_n with $2/n < d(\omega, \eta)$. Consider the normal tree $T' = \bigcup_{n\in\mathbb{N}} T_n$. We claim that each end $\omega \in \Omega(G) \setminus \partial_\Omega T'$ belongs to a component C of $G - T'$ such that $N(C)$ is finite. Otherwise $N(C)$ lies on a unique normal ray R of T belonging to some end $\eta \in \partial_\Omega T'$, but then clearly, none of the T_n would separate ω from η, a contradiction. Hence, $N(C)$ is finite, and since C contains at most one end, T' extends to an end-faithful normal tree of G. □

From the new implication (i) $\Rightarrow$ (iii) in Theorem 2.4.1 one also obtains a simple proof of Diestel's characterisation from [21] when $|G|$ is metrizable.

Corollary 2.4.2. *For every connected graph G, the following are equivalent:*

(i) *$|G|$ with* MTop *is metrizable,*
(ii) *the space $\hat{V}(G)$ is metrizable,*
(iii) *G has a normal spanning tree.*

Proof. The first implication (iii) $\Rightarrow$ (i) is routine, see e.g. [21]. The implication (i) $\Rightarrow$ (ii) is trivial. For (ii) $\Rightarrow$ (iii) apply Theorem 2.4.1 to the space $\Omega(G^+) \cong \hat{V}(G)$, noting that every end-faithful normal tree of G^+ is automatically spanning. □

To motivate our next applications, suppose that a given graph G admits a normal spanning tree. Let us call such graphs *normally spanned.* If G is normally spanned, then every component of $G - X$ is normally spanned, too, for any finite $X \subseteq V(G)$. Conversely, the question arises whether a graph admits a normal spanning tree as soon as every end ω has some basic neighbourhood $C(X, \omega)$ that is normally spanned. It turns out that the answer is yes:

Theorem 2.4.3. *If every end of a connected graph G has a normally spanned neighbourhood, then G itself is normally spanned.*

Proof. Let $\mathscr{C} = \{\, C(X_\omega, \omega) \colon \omega \in \Omega(G) \,\}$ be a selection of normally spanned neighbourhoods for all ends of G, and apply Theorem 2.1 to $\mathscr{C}$ to find a rayless normal tree T such that the collection of components of $G-T$ refines $\mathscr{C}$. By Jung's Theorem 2.2.2, each such component C of $G-T$ is the union of countably many dispersed sets, say $V(C) = \bigcup_{n\geq 1} V_n^C$. But then $V_0 = V(T)$ together with all the sets $V_n := \bigcup \{\, V_n^C \colon C \text{ a component of } G-T \,\}$, for $n \geq 1$, witnesses that $V(G)$ is a countable union of dispersed sets. Hence, G has a normal spanning tree by Jung's theorem. □

There is also a more topological viewpoint on the above result: The assumptions of Theorem 2.4.3 are by Corollary 2.4.2 equivalent to the assertion that $\hat{V}(G)$ is locally metrizable. But locally metrizable paracompact spaces are metrizable, [33, Exercise 5.4.A]. Hence, applying Corollary 2.4.2 once again to $\hat{V}(G)$ yields the desired normal spanning tree of G.

Continuing along these lines, we now address the question whether the existence of some *local end-faithful normal tree* for every end of a graph already ensures the existence of an end-faithful normal tree of the entire graph. For a graph G and an end ω, we say that ω has a *local end-faithful normal tree* if there is a normal tree T in G such that $\partial_\Omega T$ is a neighbourhood of ω in $\Omega(G)$.

Theorem 2.4.4. *If every end of a connected graph G has a local end-faithful normal tree, then G has an end-faithful normal tree.*

Proof. By Theorem 2.4.1 every end in $\Omega(G)$ has a metrizable neighbourhood. But (ultra-)paracompact spaces which are locally metrizable are metrizable, [33, Exercise 5.4.A]. Consequently, we have by Corollary 2.3.1 that $\Omega(G)$ is metrizable. Applying again Theorem 2.4.1 yields the desired end-faithful normal tree of G. □

2.5. Paracompactness in subspaces of end spaces

We conclude this chapter with an observation concerning the following fundamental problem on the structure of end spaces raised by Diestel in 1992 [27, Problem 5.1]:

Problem 2.5.1. *Which topological spaces can be represented as an end space $\Omega(G)$ for some graph G?*

In Corollary 2.3.1 we established that end spaces are always ultra-paracompact. In this section we show that also all subspaces of end spaces inherit the property of being ultra-paracompact, i.e. that end spaces are hereditarily ultra-paracompact. This significantly reduces the number of topological candidates for a solution of Problem 2.5.1, and for example shows that certain compact spaces cannot occur as end space, which Corollary 2.3.1 wouldn't do on its own.

It is known that paracompactness and ultra-paracompactness, along with a number of other properties which are not per se hereditary such as normality and collectionwise normality, have the property that they are inherited by *all* subspaces as soon as they are inherited by all *open* subspaces. For the easy proof in case of paracompactness see e.g. Dieudonné's original paper [32, p. 68]. Hence, our assertion follows at once from Corollary 2.3.1 given the following observation:

Lemma 2.5.2. *Open subsets of end spaces are again end spaces.*

Proof. Let G be any graph, and consider some open, non-empty set $\Gamma \subseteq \Omega(G)$. Write $\Gamma^\complement$ for its complement in $\Omega(G)$. Using Zorn's lemma, pick a maximal collection $\mathcal{R}$ of disjoint rays all belonging to ends in $\Gamma^\complement$, and let W be the union $\bigcup \{ V(R) \colon R \in \mathcal{R} \}$ of their vertex sets. Note that $\partial_\Omega W \subseteq \Gamma^\complement$ because $\Gamma^\complement$ is closed. We claim that Γ is homeomorphic to the end space of the graph $G' := G - W$.

In order to find a homeomorphism $\varphi \colon \Omega(G') \to \Gamma$, note first that, due to the maximality of $\mathcal{R}$, every ray in G' is (as a ray of G) contained in an end of Γ. Consequently, every end ω' of G' is contained in a unique end ω of Γ and we define φ via this correspondence.

To see that φ is surjective, consider an open neighbourhood $\Omega(X, \omega) \subseteq \Gamma$, for a given $\omega \in \Gamma$. Then W has only finite intersection with $C(X, \omega)$, as only finitely many rays from $\mathcal{R}$ can intersect $C(X,\omega)$, but do not have a tail in $C(X,\omega)$. So we may assume that $C(X,\omega)$ is contained in G', by extending X. Now, every ray of ω contained in $C(X,\omega)$ gives an end in G' that is mapped to ω.

To see that φ is injective, suppose there are two rays R_1, R_2 in G' that are not equivalent in G' but equivalent in G. Then, there are infinitely many pairwise disjoint R_1-R_2 paths in G and all but finitely many of these paths hit W. Then the end ω of G containing R_1 and R_2 is an end in Γ which lies in the closure of $\Gamma^{\complement}$, contradicting the fact that $\Gamma^{\complement}$ is closed.

Finally, let us show that φ is continuous and open. For the continuity of φ remember that for any open set $\Omega(X,\omega) \subseteq \Gamma$ we may assume that $C(X,\omega)$ is contained in G'. In particular the preimage of $\Omega(X,\omega)$ is open in G'.

For φ being open, consider an open set $\Omega(X,\omega') \subseteq \Omega(G')$. Now, $C(X,\omega') \subseteq G' - X$ might not be a component of $G - X$. However, the set of vertices in $C(X,\omega')$ having a neighbour in W is dispersed. Again by extending X, we may assume that $C(X,\omega')$ is a component of $G - X$. Consequently, its image is open in $\Omega(G)$. □

Corollary 2.5.3. *All end spaces are hereditarily ultra-paracompact.* □

Interestingly, a careful reading of Sprüssel's proof that spaces $|G|$ are normal from [67] establishes that every end space $\Omega(G)$ is in fact *completely normal*, i.e. that subsets with $\overline{A} \cap B = \emptyset = A \cap \overline{B}$ can be separated by disjoint open sets – a property which is equivalent to hereditary normality, see [33, Theorem 2.1.7]. In any case, also this stronger result of hereditary normality is implied by our paracompactness result in Corollary 2.5.3.

Future steps. The results in this chapter narrow down which topological spaces occur as end spaces. Still, a complete solution to Problem 2.5.1 seems currently out of reach without significant new insights. We remark that similar to Theorem 2.4.1, one can show that end spaces $\Omega(G)$ have the property that every separable subspace $X \subseteq \Omega(G)$, i.e. every such X with a countable dense subset, must be metrizable. We currently do not know of an example of a hereditarily ultra-paracompact space where every separable subspace is metrizable that does not occur as the end space $\Omega(G)$ of some graph G.

3. Countably determined ends and graphs

3.1. Introduction

Halin [45] defined the ends of an infinite graph 'from below' as equivalence classes of rays in the graph, where two rays are equivalent if no finite set of vertices separates them. As a complementary description of Halin's ends, Diestel and Kühn [28] introduced the notion of directions of infinite graphs. These are defined 'from above': A *direction* of a graph G is a map f, with domain the collection $\mathcal{X} = \mathcal{X}(G)$ of all finite vertex sets of G, that assigns to every $X \in \mathcal{X}$ a component $f(X)$ of $G - X$ such that $f(X) \supseteq f(X')$ whenever $X \subseteq X'$.

Every end ω of G defines a direction f_ω of G by letting $f_\omega(X)$ be the component of $G - X$ that contains a subray of every ray in ω. Diestel and Kühn showed that the natural map $\omega \mapsto f_\omega$ is in fact a bijection between the ends of G and its directions. This correspondence is now well known and has become a standard tool in the study of infinite graphs. See [12, 21, 23, 54–56] for examples.

The domain of the directions of G might be arbitrarily large as its size is equal to the order of G. This contrasts with the fact that every direction of G is induced by a ray of G and rays have countable order. Hence the question arises whether every direction of G is 'countably determined' in G also by a countable subset of its choices. A *directional choice* in G is a pair (X, C) of a finite vertex set $X \in \mathcal{X}$ and a component C of $G - X$. We say that a directional choice (X, C) in G *distinguishes* a direction f from another direction h if $f(X) = C$ and $h(X) \neq C$. A direction f of G is *countably determined* in G if there is a countable set of directional choices in G that distinguish f from every other direction of G.

Curiously, the answer to this question is in the negative: Consider the graph G that arises from the uncountable complete graph $K^{\aleph_1}$ by adding a new ray R_v for every vertex $v \in K^{\aleph_1}$ so that R_v meets $K^{\aleph_1}$ precisely in its first vertex v and R_v is disjoint from all the other new rays $R_{v'}$. Then $K^{\aleph_1} \subseteq G$ induces a direction of G that is not countably determined in G.

This example raises the question of which directions of a given graph G are countably determined. In the first half of this chapter we answer this question: we characterise for every graph G, by unavoidable substructures, both the countably determined directions of G and its directions that are not countably determined.

If $R \subseteq G$ is any ray, then every finite initial segment X of R naturally defines a directional choice in G, namely (X, C) for the component C that contains $R - X$. Let us call R *directional* in G if its induced direction is distinguished from every other direction of G by the directional choices that are defined by R. By definition, every direction of G that is induced by a directional ray is countably determined in G. Surprisingly, our characterisation implies that the converse holds as well: if a direction of G is distinguished from every other direction by countably many directional choices (X, C), then no matter how the vertex sets X lie in G we can always assume that the sets X are the finite initial segments of a directional ray.

Theorem 3.1. *For every graph G and every direction f of G the following assertions are equivalent:*

(i) *The direction f is countably determined in G.*
(ii) *The direction f is induced by a directional ray of G.*

As our second main result we characterise by unavoidable substructures the directions of any given graph that are not countably determined in that graph, and thereby complement our first characterisation. Our theorem is phrased in terms of substructures that are uncountable star-like combinations either of rays or of double rays. Recall that a vertex v of a graph G *dominates* a ray $R \subseteq G$ if there is an infinite v–R fan in G. An end of G is *dominated* if one (equivalently: each) of its rays is dominated, see [25]. Given a direction f of G we write ω_f for the unique end ω of G whose rays induce f, i.e., which satisfies $f_\omega = f$. If G is a graph and $(T, \mathcal{V})$ is a tree-decomposition of G that has finite adhesion, then every direction of G either corresponds to a direction of T or lives in a part of $(T, \mathcal{V})$; see Section 3.2.3. An *uncountable star-decomposition* is a tree-decomposition whose decomposition tree is a star $K_{1,\kappa}$ for some uncountable cardinal κ.

Theorem 3.2. *For every graph G and every direction f of G the following assertions are equivalent:*

(i) *The direction f is not countably determined in G.*
(ii) *The graph G contains either*
   - *uncountably many disjoint pairwise inequivalent rays all of which start at vertices that dominate ω_f, or*
   - *uncountably many disjoint double rays, all having one tail in ω_f and another not in ω_f, so that the latter tails are inequivalent for distinct double rays.*

Moreover, if (ii) *holds, we can find the (double) rays together with an uncountable star-decomposition of G of finite adhesion such that f lives in the central part and each (double) ray has a tail in its own leaf part.*

Note that (ii) clearly implies (i).

Does the local property that every direction of G is countably determined in G imply the stronger global property that there is one countable set of directional choices that distinguish every two directions of G from each other? We answer this question in the negative; see Example 3.4.1. Let us call a graph G *countably determined* if there is a countable set of directional choices in G that distinguish every two directions of G from each other.

In the second half of this chapter we structurally characterise both the graphs that are countably determined and the graphs that are not countably determined. A rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T-path in G are comparable in the tree-order of T, cf. [25]. (A *T-path* in G is a non-trivial path that meets T exactly in its endvertices.)

Theorem 3.3. *For every connected graph G the following assertions are equivalent:*

(i) *G is countably determined.*
(ii) *G contains a countable normal tree that contains a ray from every end of G.*

Complementing this characterisation we structurally characterise, as our fourth main result of this chapter, the graphs that are not countably determined.

Theorem 3.4. *For every connected graph G the following assertions are equivalent:*

(i) *G is not countably determined.*
(ii) *G has an uncountable star-decomposition of finite adhesion such that in every leaf part there lives a direction of G.*

Interestingly, countably determined directions and countably determined graphs admit natural topological interpretations. Over the course of the last two decades, the topological properties of end spaces have been extensively investigated, see e.g. [21, 28, 60, 61, 67]. However, not much is known about such fundamental properties as countability axioms. Recall that a topological space is *first countable* at a given point if it has a countable neighbourhood base at that point. A direction of a graph G is countably determined in G if and only if it is defined by an end that has a countable neighbourhood base in the end space of G (Theorem 3.3.2). Thus, Theorems 3.1 and 3.2 characterise combinatorially when the end space of a graph is first countable or not first countable at a given end, respectively. Similarly, a graph is countably determined if and only if its end space is *second countable* in that its entire topology has a countable base (Theorem 3.4.7). Therefore, Theorems 3.3 and 3.4 characterise combinatorially the infinite graphs whose end spaces are second countable or not second countable, respectively. Furthermore, our four theorems imply similar results for the space $|G|$ formed by a graph G together with its end space; see Section 3.5.

This chapter is organised as follows: In the next section we give a reminder on end spaces and recall all the results from graph theory and general topology that we need. We prove Theorems 3.1 and 3.2 in Section 3.3 and we prove Theorems 3.3 and 3.4 in Section 3.4. Finally, in Section 3.5 we consider the spaces $|G|$.

3.2. Preliminaries

For graph theoretic terms we follow the terminology in [25]. For topological notions we follow the terminology in [33].

3.2.1. Ends of graphs

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. The set of ends of a graph G is denoted by $\Omega = \Omega(G)$. If $X \subseteq V$ is finite and $\omega \in \Omega$, there is a unique component of $G - X$ that contains a tail of every ray in ω, which we denote by $C(X, \omega)$. If C is any component of $G - X$, we write $\Omega(X, C)$ for the set of ends ω of G with $C(X, \omega) = C$, and abbreviate $\Omega(X,\omega) := \Omega(X, C(X,\omega))$. The ends in $\Omega(X,\omega)$ are said to *live* in $C(X,\omega)$. For two ends ω_1 and ω_2 a finite vertex set $X \subseteq V$ *separates* ω_1 and ω_2 if they live in distinct components of $G - X$.

The collection of sets $\Omega(X,C)$ with $X \subseteq V$ finite and C a component of $G - X$ form a basis for a topology on Ω. This topology is Hausdorff, and it has a basis consisting of closed-and-open sets. The space $\Omega(G)$ with this topology is called the *end space* of G.

A vertex v of G *dominates* a ray $R \subseteq G$ if there is an infinite v–R fan in G. An end of G is *dominated* by v if one (equivalently: each) of its rays is dominated by v. If a vertex v of G dominates an end ω of G, then $v \in C(X,\omega)$ for all finite sets $X \subseteq V(G)$ with $v \notin X$.

Recall that a *comb* is the union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R. The last vertices of those paths are the *teeth* of this comb. Given a vertex set U, a *comb attached to U* is a comb with all its teeth in U, and a *star attached to U* is a subdivided infinite star with all its leaves in U. Then the set of teeth is the *attachment set* of the comb, and the set of leaves is the *attachment set* of the star. The following lemma is [25, Lemma 8.2.2]:

Lemma 3.2.1 (Star-Comb Lemma). *Let U be an infinite set of vertices in a connected graph G. Then G contains either a comb attached to U or a star attached to U.*

Let us say that an end ω of G is contained *in the closure* of M, where M is either a subgraph of G or a set of vertices of G, if for every $X \in \mathcal{X}$ the component $C(X,\omega)$ meets M. Equivalently, ω lies in the closure of M if and only if G contains a comb attached to M with its spine in ω. We write $\partial_\Omega M$ for the subset of Ω that consists of the ends of G lying in the closure of M. A vertex set $U \subseteq V(G)$ is *dispersed* in G if $\partial_\Omega U = \emptyset$.

3.2.2. Normal trees

The *tree-order* of a rooted tree $T = (T, r)$ 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* of T is the set of vertices at distance n from r in T. The *down-closure* of a vertex v is the set $\lceil v \rceil := \{\, u\colon u \le v \,\}$; its *up-closure* is the set $\lfloor v \rfloor := \{\, w\colon v \le w \,\}$. The down-closure of v is always a finite chain, the vertex set of the path rTv. A ray $R \subseteq T$ starting at the root is called a *normal ray* of T.

A rooted tree T contained in a graph G is *normal* in G if the endvertices of every T-path in G are comparable in the tree-order of T. Here, for a given graph H, a path P is said to be an *H-path* if P is non-trivial and meets H exactly in its endvertices. We remark that for a normal tree $T \subseteq G$ the neighbourhood $N(C)$ of every component C of $G - T$ forms a chain in T.

The *generalised up-closure* $\lfloor\!\lfloor x \rfloor\!\rfloor$ of a vertex $x \in T$ is the union of $\lfloor x \rfloor$ with the vertex set of $\bigcup \mathscr{C}(x)$, where the set $\mathscr{C}(x)$ consists of those components of $G - T$ whose neighbourhoods meet $\lfloor x \rfloor$. Every graph G reflects the separation properties of each normal tree $T \subseteq G$:

Lemma 3.2.2 ([12, Lemma 2.10])**.** *Let G be any graph and let $T \subseteq G$ be any normal tree.*

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

As a consequence, the normal rays of a normal spanning tree $T \subseteq G$ reflect the end structure of G in that every end of G contains exactly one normal ray of T, [25, Lemma 8.2.3]. More generally:

Lemma 3.2.3 ([12, Lemma 2.11])**.** *If G is any graph and $T \subseteq G$ is any normal tree, then every end of G in the closure of T contains exactly one normal ray of T. Moreover, sending these ends to the normal rays they contain defines a bijection between $\partial_\Omega T$ and the normal rays of T.*

Not every connected graph has a normal spanning tree. However, every countable connected graph does. More generally:

Lemma 3.2.4 (Jung [50], [12, Corollary 3.3])**.** *Let G be any graph and let $U \subseteq V(G)$ be any vertex set. If U is countable and v is any vertex of G, then G contains a normal tree that contains U cofinally and is rooted in v.*

If H is a subgraph of G, then rays equivalent in H remain equivalent in G; in other words, every end of H can be interpreted as a subset of an end of G, so the natural inclusion map $\iota\colon \Omega(H) \to \Omega(G)$ is well-defined. A subgraph $H \subseteq G$ is *end-faithful* if this inclusion map ι is a bijection. The terms *end-injective* and *end-surjective* are defined accordingly. Normal trees are always end-injective; hence, normal trees are end-faithful as soon as they are end-surjective.

3.2.3. Tree-decompositions, separations and ends

Recall from [25, Section 12.5] that a *tree-decomposition* of a graph G is an ordered pair $(T, \mathcal{V})$ of a tree T and a family $\mathcal{V} = (V_t)_{t \in V(T)}$ of *parts* $V_t \subseteq V(G)$ such that:

(i) $V(G) = \bigcup_{t \in T} V_t$;
(ii) every edge of G has both endvertices in V_t for some t;
(iii) $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever $t_2 \in t_1 T t_3$.

When we introduce a tree-decomposition as $(T, \mathcal{V})$ we tacitly assume $\mathcal{V} = (V_t)_{t \in T}$. Every edge of T induces a 'separation' of G, as follows.

A *separation* of a graph G is an unordered pair $\{A, B\}$ with $A \cup B = V(G)$ and such that no edge of G 'jumps' the *separator* $A \cap B$, meaning that no edge of G runs between $A \setminus B$ and $B \setminus A$. Every edge $t_1 t_2 \in T$ *induces* the separation $\{U_1, U_2\}$ of G that is defined by $U_1 := \bigcup_{t \in T_1} V_t$ and $U_2 := \bigcup_{t \in T_2} V_t$, where T_1 and T_2 are the components of $T - t_1 t_2$ containing t_1 and t_2 respectively. Then the separator $U_1 \cap U_2 = V_{t_1} \cap V_{t_2}$ is an *adhesion set* of $(T, \mathcal{V})$. We usually refer to the adhesion sets as separators. The tree-decomposition $(T, \mathcal{V})$ has *finite adhesion* if all its adhesion sets are finite.

Suppose now that $(T,\mathcal{V})$ has finite adhesion. Then every end ω of G either 'lives' in a unique part V_t or 'corresponds' to a unique end η of T. Here, ω *lives* in V_t if some (equivalently: every) ray in ω has infinitely many vertices in V_t. And ω *corresponds* to η if some (equivalently: every) ray $R\in\omega$ follows the course of some (equivalently: every) ray $S\in\eta$ (in that for every tail $S'\subseteq S$ the ray R has infinitely many vertices in $\bigcup_{t\in S'}V_t$). If f is a direction of G, then f *lives* in the part of $(T,\mathcal{V})$ or *corresponds* to the end of T that ω_f lives in or corresponds to, respectively.

If $(T,\mathcal{V})$ is a star-decomposition and s is the centre of the star T, then the induced separations of $(T,\mathcal{V})$ form a 'star of separations', as follows. Every separation $\{A,B\}$ has two *orientations*, the *oriented separations* (A,B) and (B,A). A partial ordering $\le$ is defined on the set of all oriented separations of G by letting

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

A *star (of separations)* is a set $\{\,(A^i,B^i)\colon i\in I\,\}$ of oriented separations (A^i,B^i) of G such that $(A^i,B^i)\le(B^j,A^j)$ for every two distinct indices $i,j\in I$. If we write the edge set of T as $\{t_is\colon i\in I\}$, then the induced separations $\{U_1^i,U_2^i\}$ defined by $U_1^i:=V_{t_i}$ and $U_2^i:=\bigcup_{t\in T-t_i}V_t$ form a star of separations $\{\,(U_1^i,U_2^i)\colon i\in I\,\}$. Conversely, every star of separations $\{\,(A^i,B^i)\colon i\in I\,\}$ defines a star-decomposition with leaf parts A^i $(i\in I)$ and central part $\bigcap_{i\in I}B^i$. If this star is $\{\,(U_1^i,U_2^i)\colon i\in I\,\}$, we retrieve $(T,\mathcal{V})$.

The following non-standard notation often will be useful as an alternative perspective on separations of graphs. For a vertex set $X\subseteq V(G)$ we denote the collection of the components of $G-X$ by $\mathscr{C}_X$. If any $X\subseteq V(G)$ and $\mathscr{C}\subseteq\mathscr{C}_X$ are given, then these give rise to a separation of G which we denote by

$$\{X,\mathscr{C}\}:=\big\{\ V\smallsetminus V[\mathscr{C}]\ ,\ X\cup V[\mathscr{C}]\ \big\}$$

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

$$(X,\mathscr{C}):=\big(\ V\smallsetminus V[\mathscr{C}]\ ,\ X\cup V[\mathscr{C}]\ \big)\quad\text{and}\quad(\mathscr{C},X):=\big(\ V[\mathscr{C}]\cup X\ ,\ V\smallsetminus V[\mathscr{C}]\ \big).$$

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

We conclude this section with a paragraph on how tree-decompositions can be used to distinguish ends of graphs. The *order* of a separation is the cardinality of its separator. A finite-order separation $\{A,B\}$ of G *distinguishes* two ends ω_1 and ω_2 of G if $C(A\cap B,\omega_1)\subseteq G[A\smallsetminus B]$ and $C(A\cap B,\omega_2)\subseteq G[B\smallsetminus A]$ (or vice versa). If $\{A,B\}$ distinguishes ω_1 and ω_2 and has minimal order among all the separations of G that distinguish ω_1 and ω_2, then $\{A,B\}$ distinguishes ω_1 and ω_2 *efficiently.* The ends ω_1 and ω_2 are said to be *k-distinguishable* for an integer $k\ge 0$ if there is a separation of G of order at most k that distinguishes ω_1 and ω_2. We say that $(T,\mathcal{V})$ *distinguishes* ω_1 and ω_2 if some edge of T induces a separation of G that distinguishes ω_1 and ω_2; it distinguishes ω_1 and ω_2 *efficiently* if the induced separation can be chosen to distinguish ω_1 and ω_2 efficiently. The following theorem is an immediate consequence of [18, Corollary 6.6] and the construction in [18, Theorem 6.2] that leads to the corollary. (The construction in the proof of Theorem 6.2 ensures that the tree of tree-decompositions obtained in Corollary 6.6

has finite height, and it is straightforward to combine all the tree-decompositions into a single one.)

Theorem 3.2.5. *Every connected graph G has for every number $k \in \mathbb{N}$ a tree-decomposition that efficiently distinguishes all the k-distinguishable ends of G.*

3.3. Countably determined directions and the first axiom of countability

In this section we characterise for every graph G, by unavoidable substructures, both the countably determined directions of G and its directions that are not countably determined.

Given a graph G we call a ray $R \subseteq G$ *topological* in G if the end ω of G that contains R has a countable neighbourhood base $\{\, \Omega(X_n, \omega) \colon n \in \mathbb{N} \,\}$ in $\Omega(G)$ where each vertex set X_n consists of the first n vertices of R. Our first lemma shows that rays are directional if and only if they are topological:

Lemma 3.3.1. *For every graph G and every ray $R \subseteq G$ the following assertions are equivalent:*

(i) *R is directional in G.*
(ii) *R is topological in G.*

Proof. Let us write ω for the end of G that is represented by R, and let us denote by X_n the set of the first n vertices of R.

(ii)⇒(i) By assumption, $\{\, \Omega(X_n, \omega) \colon n \in \mathbb{N} \,\}$ is a countable neighbourhood base for $\omega \in \Omega(G)$. Then f_ω is countably determined by the directional choices $(X_n, f_\omega(X_n))$ because the end space is Hausdorff.

(i)⇒(ii) We claim that $\{\, \Omega(X_n, \omega) \colon n \in \mathbb{N} \,\}$ is a neighbourhood base for $\omega \in \Omega(G)$. Now, suppose for a contradiction that there is a basic open neighbourhood $\Omega(X, \omega)$ of ω in $\Omega(G)$ that contains none of the sets $\Omega(X_n, \omega)$. We recursively construct a sequence of pairwise disjoint rays R_n all having precisely their first vertex on R and belonging to ends not in $\Omega(X, \omega)$. Having these rays at hand will give the desired contradiction; as X is finite one of these rays lies in $C(X, \omega)$ contradicting that its end is not in $\Omega(X, \omega)$.

So suppose we have found $R_0, \ldots, R_{n-1}$. In order to define R_n, choose $k \in \mathbb{N}$ large enough that $(X_k, f_\omega(X_k))$ distinguishes f_ω from all directions induced by the rays $R_0, \ldots, R_{n-1}$ (if $n = 0$ pick $k = 0$). Such a k exists because R is directional. Since the rays $R_0, \ldots, R_{n-1}$ have precisely their first vertex on R and X_k consists of vertices of R, none of the rays $R_0, \ldots, R_{n-1}$ meets the component $f_\omega(X_k) = C(X_k, \omega)$. By our assumption there is an end, η say, that is contained in $\Omega(X_k, \omega)$ but not in $\Omega(X, \omega)$. We choose any ray of η in $C(X_k, \omega)$ having precisely its first vertex on R to be the nth ray R_n. □

Theorem 3.3.2. *For every graph G and every end ω of G the following assertions are equivalent:*

(i) *The end space of G is first countable at ω.*
(ii) *The direction f_ω is countably determined in G.*
(iii) *The end ω is represented by a directional ray.*
(iv) *The end ω is represented by a topological ray.*

This theorem clearly implies Theorem 3.1:

Proof of Theorem 3.1. Theorem 3.3.2 (ii)⇔(iii) is the statement of Theorem 3.1. □

Proof. (i)⇒(ii) Let $\{\, \Omega(X_n, \omega) \colon n \in \mathbb{N} \,\}$ be a countable neighbourhood base of basic open sets for $\omega \in \Omega(G)$. Then f_ω is countably determined by its countably many directional choices $(X_n, f_\omega(X_n))$ because the end space is Hausdorff.

(ii)⇒(iii) Let $\{\, \Omega(X_n, f_\omega(X_n)) \colon n \in \mathbb{N} \,\}$ be a countable set of directional choices that distinguish f_ω from every other direction. Fix any ray $R \in \omega$ and denote by U the union of $V(R)$ and all the X_n. By Lemma 3.2.4 there is a normal tree $T \subseteq G$ that contains U cofinally. As $V(R) \subseteq T$ we have that $\omega \in \partial_\Omega T$. Now, by Lemma 3.2.3, there is a normal ray R_ω in T belonging to ω.

We claim that R_ω is directional in G. For this, it suffices to show that for every other end $\eta \neq \omega$ of G there is a finite initial segment of R_ω separating ω and η. By assumption, there is $n \in \mathbb{N}$ such that X_n separates ω and η.

Let v be any vertex of the ray $R_\omega - \lceil X_n \rceil$ where the down-closure is taken in T. Since T is normal in G, we have $C(\lceil v \rceil, \omega) \subseteq C(X_n, \omega)$ by Lemma 3.2.2. In particular, the initial segment $\lceil v \rceil$ of R_ω separates ω from η.

(iii)⇒(iv) This is Lemma 3.3.1 (i)⇒(ii).

(iv)⇒(i) This holds by the definition of a topological ray. □

Our second main result, the characterisation by unavoidable substructures of the directions of any given graph that are not countably determined in that graph, needs some preparation.

Definition 3.3.3 (Generalised paths)**.** For a graph G a *generalised path* in G with *endpoints* $\omega_1 \neq \omega_2 \in \Omega(G)$ is an ordered pair $(P, \{\omega_1, \omega_2\})$ where $P \subseteq G$ is one of the following:

- a double ray with one tail in the end ω_1 and another tail in the end ω_2;
- a finite path $v_0 \dots v_k$ such that v_0 dominates the end ω_1 and v_k dominates the end ω_2;
- a ray in ω_1 whose first vertex dominates ω_2.

Two generalised paths (P, Ψ) and (P', Ψ') are *vertex-disjoint* if P and P' are disjoint. Two generalised paths (P, Ψ) and (P', Ψ') are *disjoint* if they are vertex-disjoint and $\Psi \cap \Psi' = \emptyset$.

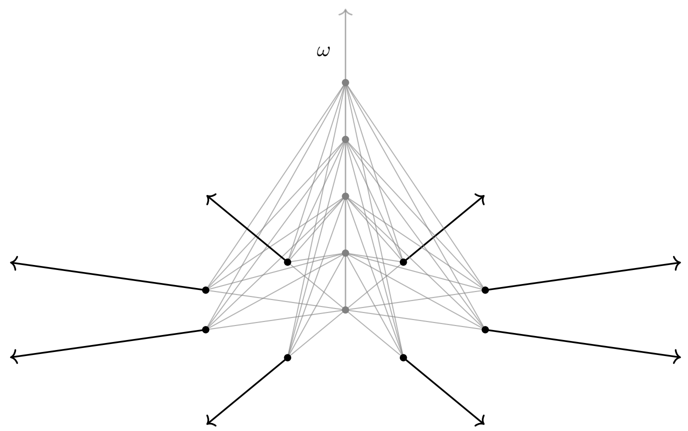

Figure 3.3.1.: The black rays form a sun centred at ω

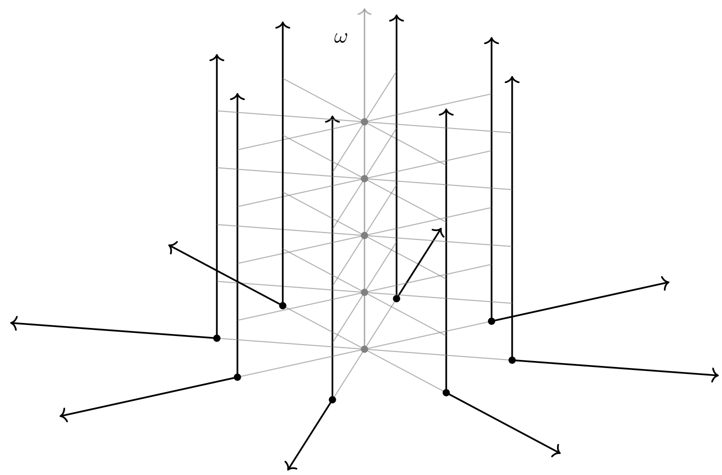

Figure 3.3.2.: The black double rays form a sun centred at ω

Definition 3.3.4 (Generalised star and sun)**.** For a graph G a *generalised star* in G with *centre* $\omega \in \Omega(G)$ is a collection of pairwise vertex-disjoint generalised paths $\{\, (P^i, \{\omega, \omega^i\}) : i \in I \,\}$ such that each end ω^i is distinct from all other ends ω^j with $j \neq i \in I$. Then the ends ω^i with $i \in I$ are the *leaves* of the generalised star.

A generalised star $\{\, (P^i, \{\omega, \omega^i\}) : i \in I \,\}$ is *proper* if either every path P^i is a double ray or every path P^i is a ray in ω^i whose first vertex dominates ω. A proper generalised star in G with centre ω is also called a *sun* in G with *centre* ω.

In Figures 3.3.1 and 3.3.2 we see two examples of suns of size eight centred at an end ω. If we increase their size from eight to $\aleph_1$ in the obvious way, then the direction f_ω is no longer countably determined.

Theorem 3.3.5. *For every graph G and every end ω of G the following assertions are equivalent:*

(i) *The end space of G is not first countable at ω.*
(ii) *There is an uncountable sun in G centred at ω.*

Moreover, if (ii) *holds, then we find an uncountable sun* $\{ (P^i, \{\omega, \omega^i\}) \mid i \in I \}$ *in* G *and an uncountable star-decomposition* $(T, \mathcal{V})$ *of* G *of finite adhesion such that* ω *lives in the central part and every* ω^i *lives in its own leaf part.*

This theorem clearly implies Theorem 3.2:

Proof of Theorem 3.2. Combine Theorem 3.3.5 with Theorem 3.3.2 (i)⇔(ii). □

Proof of Theorem 3.3.5. (ii)⇒(i) Suppose for a contradiction that (ii) and ¬(i) hold. Let $\{(P^i, \{\omega, \omega^i\}) \colon i \in I\}$ be an uncountable sun in G centred at ω, and let $\{ \Omega(X_n, \omega) \colon n \in \mathbb{N} \}$ be a countable open neighbourhood base for ω in $\Omega(G)$. As all the vertex sets X_n are finite and the P^i are pairwise disjoint, there is an index $i \in I$ such that P^i misses all of the vertex sets X_n. Hence, the leaf ω^i is contained in all of the neighbourhoods $\Omega(X_n, \omega)$ contradicting that $\Omega(G)$ is Hausdorff and that $\{ \Omega(X_n, \omega) \colon n \in \mathbb{N} \}$ is a neighbourhood base for $\omega \in \Omega(G)$.

(i)⇒(ii) Suppose that the end space $\Omega(G)$ is not first countable at ω. Our aim is to construct an uncountable sun in G centred at ω. Let Δ be the set of vertices that dominate ω. By Zorn's lemma there is an inclusionwise maximal set $\mathcal{R}$ of pairwise disjoint rays all belonging to ω. Denote by $V[\mathcal{R}]$ the union of all the vertex sets of the rays contained in $\mathcal{R}$, and let $A := \Delta \cup V[\mathcal{R}]$. Then $\partial_\Omega A = \{\omega\}$. By Zorn's lemma there is an inclusionwise maximal set $\mathcal{P}$ of pairwise disjoint rays all starting at A and belonging to ends of G other than ω.

First, note that $\mathcal{P}$ yields the desired sun if $\mathcal{P}$ is uncountable: Since $\partial_\Omega A = \{\omega\}$ only finitely many rays in $\mathcal{P}$ belong to the same end and every ray in $\mathcal{P}$ has a tail avoiding A. Hence, if $\mathcal{P}$ is uncountable, we pass to an uncountable subset $\mathcal{P}' \subseteq \mathcal{P}$ such that all rays in $\mathcal{P}'$ belong to pairwise distinct ends and have precisely their first vertex in A. If uncountably many rays in $\mathcal{P}'$ start at Δ we are done. So we may assume that uncountably many rays in $\mathcal{P}'$ start at $V[\mathcal{R}]$. As only countably many rays in $\mathcal{P}'$ start at the same ray in $\mathcal{R}$, we may pass to an uncountable subset $\mathcal{P}'' \subseteq \mathcal{P}'$ such that all rays in $\mathcal{P}''$ start at distinct rays in $\mathcal{R}$. Extending every ray in $\mathcal{P}''$ by a tail of the unique ray in $\mathcal{R}$ it hits yields again the desired uncountable sun.

Therefore, we may assume that $\mathcal{P}$ is countable. In the remainder of this proof, we show that this is impossible: we deduce that then there is a countable neighbourhood base for ω in $\Omega(G)$, contradicting our assumption.

We claim that if $\mathcal{P}$ is finite, then ω is an isolated point in $\Omega(G)$, that is, there is a finite vertex set separating ω from all other ends of G simultaneously. Let $\omega_0, \ldots, \omega_n$ be the ends of the rays in $\mathcal{P}$. As these are only finitely many, there is a finite vertex set $X \subseteq V(G)$ separating ω from all of the ω_i simultaneously. Now, $C(X, \omega)$ contains only finitely many vertices of the rays in $\mathcal{P}$; by possibly extending X we may assume that $C(X, \omega)$ contains no vertex from any ray in $\mathcal{P}$. Then no end of G other than ω lies in $\Omega(X, \omega)$, because any such end has a ray in $C(X, \omega)$

starting at A and avoiding all rays in $\mathcal{P}$, contradicting the maximality of $\mathcal{P}$.

Thus, $\mathcal{P}$ must be countably infinite. Then the vertex set $V[\mathcal{P}] = \bigcup_{R \in \mathcal{P}} V(R)$ is countable as well. By Lemma 3.2.4 there is a normal tree $T \subseteq G$ that contains $V[\mathcal{P}]$ cofinally. Moreover, as $\mathcal{P}$ is infinite we have $\omega \in \partial_\Omega T$, and so there is a normal ray $R_\omega \subseteq T$ belonging to ω by Lemma 3.2.3.

We claim that for any end $\eta \neq \omega$ of G there is a finite initial segment of R_ω separating ω from η in G. This suffices to derive the desired contradiction, because then Lemma 3.3.1 shows that the finite initial segments of R_ω define a countable open neighbourhood base for ω in $\Omega(G)$.

First, suppose $\eta \in \partial_\Omega T$. Then η has a normal ray R_η in T by Lemma 3.2.3. As T is normal in G, the initial segment $R_\omega \cap R_\eta$ of R_ω separates ω from η in G.

Second, suppose $\eta \notin \partial_\Omega T$. Then there is a unique component C of $G - T$ that contains a tail of every ray in η. The neighbourhood $N(C)$ of C in T is a chain. If the neighbourhood $N(C)$ of C is not cofinal in R_ω, then any finite initial segment of R_ω containing $N(C) \cap R_\omega$ separates ω from η in G. So suppose that $N(C)$ is cofinal in R_ω and denote by U the set of all the vertices in C having a neighbour in R_ω.

If some vertex $u \in U$ sends infinitely many edges to T, then u dominates ω; in particular, there is a generalised path $(P, \{\eta, \omega\})$ in G where P is a ray in η that is contained in C and starts at $u \in \Delta \subseteq A$, contradicting the maximality of $\mathcal{P}$. Therefore, we may assume that every vertex in U sends only finitely many edges to T; in particular, U is infinite. Thus, we find an independent set M of infinitely many U–T edges in G; we denote by U' the set of endvertices that these edges have in U. Then we apply the star-comb lemma (3.2.1) to U' in C. If this yields a star attached to U', then the centre of the star dominates ω and we obtain the same contradiction as above. Otherwise this yields a comb attached to U'. Then its spine, R say, is a ray belonging to ω. Thus, by the maximality of $\mathcal{R}$ there is a vertex of $V[\mathcal{R}]$ on R and in particular in C. Consequently, there is a ray in η that is contained in C and starts at $V[\mathcal{R}] \subseteq A$, contradicting the maximality of $\mathcal{P}$. This completes the proof of (i)⇒(ii).

Finally, we prove the 'moreover' part. Let $S = \{\, (P^i, \{\omega, \omega^i\}) \colon i \in I \,\}$ be any uncountable sun in G with centre ω. We say that an oriented finite-order separation (A, B) of G is *S-separating* if (A, B) separates the centre of S from some leaf ω^i of S in that $C(A \cap B, \omega) \subseteq G[B \setminus A]$ while $C(A \cap B, \omega^i) \subseteq G[A \setminus B]$. Oriented separations of the form $(C, N(C))$ with $C = C(X, \omega')$ for some finite vertex set $X \in \mathcal{X}$ and an end ω' of G are called *golden*. A star σ of finite-order separations is *golden* if every separation in σ is golden.

Consider the set Σ of all the golden stars that are formed by S-separating separations of G, partially ordered by inclusion, and apply Zorn's lemma to $(\Sigma, \subseteq)$ to obtain a maximal element $\sigma \in \Sigma$. We claim that σ must be uncountable, and assume for a contradiction that σ is countable. Let us write U for the union of the separators of the separations in σ. As U is countable, some path P^j avoids U. We consider the two cases that the end ω^j lies in the closure of U or not.

First, suppose that the end ω^j does not lie in the closure of U. It is straightforward to find an S-separating golden separation (C, X) with ω^j living in C and C avoiding U. Note that (C, X) is not contained in σ. We claim that $\sigma' := \sigma \cup \{\, (C, X) \,\}$ is again a star contained in Σ. Since all the elements of σ' are S-separating and golden, it remains to show that the separations in σ' indeed form a star. As $\sigma \subseteq \sigma'$ already is a star, it suffices to show $(C, X) \le (Y, D)$ for all separations $(D, Y) \in \sigma$. For this, let any separation $(D, Y) \in \sigma$ be given. To establish $(C, X) \le (Y, D)$ it suffices to show that C avoids $Y \cup D$, because X is the neighbourhood of C and Y is the neighbourhood of D. The component C avoids Y because it avoids U which contains Y as a subset. Therefore, the component C is contained in some component of $G-Y$. Now suppose for a contradiction that C and D meet. Then C must be contained in D. Since P^j avoids Y and contains a ray that lies C, we deduce $P^j \subseteq D$. But then ω must live in D, contradicting that (D,Y) is S-separating. Thus, σ' is again an element of Σ, contradicting the maximal choice of σ.

Second, suppose that the end ω^j lies in the closure of U. We show that this implies $\omega^j = \omega$, a contradiction. For this, let any finite vertex set $X \subseteq V(G)$ be given; we show $C(X,\omega^j) = C(X,\omega)$. To get started, we observe that all but finitely many of the paths P^i avoid X. Also, all but finitely many of the separations $(D,Y) \in \sigma$ have their component D avoid X (the components are disjoint for distinct separations in σ because σ is a star of separations). As U meets $C(X,\omega^j)$ infinitely, this allows us to find a separation $(C(Y,\omega^i),Y) \in \sigma$ that has its separator Y meet the component $C(X,\omega^j)$ while both $C(Y,\omega^i)$ and P^i avoid the finite separator X. To show $C(X,\omega^j) = C(X,\omega)$, it suffices to find a P^i–P^j path in G that avoids X. We find such a path as follows: In $C(Y,\omega^i)$ we find a path from P^i to a vertex that sends an edge to a vertex v in the non-empty intersection $Y \cap C(X,\omega^j)$. And in $C(X,\omega^j)$ we find a path from P^j to v. Then both paths avoid X, and their union contains the desired P^i–P^j path avoiding X. This concludes the proof that σ is uncountable.

Now σ is an uncountable star of separations $(C,N(C))$ such that some leaf $\omega^{i(C)}$ of the sun S lives in C and ω does not live in C. Then $J := \{\, i(C) \colon (C,N(C)) \in \sigma \,\}$ is an uncountable subset of I. In particular, the uncountable sun

$$\{\, (P^j, \{\omega,\omega^j\}) \colon j \in J \,\}$$

and the star-decomposition arising from σ are as desired. □

For the interested reader we remark that, even though end spaces of graphs are in general not first countable, it is straightforward to show that every end space is strong Fréchet–Urysohn (which is a generalisation of the first axiom of countability): A topological space X is called a *strong Fréchet–Urysohn space* if for any sequence of subsets $A_0, A_1, \ldots$ of X and every $x \in \bigcap_{n\in\mathbb{N}} \overline{A_n}$ there is a sequence of points $x_0, x_1, \ldots$ converging to x such that $x_n \in A_n$ for all $n \in \mathbb{N}$.

Lemma 3.3.6. *End spaces of graphs are strong Fréchet–Urysohn.* □

3.4. Countably determined graphs and the second axiom of countability

In this section we structurally characterise both the graphs that are countably determined and the graphs that are not countably determined.

Clearly, a graph G is countably determined if and only if every component of G is countably determined and only countably many components of G have directions. Similarly, the end space of a graph G is second countable if and only if every component of G has a second countable end space and only countably many components of G have ends. Thus, to structurally characterise the countably determined graphs and the graphs that are not countably determined, and to link this to whether or not the end space is second countable, it suffices to consider only connected graphs.

The local property that every direction of G is countably determined in G does not imply the stronger global property that G is countably determined:

Example 3.4.1. There exists a connected graph G all whose directions are countably determined in it but which is itself not countably determined. The graph G can be chosen so that its end space is compact and first countable at every end, but neither metrisable nor second countable nor separable.

Recall that a topological space is called *separable* if it admits a countable dense subset. Every second countable space is separable, but the converse is generally false. For end spaces, however, we shall see in Theorem 3.4.2 that the converse is true: the end space of any graph is second countable if and only if it is separable.

Proof of Example 3.4.1. Let T_2 be the rooted infinite binary tree. The graph G arises from T_2 by disjointly adding a new ray R' for every rooted ray $R \subseteq T_2$ such that R'_1 and R'_2 are disjoint for distinct rooted rays $R_1, R_2 \subseteq T_2$, and joining the first vertex $v_{R'}$ of each R' to all the vertices of R. Then for every rooted ray $R \subseteq T_2$ the two rays R' and $v_{R'}R$ are directional in G. Since every direction of G is induced by precisely one of these directional rays, all the directions of G are countably determined.

The graph G, however, is not countably determined: If $\{ (X_n, C_n) \colon n \in \mathbb{N} \}$ is any countable collection of directional choices in G, then there is a rooted ray $R \subseteq T_2$ such that R' avoids all X_n (because T_2 contains continuum many rooted rays and $\bigcup_n X_n$ is countable). But then, for all $n \in \mathbb{N}$, the subgraph $G - X_n$ contains a double ray formed by R' and a subray of R avoiding X_n (that is connected to R' by one of the infinitely many $v_{R'}$–R edges). These double rays then witness that none of the directional choices (X_n, C_n) distinguishes the direction induced by R from the direction induced by R' or vice versa.

The end space of G is first countable at every end because every direction of G is countably determined (Theorem 3.3.2). It is compact because the deletion of any finite set of vertices of G leaves only finitely many components, cf. [21, Theorem 4.1] or [12, Lemma 2.8]. However, the end space of G is not separable, because every dense subset of $\Omega(G)$ must contain all the continuum many ends represented by the rays R'. Thus, it its neither second countable nor metrizable. □

This is essentially a combinatorially constructed version of the Alexandroff double circle [33, Example 3.1.26].

Now we structurally characterise the countably determined graphs and the graphs that are not countably determined, and structurally characterise the graphs whose end spaces are second countable or not. Our introduction suggests that this is the order in which we prove these results, but we will prove them in a different order: First, we shall structurally characterise the graphs whose end spaces are second countable or not. Then, we shall prove that the end space of a graph is second countable if and only if the graph is countably determined. Finally, we shall use this equivalence to immediately obtain structural characterisations of the countably determined graphs and the graphs that are not countably determined. Here, then is our structural characterisation of the graphs whose end spaces are second countable:

Theorem 3.4.2. *For every connected graph G the following assertions are equivalent:*

(i) *The end space of G is second countable.*
(ii) *The end space of G is separable.*
(iii) *There is a countable end-faithful normal tree $T \subseteq G$.*
(iv) *The end space of G has a countable base that consists of basic open sets.*

Proof. (i)⇒(ii) Every second countable space is separable.

(ii)⇒(iii) Let $\Psi \subseteq \Omega(G)$ be any countable and dense subset. Pick a ray $R_\omega \in \omega$ for every end $\omega \in \Psi$ and let $U := \bigcup \{ V(R_\omega) \colon \omega \in \Psi \}$. By Lemma 3.2.4 there is a countable normal tree $T \subseteq G$ that contains U cofinally. We have to show that T is end-faithful. As mentioned in Section 3.2.2, normal trees are always end-injective. To show that T is end-surjective note that $\partial_\Omega T$ is closed in $\Omega(G)$. Hence we have $\Omega(G) = \overline{\Psi} \subseteq \overline{\partial_\Omega T} = \partial_\Omega T$. So by Lemma 3.2.3 the normal tree T contains a normal ray of every end of G.

(iii)⇒(iv) Let $T \subseteq G$ be any countable end-faithful normal tree. We claim that the collection $\mathcal{B} := \{ \Omega(\lceil t \rceil, \omega) \colon t \in T,\, \omega \in \Omega \}$ is a countable base of the topology on $\Omega(G)$. Note first that $\mathcal{B}$ is indeed countable: Since T is end-faithful and countable, the deletion of finitely many vertices of T from G results in only countably many components containing an end. Consequently, for every $t \in T$ there are only countably many distinct sets of the form $\Omega(\lceil t \rceil, \omega)$.

Now, given a basic open set of $\Omega(G)$, say $\Omega(X, \omega)$, our goal is to find a vertex $t \in T$ such that $\Omega(\lceil t \rceil, \omega) \subseteq \Omega(X, \omega)$. By Lemma 3.2.3, every end η of G in the closure of T contains a normal ray $R_\eta \subseteq T$. By the normality of T and Lemma 3.2.2, every end $\eta \neq \omega$ of G is separated from ω in G by the finite initial segment $R_\eta \cap R_\omega$ of R_ω. In particular, R_ω is directional in G. Hence, by the implication (i)⇒(ii) of Lemma 3.3.1 the ray R_ω is topological. Thus, there is a vertex $t \in R_\omega$ such that $\Omega(\lceil t \rceil, \omega) \subseteq \Omega(X, \omega)$ holds.

(iv)⇒(i) This is clear. □

Next, we structurally characterise the graphs whose end spaces are not second countable. The characterising structure is not the star-decomposition in Theorem 3.4 that one would expect; that is because this result is an auxiliary result that we will use in a second step to prove a second structural characterisation, Theorem 3.4.4, which then is phrased in terms of the desired star-decomposition.

Theorem 3.4.3. *For every connected graph G the following assertions are equivalent:*

(i) *The end space of G is not second countable.*
(ii) *The graph G contains either*
   - *an uncountable sun,*
   - *uncountably many disjoint generalised paths, or*
   - *a finite vertex set that separates uncountably many ends of G simultaneously.*

Proof. Recall that, by Theorem 3.4.2, the end space of G is second countable if and only if it has a countable base that consists of basic open sets. Then clearly (ii)⇒(i).

(i)⇒(ii) For this, suppose that G is given such that the end space of G is not second countable. We have to find one of the three substructures for G listed in (ii). By Zorn's lemma we find an inclusionwise maximal collection $\mathcal{P}$ of pairwise vertex-disjoint generalised paths in G. Our proof consists of two halves. In the first half we show that if $\mathcal{P}$ is uncountable, then we find either an uncountable sun in G or uncountably many disjoint generalised paths in G. In the second half we show that if $\mathcal{P}$ is countable, then we find a finite vertex set of G that separates uncountably many ends of G simultaneously.

First, we assume that $\mathcal{P}$ is uncountable. In this case, we consider the auxiliary multigraph that is defined on the set of ends of G by declaring every generalised path $(P, \{\omega_1, \omega_2\}) \in \mathcal{P}$ to be an edge between ω_1 and ω_2. Note that the auxiliary multigraph contains only finitely many parallel edges between any two vertices. Thus, by replacing $\mathcal{P}$ with a suitable uncountable subset we may assume that the auxiliary multigraph is in fact a graph.

If that auxiliary graph has a vertex ω of uncountable degree, then its incident edges correspond to uncountably many generalised paths that form an uncountable generalised star in G with centre ω. This uncountable generalised star need not be proper in general. However, it shows that ω has no countable neighbourhood base in $\Omega(G)$, so Theorem 3.3.5 yields an uncountable sun in G with centre ω.

Otherwise, every vertex of the auxiliary graph has countable degree. Then we greedily find an uncountable independent edge set, and this edge set corresponds to an uncountable collection of disjoint generalised paths in G.

Second, we assume that $\mathcal{P}$ is countable. Then our goal is to find a finite vertex set $X \subseteq V(G)$ that separates uncountably many ends of G simultaneously. By Lemma 3.2.4 there is a countable normal tree $T \subseteq G$ that cofinally contains the union of the vertex sets of the generalised paths in $\mathcal{P}$. Then $\Omega(G) \setminus \partial_\Omega T$ is uncountable, since otherwise applying Lemma 3.2.4 to the union of $V(T)$ with the vertex set of a ray from every end in $\Omega(G) \setminus \partial_\Omega T$ gives a countable end-faithful normal tree in G, contradicting Theorem 3.4.2.

Every ray from an end in $\Omega(G) \setminus \partial_\Omega T$ has a tail in one of the components of $G - T$ and this component is the same for any two rays in the same end. We say that an end in $\omega \in \Omega(G) \setminus \partial_\Omega T$ *lives* in the unique component of $G - T$ in which every ray in ω has a tail. By the maximality of $\mathcal{P}$, distinct ends in $\Omega(G) \setminus \partial_\Omega T$ live in distinct components of $G - T$. As $\Omega(G) \setminus \partial_\Omega T$ is uncountable, we conclude that there are uncountably many components of $G - T$ in which an end of $\Omega(G) \setminus \partial_\Omega T$ lives; we call these components *good*.

We claim that every good component of $G - T$ has finite neighbourhood. For this, assume for a contradiction that there is a good component C of $G - T$ whose neighbourhood $N(C) \subseteq T$ is infinite. Write ω for the end in $\Omega(G) \setminus \partial_\Omega T$ that lives in C. The down-closure of $N(C)$ in T forms a ray and we denote by η the end in $\partial_\Omega T$ represented by this ray. Consider the set U of all the vertices in C sending an edge to T. If some vertex $u \in U$ sends infinitely many edges to T, then u dominates η; in particular, there is a generalised path $(P, \{\omega, \eta\})$ in G where P is a ray in ω that is contained in C and starts at u, contradicting the maximality of $\mathcal{P}$. Therefore, we may assume that every vertex in U sends only finitely many edges to T; in particular, U is infinite. Thus, we find an independent set M of infinitely many U–T edges in G; we denote by U' the set of the endvertices that these edges have in U. Applying the star-comb lemma 3.2.1 in C to U' gives either a star attached to U' or a comb attached to U'. The centre of a star attached to U' would dominate η, yielding the same contradiction that would be caused by a vertex in U sending infinitely many edges to T. Hence we obtain a comb attached to U'. The comb's spine represents η, because of the edges in M. Consequently, there is a double ray $P \subseteq C$ defining a generalised path $(P, \{\omega, \eta\})$ vertex-disjoint from all generalised paths in $\mathcal{P}$, contradicting the maximality of $\mathcal{P}$. This completes the proof of the claim that every good component of $G - T$ has finite neighbourhood.

Finally, as all of the uncountably many good components of $G - T$ have a finite neighbourhood in T and T is countable, there are uncountably many such components having the same finite neighbourhood $X \subseteq V(T)$. Then X is a finite vertex set of G that separates uncountably many ends of G simultaneously, as desired. □

Next, we will prove the structural characterisation of the graphs whose end spaces are not second countable, in terms of the desired star-decomposition:

Theorem 3.4.4. *For every connected graph G the following assertions are equivalent:*

(i) *The end space of G is not second countable.*
(ii) *G has an uncountable star-decomposition of finite adhesion such that in every leaf part there lives an end of G. In particular, the end space of G contains uncountably many pairwise disjoint open sets.*

Curiously, the star-decomposition in (ii) cannot be replaced with a star of pairwise inequivalent rays:

Example 3.4.5. There is a graph G with second countable end space that contains uncountably many rays meeting precisely in their first vertex and representing distinct ends of G.

Proof. We start with T_2, the rooted binary tree. Every end $\omega \in \Omega(T_2)$ is represented by a ray R_ω starting at the root. We obtain the graph H from T_2 by adding, for every end $\omega \in \Omega(T_2)$, a new ray R'_ω and joining the nth vertex of R'_ω to the nth vertex of R_ω. Note that the natural inclusion $\Omega(T_2) \subseteq \Omega(H)$ is a homeomorphism, so $\Omega(H)$ is second countable. Finally, we obtain the graph G from H by adding a single new vertex v^* that we join to the first vertex of every ray R'_ω. The addition of v^* did not affect the end space. However, now v^* together with its incident edges and all the rays R'_ω gives the desired substructure. □

We need the next lemma for the proof of Theorem 3.4.4. Recall that oriented separations of the form $(C, N(C))$ with $C = C(X, \omega)$ for some finite vertex set $X \in \mathcal{X}$ and an end ω of G are called *golden*. A star σ of finite-order separations is *golden* if every separation in σ is golden.

Lemma 3.4.6. *Let G be any connected graph. If there exist uncountably many pairwise vertex-disjoint generalised paths in G, then G admits an uncountable golden star of separations.*

Proof. Let $\{\, (P^i, \{\omega_1^i, \omega_2^i\}) : i \in I \,\}$ be any uncountable collection of pairwise vertex-disjoint generalised paths in G. By the pigeonhole principle there exists a number $k \in \mathbb{N}$ and an uncountable subset $J \subseteq I$ such that for all $j \in J$ the ends ω_1^j and ω_2^j are k-distinguishable. Without loss of generality $J = I$. By Theorem 3.2.5 we find a tree-decomposition $(T, \mathcal{V})$ of G that efficiently distinguishes all the k-distinguishable ends of G.

Fix an arbitrary root $r \in T$ and write F for the collection of all the edges $e \in T$ whose induced separation distinguishes two ends ω_1^i and ω_2^i. Then let $T' \subsetneq T$ be the subtree that is induced by the down-closure of the endvertices of the edges in F in the rooted tree T. If T' has a vertex t of uncountable degree, then we find an uncountable subset $\Psi \subseteq \{\, \omega_1^i, \omega_2^i \colon i \in I \,\}$ such that every end $\omega \in \Psi$ lives in its own component C_ω of $G - V_t$ with finite neighbourhood; in particular, $\{\, (C_\omega, N(C_\omega)) \colon \omega \in \Psi \,\}$ is the desired uncountable golden star of separations. We claim that T' must have a vertex of uncountable degree. Otherwise, T' is countable. Then the union U of the separators of the separations induced by the edges of $T' \subsetneq T$ is a countable vertex set. In order to obtain a contradiction note that every P^i meets U in at least one vertex and these vertices are distinct for distinct P^i. □

Proof of Theorem 3.4.4. Recall that, by Theorem 3.4.2, the end space of G is second countable if and only if it has a countable base that consists of basic open sets. Then clearly (ii)⇒(i).

(i)⇒(ii) Suppose that the end space of G is not second countable. We are done if there is a finite vertex set separating uncountably many ends of G simultaneously. Therefore, we may assume by Theorem 3.4.3 that either there is an uncountable sun in G or G contains uncountably many disjoint generalised paths. In either case we are done by Lemma 3.4.6. □

Theorems 3.4.2 and 3.4.4 structurally characterise the graphs whose end spaces are second countable or not, by the structures in terms of which Theorems 3.3 and 3.4 are phrased. The next theorem allows us to deduce Theorems 3.3 and 3.4 from Theorems 3.4.2 and 3.4.4 immediately.

Theorem 3.4.7. *The end space of a graph is second countable if and only if the graph is countably determined.*

Proof. Let G be any graph. For the forward implication, suppose that the end space of G is second countable. Then, by Theorem 3.4.2, the end space of G has a countable base $\{\, \Omega(X_n, C_n) \colon n \in \mathbb{N} \,\}$ that consists of basic open sets $\Omega(X_n, C_n) \subseteq \Omega(G)$. We claim that the directional choices (X_n, C_n) distinguish every two directions of G from each other. For this, let any two distinct directions f, h of G be given. Since the end space of G is Hausdorff, there is a number $n \in \mathbb{N}$ such that $\Omega(X_n, C_n)$ contains ω_f but not ω_h; in particular, $f(X_n) = C_n$ and $h(X_n) \neq C_n$ as desired.

For the backward implication suppose that G is countably determined, and let $\{\, (X_n, C_n) \colon n \in \mathbb{N} \,\}$ be any countable set of directional choices (X_n, C_n) in G that distinguish every two directions of G from each other. Let us assume for a contradiction that the end space of G is not second countable. Then, by Theorem 3.4.3, the graph G contains either

- an uncountable sun,
- uncountably many disjoint generalised paths, or
- a finite vertex set that separates uncountably many ends of G simultaneously.

If G contains an uncountable sun or uncountably many disjoint generalised paths, then in either case G contains a generalised path $(P, \{\omega_1, \omega_2\})$ such that P avoids all the countably many finite vertex sets X_n. But then no directional choice (X_n, C_n) distinguishes f_{ω_1} from f_{ω_2} or vice versa, a contradiction. Thus, there must be a finite vertex set $X \subseteq V(G)$ that separates uncountably many ends ω_i $(i \in I)$ of G simultaneously. We abbreviate $C(X, \omega_i)$ as D_i. By the pigeonhole principle we may assume that $X = N(D_i)$ for all $i \in I$.

Let us consider the subset $N \subseteq \mathbb{N}$ of all indices $n \in \mathbb{N}$ whose directional choice (X_n, C_n) distinguishes some f_{ω_i} from some f_{ω_j}. Every component C_n with $n \in N$ meets some component D_i because some end ω_i lives in C_n. Then, for all $n \in N$, either C_n is contained in some component D_i entirely, or C_n meets X and contains all of the components D_i except possibly for the finitely many components D_i that meet X_n. This means that every directional choice (X_n, C_n) with $n \in N$ either distinguishes finitely many directions f_{ω_i} from uncountably many directions f_{ω_j} or vice versa. Thus, for every $n \in N$ there is a cofinite subset $I_n \subseteq I$ such that no f_{ω_i} with $i \in I_n$ is distinguished by (X_n, C_n) from any other f_{ω_j} with $j \in I_n$. But then the uncountably many directions f_{ω_i} with $i \in \bigcap_{n \in N} I_n$ are not distinguished from each other by any directional choices (X_n, C_n), a contradiction. □

Proof of Theorem 3.3. Theorem 3.4.2 and Theorem 3.4.7 together imply Theorem 3.3. □

Proof of Theorem 3.4. Theorem 3.4.4 and Theorem 3.4.7 together imply Theorem 3.4. □

3.5. First and second countability for $|G|$

In this section, we employ our results to characterise when the spaces $|G|$ formed by a graph and its ends are first countable or second countable.

First, we describe a common way to extend the topology on $\Omega(G)$ to a topology on $|G| = G \cup \Omega(G)$. The 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 two points of the graph is the length of a shortest arc in G between them), together with basic open neighbourhoods for ends of the form

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

where $\mathring{E}_\varepsilon(X, C(X, \omega))$ denotes the open ball around $C(X, \omega)$ in G of radius ε. Polat observed that the subspace $V(G) \cup \Omega(G)$ is homeomorphic to $\Omega(G^+)$, where G^+ denotes the graph obtained from G by gluing a new ray R_v onto each vertex v of G so that R_v meets G precisely in its first vertex v and R_v is distinct from all other $R_{v'}$, cf. [60, §4.16].

Note that $|G|$ with MTop is first countable at every vertex of G and at inner points of edges.

Lemma 3.5.1. *For every graph G and every end ω in the space $|G|$ with* MTop*, the following assertions are equivalent:*

(i) *$|G|$ is first countable at ω.*
(ii) *The end ω is represented by a ray R such that every component of $G - R$ has finite neighbourhood.*

Note that every ray as in (ii) is directional and contains all the vertices that dominate the end.

Proof. (i)⇒(ii) If $|G|$ is first countable at ω, then $\Omega(G^+)$ is also first countable at ω. So, by Theorem 3.3.2, the end ω of G considered as an end of G^+ is represented by a directional ray R in G^+. We may assume that $R \subseteq G$. Now, if there is a component C of $G - R$ such that $N(C)$ is infinite, then C is contained in $C(X, \omega)$ for every finite initial segment X of R. Consider a vertex $v \in C$. The end ω_v of G^+ represented by R_v is contained in all of the sets $\Omega_{G^+}(X, \omega)$, contradicting the fact that R is directional.

(ii)⇒(i) Let $R \in \omega$ be as in (ii). First note that for every vertex $v \in V(G)$ there is a finite initial segment X of R that separates v from ω in that $v \notin C(X, \omega)$. Now, if $X' \subseteq V(G)$ is a finite vertex set, chose a finite initial segment X of R that separates every vertex in X' from ω. Then $C(X, \omega) \subseteq C(X', \omega)$. Hence the sets $\hat{C}_{\frac{1}{n}}(X, \omega)$, with $n \in \mathbb{N}$ and X a finite initial segment of R, form a countable neighbourhood base for ω. □

An end is called *fat* if there are uncountably many disjoint rays that represent the end.

Lemma 3.5.2. *For every graph G and every end ω in the space $|G|$ with* MTop*, the following assertions are equivalent:*

(i) *$|G|$ is not first countable at ω.*
(ii) *The end ω is fat or dominated by uncountably many vertices.*

Proof. (i)⇒(ii) If $|G|$ is not first countable at ω, then ω considered as end of G^+ has no countable neighbourhood base in $\Omega(G^+)$ either. By Theorem 3.3.5 there is an uncountable sun in G^+, which gives either uncountably many disjoint rays in G that represent ω or uncountably many vertices in G that dominate ω.

(ii)⇒(i) Let ω be an end of G. If there are uncountably many disjoint rays that represent ω or uncountably many vertices that dominate ω, then any countable collection $\{\hat{C}_{\varepsilon_n}(X_n, \omega) \mid n \in \mathbb{N}\}$ fails to separate ω from a vertex on one of these rays or from a dominating vertex. □

Second countability for $|G|$ is simply decided by the order of G:

Lemma 3.5.3. *For every connected graph G and the space $|G|$ with* MTop *the following assertions are equivalent:*

(i) *$|G|$ is second countable.*
(ii) *$V(G)$ is countable.*

Proof. (i)⇒(ii) If G has uncountably many vertices, then the uniform stars with radius $\frac{1}{2}$ around every vertex form an uncountable collection of disjoint open sets. Hence, $|G|$ is not second countable.

(ii)⇒(i) If $V(G)$ is countable, we take an enumeration $V(G) = \{v_1, v_2, \ldots\}$ of the vertex set of G and write X_n for the set of the first n vertices. For every vertex v the uniform stars with radius $\frac{1}{n}$ around v form a countable neighbourhood base. There are only countably many edges and each is homeomorphic to the unit interval and therefore has a countable base of its topology. Finally, all the sets of the form $\hat{C}_{\frac{1}{n}}(X_n, \omega)$ for all $n \in \mathbb{N}$ and all $\omega \in \Omega(G)$ give a countable neighbourhood base for all the ends of G. Note that $G - X_n$ has only countably many components, for all $n \in \mathbb{N}$; therefore there are indeed only countably many sets of the form $\hat{C}_{\frac{1}{n}}(X_n, \omega)$. □

Part II.

The degree of an end

4. A strengthening of Halin's grid theorem

4.1. Introduction

An *end* of a graph G is an equivalence class of rays, where two rays of G are *equivalent* if there are infinitely many vertex-disjoint paths between them in G. The *degree* $\deg(\omega) \in \mathbb{N} \cup \{\infty\}$ of an end ω of G is the maximum size of a collection of pairwise disjoint rays in ω, see Halin [44]. Ends of infinite degree are also called *thick*. The *half-grid*, the graph on $\mathbb{N}^2$ in which two vertices (n, m) and (n', m') are adjacent if and only if $|n - n'| + |m - m'| = 1$, and its sibling the *hexagonal half-grid*, where one deletes every other rung from the half-grid as shown in Figure 4.1.1, are examples of graphs which have only one end, which is thick.

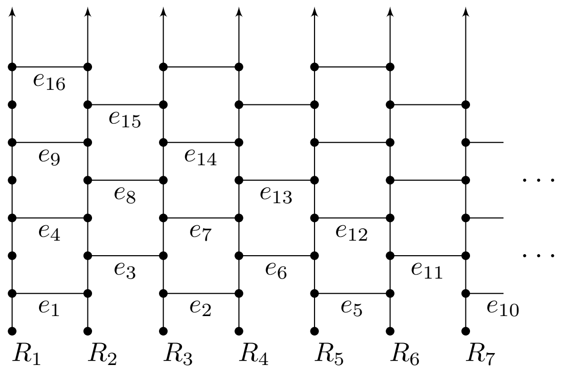

Figure 4.1.1.: The hexagonal half-grid with vertical rays R_i.

One of the cornerstones of infinite graph theory, *Halin's grid theorem* [44], says that grid-like graphs are the prototypes for ends of infinite degree:

Halin's grid theorem. *Every graph with an end of infinite degree contains a subdivision of the hexagonal half-grid whose rays belong to that end.*

Halin's theorem is a precursor of the work by Robertson, Seymour and Thomas on excluding infinite grid or clique minors [64] and has further influenced research in [4, 39, 43, 47]. It is curious, however, that Halin's theorem does not mention any specific ray families, that is if one chooses a specific infinite collection $\mathcal{R}$ of disjoint rays witnessing that the end is thick, then neither the assertion of Halin's theorem nor its available proofs by Halin [44, Satz 4] and by Diestel [24, 25] make any assertion on how the resulting subdivided hexagonal half-grid relates to the collection of rays $\mathcal{R}$ one started with. Furthermore, in recent work of ours on an extension of Halin's grid theorem to higher cardinals, see [39] and Chapter 5, it became quite important to achieve more control of specific uncountable ray families, and the question arose whether this can be done also in the countable case. And indeed, the main result of this chapter is that this is in fact possible:

Theorem 4.1. *For every infinite collection $\mathcal{R}$ of disjoint equivalent rays in a graph G there is a subdivision of the hexagonal half-grid in G such that all its vertical rays belong to $\mathcal{R}$.*

The known proofs of Halin's grid theorem are rather involved and require an elaborate recursive construction that runs close to five pages in Diestel's textbook. Our stronger result in Theorem 4.1 requires a different approach – which coincidentally provides a much shorter proof of Halin's original grid theorem.

4.2. The proof

Suppose we are handed a countably infinite collection $\mathcal{R}$ of disjoint equivalent rays in a graph G. By routine arguments, we fix, for the remainder of this chapter, a ray S in G that meets each ray in $\mathcal{R}$ infinitely often. An *$\mathcal{R}$-segment (of S)* is any maximal subpath of S which is internally disjoint from $\bigcup \mathcal{R}$. We say that an $\mathcal{R}$-segment is *between* two rays $R_1, R_2 \in \mathcal{R}$ if it has its endpoints on R_1 and R_2 respectively. Let $M(\mathcal{R})$ denote the auxiliary multigraph with vertex set $\mathcal{R}$ where the multiplicity of an edge R_1R_2 is equal to the number of $\mathcal{R}$-segments between R_1 and R_2. Finally, let $M_\infty(\mathcal{R})$ denote the spanning subgraph of $M(\mathcal{R})$ obtained by removing all edges of finite multiplicity.

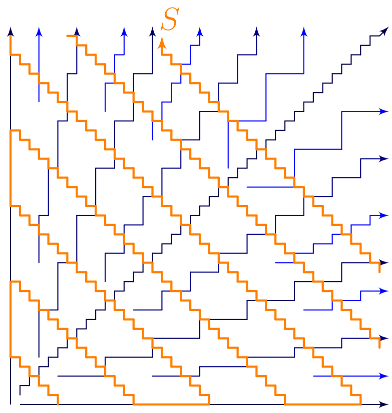

Figure 4.2.1.: A configuration of rays in the half-grid with edge-less $M_\infty(\mathcal{R})$.

Since S meets every ray in $\mathcal{R}$ infinitely often, it follows that $M(\mathcal{R})$ is infinitely edge-connected. Further, if $M_\infty(\mathcal{R})$ has a component with infinitely many vertices, then Theorem 4.1 follows at once: In this case, $M_\infty(\mathcal{R})$ either contains a ray $R_1, R_2, \dots$ or an infinite star with centre R and leaves $R_1, R_2, \dots$. In the ray case, one recursively selects sufficiently late $\mathcal{R}$-segments to represent subdivided edges $e_1, e_2, e_3, \dots$ of the hexagonal half-grid in the order indicated in Figure 4.1.1; in the star case, edges between R_i and R_j are represented by two sufficiently late $\mathcal{R}$-segments between R_i, R_j and R together with the subpath on R connecting the endpoints of those segments.

However, $M_\infty(\mathcal{R})$ might have no edges at all: Consider for example a collection of radial rays in the half-grid such that between any two rays there lies a third, see

Figure 4.2.1. Still, by moving to an infinite subcollection $\mathcal{R}' \subseteq \mathcal{R}$ and considering auxiliary multigraphs $M(\mathcal{R}')$ and $M_\infty(\mathcal{R}')$ instead, the connectivity properties of $M_\infty(\mathcal{R}')$ might improve. Indeed, the auxiliary multigraphs for $\mathcal{R}'$ remain well-defined as the same S still meets every ray in $\mathcal{R}'$ infinitely often. Note, however, that $\mathcal{R}$-segments of S may now be properly contained in $\mathcal{R}'$-segments of S. Our preceding discussion can be summarized as:

(i) The auxiliary multigraph $M(\mathcal{R}')$ is infinitely edge-connected for any infinite subcollection $\mathcal{R}' \subseteq \mathcal{R}$.
(ii) If for some $\mathcal{R}' \subseteq \mathcal{R}$ the auxiliary multigraph $M_\infty(\mathcal{R}')$ has an infinite component, there is a subdivision of the hexagonal half-grid in G such that all its vertical rays belong to $\mathcal{R}'$.

Our next observation provides a sufficient condition for $M_\infty(\mathcal{R}')$ to have an infinite component. Recall that the degree of a vertex in a multigraph denotes the number of its neighbours.

(iii) If $\mathcal{R}' \subseteq \mathcal{R}$ is infinite such that $M(\mathcal{R}')$ has only finitely many vertices of infinite degree, then $M_\infty(\mathcal{R}')$ has an infinite component.

Indeed, suppose for a contradiction that all components of $M_\infty(\mathcal{R}')$ are finite. Then there is also a finite component C of $M_\infty(\mathcal{R}')$ that contains none of the finitely many vertices that have infinite degree in $M(\mathcal{R}')$. Since $M(\mathcal{R}')$ is infinitely edge-connected by (i), there are infinitely many edges in $M(\mathcal{R}')$ from C to its complement. And since C consists of vertices of finite degree only, the neighbourhood of C in $M(\mathcal{R}')$ is finite. Thus, there is a vertex in C that sends infinitely many edges to some vertex outside of C, contradicting the choice of C. This establishes (iii).

The idea of the proof of Theorem 4.1 is now as follows: If all vertices of $M(\mathcal{R})$ have finite degree, then $M_\infty(\mathcal{R})$ has an infinite component by (iii) and we are done by (ii). Otherwise, there is a ray $R_1 \in \mathcal{R}$ that has infinitely many neighbours $N(R_1)$ in $M(\mathcal{R})$, and we may restrict our collection of rays to $\mathcal{R}_1 := \{R_1\} \cup N(R_1)$. Next, if all but finitely many rays of $\mathcal{R}_1$ have finite degree in $M(\mathcal{R}_1)$, then $M_\infty(\mathcal{R}_1)$ has an infinite component by (iii) and we are again done by (ii). Thus, we may pick a second ray R_2 in $\mathcal{R}_1$ distinct from R_1 such that $N(R_2)$ is infinite in $M(\mathcal{R}_1)$, and restrict our collection of rays to $\mathcal{R}_2 := \{R_1, R_2\} \cup N(R_2)$. Repeating this step as often as possible gives rise to a sequence of rays $R_1, R_2, R_3, \ldots$. If this procedure ever stops because there are no more vertices of infinite degree to choose, then we are done by (iii) and (ii). Thus, the question becomes what to do when this procedure does not terminate.

Informally, the solution is to modify our construction so that besides the first n rays $R_1, \ldots, R_n$ we will also have chosen suitable paths $P_1, \ldots, P_{n-2}$ between them representing the subdivided edges $e_1, \ldots, e_{n-2}$ in the copy of the hexagonal half-grid from Figure 4.1.1.[1] Then, in the case where our procedure never stops, the chosen rays $R_1, R_2, \ldots$ become the vertical rays of a hexagonal half-grid where the subdivided paths corresponding to an edge e_i are given by the path P_i.

Formally, suppose that at step n we have chosen n distinct rays $R_1, \ldots, R_n$ from $\mathcal{R}$ and an infinite subcollection $\mathcal{R}_n \subseteq \mathcal{R}$ containing all chosen R_i such that in $M(\mathcal{R}_n)$ every R_i

[1]The index shift just has the purpose that when choosing a path for e_2 we have already selected R_3 and R_4.

is adjacent to all rays in $\mathcal{R}_n \setminus \{R_1, \dots, R_n\}$. Further, suppose that we have chosen $n-2$ disjoint paths $P_1, \dots, P_{n-2}$ internally disjoint from $\bigcup \mathcal{R}_n$, such that each P_i connects the same two rays from $\{R_1, \dots, R_n\}$ as e_i in Figure 4.1.1, in a way such that whenever two paths P_i, P_j with $i < j$ have endvertices on the same ray R_k, then the endvertex of P_i comes before the endvertex of P_j on R_k.

Now if all but finitely many rays in $\mathcal{R}_n$ have finite degree in $M(\mathcal{R}_n)$, then we are done by (iii) and (ii). Hence, we may assume that there is a ray R_{n+1} in $\mathcal{R}_n \setminus \{R_1, \dots, R_n\}$ that has infinitely many neighbours $N(R_{n+1})$ in $M(\mathcal{R}_n)$. Let $\mathcal{R}'_{n+1} := \{R_1, \dots, R_{n+1}\} \cup N(R_{n+1})$, and note that in $M(\mathcal{R}'_{n+1})$, every R_i for $i = 1, \dots, n+1$ is adjacent to all other rays in $\mathcal{R}'_{n+1}$. Now let i and j denote the indices of the rays in Figure 4.1.1 containing the endvertices of the edge e_{n-1}. Note that $i, j \leq n+1$. Since both R_i and R_j have infinitely many common neighbours in $M(\mathcal{R}'_{n+1})$, we also find a common neighbour Q_{n-1} in $\mathcal{R}'_{n+1} \setminus \{R_1, \dots, R_{n+1}\}$ such that the corresponding $\mathcal{R}'_{n+1}$-segments of S between R_i, R_j and Q_{n-1} are disjoint from all earlier paths $P_1, \dots, P_{n-2}$ and also have their endvertices on R_i, R_j later than the endvertices of any previous path $P_1, \dots, P_{n-2}$. Then we may pick a new path P_{n-1} consisting of both these $\mathcal{R}'_{n+1}$-segments of S between R_i, R_j and Q_{n-1} together with a suitable subpath of Q_{n-1}. Finally, set $\mathcal{R}_{n+1} := \mathcal{R}'_{n+1} \setminus \{Q_{n-1}\}$. This completes the induction step, and the proof is complete. □

We remark that only in the case where our procedure stops and (ii) yields a ray one can just build a grid, in all other cases one can build a clique of rays, i.e. one finds an infinite $\mathcal{R}' \subseteq \mathcal{R}$ and a family of internally disjoint $\bigcup \mathcal{R}'$-paths witnessing that any two rays in $\mathcal{R}'$ are equivalent.

5. Halin's end degree conjecture

5.1. Overview

5.1.1. Halin's end degree conjecture

An *end* of a graph G is an equivalence class of rays, where two rays of G are *equivalent* if there are infinitely many vertex-disjoint paths between them in G. The *degree* $\deg(\varepsilon)$ of an end ε is the maximum cardinality of a collection of pairwise disjoint rays in ε, see Halin [44].

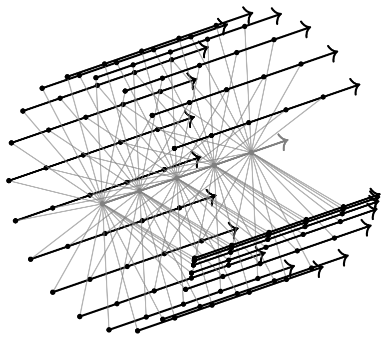

Figure 5.1.1.: The Cartesian product of a star and a ray.

However, for many purposes a degree-witnessing collection $\mathcal{R} \subseteq \varepsilon$ on its own forgets significant information about the end, as it tells us nothing about how G links up the rays in $\mathcal{R}$; in fact $G[\bigcup \mathcal{R}\,]$ is usually disconnected. Naturally, this raises the question of whether one can describe typical configurations in which G must link up the disjoint rays in some degree-witnessing subset of a pre-specified end.

Observing that prototypes of ends of any prescribed degree are given by the Cartesian product of a sufficiently large connected graph with a ray (see e.g. Figure 5.1.1), Halin [44] made this question precise by introducing the notion of a 'ray graph', as follows.

Given a set $\mathcal{R}$ of disjoint equivalent rays in a graph G, we call a graph H with vertex set $\mathcal{R}$ a *ray graph* in G if there exists a set $\mathcal{P}$ of independent $\mathcal{R}$-paths (independent paths with precisely their endvertices on rays from $\mathcal{R}$) in G such that for each edge RS of H there are infinitely many disjoint R–S paths in $\mathcal{P}$. Given an end ε in a graph G, a *ray graph for* ε is a connected ray graph in G on a degree-witnessing subset of ε. The precise formulation of the question reads as follows:

Does every graph contain ray graphs for all its ends?

For ends of finite degree it is straightforward to answer the question in the affirmative. For ends of countably infinite degree the answer is positive too. The constructions by Halin [44, Satz 4], Diestel [24, 25] and ours, see [53] and Chapter 4, show that in this case the ray graph itself can always be chosen as a ray:

Theorem 5.1.1 (Halin's grid theorem)**.** *Every graph with an end of infinite degree contains a subdivision of the hexagonal half-grid whose rays belong to that end.*

For ends of uncountable degree, however, the question is a 20-year-old open conjecture that Halin stated in his legacy collection of problems [42, Conjecture 6.1]:

Halin's conjecture. *Every graph contains ray graphs for all its ends.*

In this chapter, we settle Halin's conjecture: partly positively, partly negatively, with the answer essentially only depending on the degree of the end in question.

5.1.2. Our results

If the degree in question is $\aleph_1$, then any ray graph for such an end contains a vertex of degree $\aleph_1$, which together with its neighbours already forms a ray graph for the end in question, namely an '$\aleph_1$-star of rays'. Thus, finding in G a ray graph for an end of degree $\aleph_1$ reduces to finding such a star of rays. Already this case has remained open.

Let $\mathrm{HC}(\kappa)$ be the statement that Halin's conjecture holds for all ends of degree κ in any graph. As our first main result, we show that

$$\mathrm{HC}(\aleph_1) \textit{ fails.}$$

So, Halin's conjecture is not true after all. But we do not stop here, for the question whether $\mathrm{HC}(\kappa)$ holds remains open for end degrees $\kappa > \aleph_1$. And surprisingly, we show that $\mathrm{HC}(\aleph_2)$ holds. In fact, we show more generally that

$$\mathrm{HC}(\aleph_n) \textit{ holds for all } n \textit{ with } 2 \leq n \leq \omega.$$

Interestingly, this includes the first singular uncountable cardinal $\aleph_\omega$. Having established these results, it came as a surprise to us that

$$\mathrm{HC}(\aleph_{\omega+1}) \textit{ fails.}$$

How does this pattern continue? It turns out that from this point onward, set-theoretic considerations start playing a role. Indeed

$$\mathrm{HC}(\aleph_{\omega+n}) \textit{ is undecidable for all } n \textit{ with } 2 \leq n \leq \omega, \textit{while } \mathrm{HC}(\aleph_{\omega\cdot 2+1}) \textit{ fails.}$$

The following theorem solves Halin's problem for all end degrees:

Theorem 5.1.2. *The following two assertions about* $\mathrm{HC}(\kappa)$ *are provable in* ZFC*:*

(1) $\mathrm{HC}(\aleph_n)$ *holds for all* $2 \leq n \leq \omega$*,*
(2) $\mathrm{HC}(\kappa)$ *fails for all* κ *with* $\mathrm{cf}\,(\kappa) \in \{\mu^+ \colon \mathrm{cf}\,(\mu) = \omega\}$*; in particular,* $\mathrm{HC}(\aleph_1)$ *fails.*

Furthermore, the following assertions about $\mathrm{HC}(\kappa)$ *are consistent:*

(3) *Under* GCH, $\mathrm{HC}(\kappa)$ *holds for all cardinals not excluded by (2).*
(4) *However, for all* κ *with* $\mathrm{cf}\,(\kappa) \in \{\aleph_\alpha \colon \omega < \alpha < \omega_1\}$ *it is consistent with* ZFC+CH *that* $\mathrm{HC}(\kappa)$ *fails, and similarly also for all* κ *strictly greater than the least cardinal* μ *with* $\mu = \aleph_\mu$*.*

For the consistency results in (4) we do not presuppose any advanced set-theoretic knowledge beyond the usual concepts of ordinals, order trees, cardinals and cofinality. Rather, we identify in the literature a suitable combinatorial statement about tree branches that is known to be consistent, and then set out from there to construct our counterexamples for Halin's conjecture in the cases of (4). The proof sketch below offers a flavour which statements we need precisely.

The affirmative results in (1) and (3) are proved in the first half of this chapter, up to Section 5.5. The counterexample for HC($\aleph_1$), which can be read independently of the other chapters, is constructed in Section 5.6. The remaining counterexamples in (2) and (4) are constructed in Sections 5.7 to 5.9.

5.1.3. Proof sketch

The first step behind our affirmative results for HC(κ) is the observation that it suffices to find some countable set of vertices U for which there is a set $\mathcal{R}$ of κ many rays in ε, all disjoint from each other and from U, such that each $R \in \mathcal{R}$ comes with an infinite family $\mathcal{P}_R$ of disjoint R–U paths which, for distinct R and R', meet only in their endpoints in U: then it is not difficult to find a ray R^* that contains enough of U to include the endpoints of almost all path families $\mathcal{P}_R$, yielding a κ-star of rays on $\{R^*\} \cup \mathcal{R}'$ for some suitable $\mathcal{R}' \subseteq \mathcal{R}$ (Lemma 5.2.1).

While it may be hard to identify a countable such set U directly, for κ of cofinality at least $\aleph_2$ there is a neat greedy approach to finding a similar set U' of cardinality just $< \kappa$ rather than $\aleph_0$ (Lemma 5.2.2).

Let us illustrate the idea in the case of $\kappa = \aleph_2$: Starting from an arbitrary ray R_0 in ε, does $U_0 = V(R_0)$ already do the job? That is to say, are there $\aleph_2$ disjoint rays in ε that are independently attached to U_0 as above? If so, we have achieved our goal. If not, take a maximal set of disjoint rays $\mathcal{R}_0$ in ε all whose rays are independently attached to U_0 as above, and define U_1 to be the union of U_0 together with the vertices from all the rays in $\mathcal{R}_0$ and all their selected paths to U_0. Then $|U_1| \leq \aleph_1$. Does U_1 do the job? If so, we have achieved our goal. If not, continue as above. We claim that when continuing transfinitely and building sets $U_0 \subsetneq U_1 \subsetneq \ldots \subsetneq U_\omega \subsetneq U_{\omega+1} \subsetneq \ldots$, we will achieve our goal at some countable ordinal $< \omega_1$. For suppose not. Then $U'' := \bigcup \{U_i : i < \omega_1\}$ meets all the rays in ε. Indeed, any ray R from ε outside of U'' could be joined to U'' by an infinite family of disjoint R–U'' paths. But then their countably many endvertices already belong to some U_i for $i < \omega_1$, contradicting the maximality of $\mathcal{R}_i$ in the definition of U_{i+1}. Hence, $|U''| = \aleph_2$. For cofinality reasons there is a first index $j = i + 1$ with $|U_j| = \aleph_2$. Now $U' = U_i$ is as required.

Having identified a $<\kappa$-sized set U' together with κ disjoint rays all independently attached to it, we aim to restrict U' to a countable set U while keeping κ many rays attached to U (Section 5.4). For $\kappa = \aleph_2$ this is straightforward, since if U' is written as an increasing $\aleph_1$-union of countable sets, one of them already contains all the endpoints of the path systems for some $\aleph_2$-sized subcollection $\mathcal{R}' \subseteq \mathcal{R}$. Take this countable set as the set U originally sought. This completes the proof of HC($\aleph_2$); the other affirmative results for HC(κ) are similar (Corollary 5.5.2 and Theorem 5.5.4).

What about general cardinalities κ? The above strategy can fail in two different ways: First, if $\mathrm{cf}(\kappa) = \aleph_1$, the greedy approach may not terminate: for example, it may well be possible that $|U'| = \aleph_1$ while $|U_i| = \aleph_0$ for all $i < \omega_1$. And indeed, we will show that rays in ends of degree $\aleph_1$ may be 'arranged like an Aronszajn tree', witnessing the failure of HC($\aleph_1$). This idea can be captured as follows (see Definition 5.6.1 for precise details): For an Aronszajn tree T, consider first a disjoint family of rays $\{R_t \colon t \in T\}$ indexed by the nodes of the tree. If t is a successor of s in T, add an infinite matching between the rays R_t and R_s. And if t is a limit, pick a cofinal ω-sequence $t_0 < t_1 < \ldots < t$ of nodes below t, and add an edge from the nth vertex of R_t to the nth vertex of R_{t_n}, for all $n \in \mathbb{N}$. If these cofinal sequences $t_0 < t_1 < \ldots < t$ below each limit t are chosen carefully (for this we rely on a trick by Diestel, Leader and Todorčević, see Theorem 5.6.4), the resulting graph, which we call the *ray inflation of* T, is one-ended of degree $\aleph_1$ but contains no $\aleph_1$-star of rays. This refutes HC($\aleph_1$). See Theorem 5.6.6 for details.

What about the remaining cardinals κ with $\mathrm{cf}\,(\kappa) = \aleph_1$? Also there, counterexamples to $HC(\kappa)$ exist, and we have a machinery that produces a multitude of such examples: Any counterexample for HC(κ) for regular κ may be turned canonically into a counterexample for HC(λ) for all λ with $\mathrm{cf}\,(\lambda) = \kappa$, see Theorem 5.9.1.

The second way in which our above strategy can fail is that even if our greedy algorithm terminates and provides a $<\kappa$-sized U' to which there are κ disjoint rays independently attached, it may not be possible to further reduce U' as earlier to some countable subset U. And indeed, using our idea of ray inflations of order trees, also this constellation can be exploited to construct counterexamples to Halin's conjecture. However, the trees that work now are quite different from the earlier Aronszajn trees: Generalising the concept of *binary trees with tops* introduced by Diestel and Leader in [30], we consider the class of λ-*regular trees with tops*, where λ is any singular cardinal of countable cofinality. These are order trees of height $\omega + 1$ in which every point of finite height has exactly λ successors, and above some $\kappa > \lambda$ many selected branches we add further nodes to the tree at height ω, called *tops*.

There is a reason why we take λ to be singular of countable cofinality: Just like the binary tree has uncountably many branches, these λ's are the only other cardinals for which an uncountable regular tree is guaranteed to have strictly more than λ branches. And just like the precise number of branches of the binary tree is not determined in ZFC alone (it is $2^{\aleph_0}$, which may be $\aleph_1$ if CH holds, or may be arbitrarily large), also the precise number of branches of the λ-regular tree is λ^+ if GCH holds, but it also may be much larger.

Now the starting point for our consistent counterexamples of (4) in Theorem 5.1.2 are simply models of ZFC + CH in which the two λ-regular trees for $\lambda = \aleph_\omega$ and λ equal to the first fixed point of the $\aleph$-function have a lot more branches than nodes. In these cases, any λ-regular tree with tops gives rise to counterexamples for Halin's conjecture (Theorem 5.8.1). What happens if one looks for ZFC-counterexamples, not just consistent ones? With significantly more effort, and building on the concept of a *scale* from Shelah's pcf-theory, we will show in Theorem 5.7.1 that for any singular λ of countable cofinality one can directly select a suitable set of λ^+ many branches so that the λ-regular tree with corresponding tops gives rise to the counterexamples for Halin's conjecture, settling the remaining cases of (2) in Theorem 5.1.2.

5.1.4. Open problems

We suspect that (1) and (2) in Theorem 5.1.2 capture all the cases of Halin's conjecture that can be proved or disproved in ZFC alone. This is certainly true up to $\aleph_{\omega_1}$, as for each $\kappa \leq \aleph_{\omega_1}$ our main Theorem 5.1.2 provides either a ZFC or an independence result regarding the truth value of HC(κ). While for all remaining cardinals assertion (3) of Theorem 5.1.2 gives consistent affirmative results, we do not know whether any of these can be established in ZFC.

Question 1. *Is Halin's conjecture true for any $\kappa > \aleph_{\omega_1}$?*

Question 2. *Is it true that any end of degree $\aleph_1$ either contains an $\aleph_1$-star of rays or a subdivision of a ray inflation of an Aronszajn tree, as in Theorem 5.6.6?*

Question 3. *Is it true that for every cardinal κ there is $f(\kappa) \geq \kappa$, such that every end ε of degree $f(\kappa)$ contains a connected ray graph of size κ?*

5.2. Ray collections with small core

For general notions in graph theory and in set theory we generally follow the textbooks by Diestel [25] and Jech [49]. Let ε be an end of a graph G, and U be a set of vertices in G. An *ε–U comb* is a subgraph $C = R \cup \bigcup \mathcal{P}$ of G that consist of a ray R disjoint from U that represents ε and an infinite family $\mathcal{P}$ of disjoint R–U paths. The vertices in $C \cap U$ are the *teeth* of the comb. We write $\mathring{C}$ for $C - U$, the *interior* of the ε–U comb, which is disjoint from U. We call two ε–U combs *internally disjoint* if they have disjoint interior.

Lemma 5.2.1. *Let ε be an end of a graph G and U a countable set of vertices. If there is an uncountable collection $\mathcal{C}$ of internally disjoint ε–U combs in G, then ε contains a $|\mathcal{C}|$-star of rays whose leaf rays are the spines of (a subset of) combs in $\mathcal{C}$.*

As this lemma is fundamental to our affirmative results, we provide two proofs: one relying on the theory of normal trees, and a second, more elementary proof, using dominating vertices. Recall that a rooted tree $T \subseteq G$ is *normal (in G)* if the endvertices of any T-path in G (a non-trivial path in G with endvertices in T but all edges and inner vertices outside of T) are comparable in the tree order of T. Thus, if $T \subseteq G$ is normal, then the neighbourhoods of the components of $G - T$ form chains in T. The rays of T starting in the root are called *normal rays*.

First proof of Lemma 5.2.1. We may assume that G is connected. Since U is countable, by Jung's Theorem (see [50, Satz 6], or the proof in [25, Theorem 8.2.4]), there is a countable normal tree $T \subseteq G$ that includes U. As T is countable, $\mathcal{C}$ contains a subcollection $\mathcal{C}'$ of size $|\mathcal{C}|$ such that the interior of every comb in $\mathcal{C}'$ is disjoint from T. As T is normal, the teeth of any comb in $\mathcal{C}'$ lie on the unique normal ray R in T that represents ε, see [25, Lemma 8.2.3]. Consequently, R with all combs in $\mathcal{C}'$ forms the desired $|\mathcal{C}|$-star of rays. $\square$

A vertex $v \in G$ *dominates* a ray $R \subseteq G$ if v is the centre of a subdivided infinite star with all its leaves in R. It *dominates* an end ε of G if it dominates some (equivalently: each) ray in ε.

Second proof of Lemma 5.2.1. Denote by $U' \subseteq U$ the vertices of U that are teeth of only finitely many combs in $\mathcal{C}$, and let $\mathcal{C}' \subseteq \mathcal{C}$ be the subcollection of combs with a tooth in U'. Then $\mathcal{C}'$ is countable, so $\mathcal{C} \setminus \mathcal{C}'$ is of size $|\mathcal{C}|$, and every comb in $\mathcal{C} \setminus \mathcal{C}'$ has all its teeth contained in $U \setminus U'$. Since all vertices in $U \setminus U'$ dominate the end ε, there is a ray R in G with $U \setminus U' \subseteq V(R)$, cf. [25, Ex. 8.29 & 8.30]. Since R is countable, $\mathcal{C} \setminus \mathcal{C}'$ contains a subcollection $\mathcal{C}''$ of size $|\mathcal{C}|$ such that the interior of every comb in $\mathcal{C}''$ is disjoint from R. Consequently, R together with all the combs in $\mathcal{C}''$ forms the desired $|\mathcal{C}|$-star of rays. □

Given an end ε of a graph G with uncountable degree κ, finding a countable vertex set $U \subseteq V(G)$ so that G admits κ internally disjoint ε–U combs (i.e. so that U satisfies the premise of Lemma 5.2.1) may be hard, and sometimes even impossible by Theorem 5.1.2(1). Perhaps surprisingly, U can be found greedily if κ has large cofinality and the condition $|U| = \aleph_0$ is relaxed to $|U| < \kappa$.

Lemma 5.2.2 (The greedy lemma). *Let $\mathcal{R}$ be any κ-sized collection of disjoint rays belonging to an end ε of G. If $\operatorname{cf}(\kappa) > \aleph_1$, then there exist a set of vertices U in G with $|U| < \kappa$ and a κ-sized collection $\mathcal{C}$ of internally disjoint ε–U combs in G with all spines in $\mathcal{R}$.*

Proof. Let $U_0 := V(R)$ for an arbitrarily chosen ray $R \in \mathcal{R}$. Recursively construct a sequence $(U_i : i < \omega_1)$ of vertex sets of G as follows: If U_i is already defined, use Zorn's lemma to choose a maximal collection $\mathcal{C}_i$ of internally disjoint ε–U_i combs with spines in $\mathcal{R}$, and let $U_{i+1} := U_i \cup \bigcup \mathcal{C}_i$. For a limit $\ell < \omega_1$, simply define $U_\ell := \bigcup_{i<\ell} U_i$.

Consider $U' := \bigcup_{i<\omega_1} U_i$. If $|U'| < \kappa$, then since $|\mathcal{R}| = \kappa$ there still exists an ε–U' comb C in G with spine in $\mathcal{R}$. However, the countably many teeth from $C \cap U'$ belong already to some U_i for $i < \omega_1$. But then the existence of C contradicts the maximality of $\mathcal{C}_i$.

Hence, $|U'| = \kappa$, and $\operatorname{cf}(\kappa) > \omega_1$ implies that there is a first $i < \omega_1$ such that $|U_i| = \kappa$ which must be a successor, say $i = j + 1$. Then $U := U_j$ satisfies $|U| < \kappa$ and $\mathcal{C} := \mathcal{C}_j$ is a κ-sized collection of ε–U combs as desired. □

For which cardinals κ it is possible to bridge the gap between Lemma 5.2.1 and Lemma 5.2.2 is discussed in Section 5.4 below. Before doing so, however, we use Lemma 5.2.1 to classify the minimal types of connected ray graphs in the next section.

5.3. Typical types of ray graphs

It is well known that every countable connected graph contains a vertex of infinite degree or a ray. This carries over to larger cardinals as follows. Every connected graph on κ many vertices, for regular uncountable κ, has a vertex of degree κ. Indeed, consider the distance classes from any fixed vertex of the graph. Then there is a first distance class with κ many vertices and the prior distance class contains a vertex of degree κ. We call a star with κ many leafs a *κ-star*.

A *frayed star* is a rooted tree in which all vertices are within distance 2 from the root. The root is also called the *centre* and the neighbours of the root are called the *distributor vertices* of the frayed star. For a cardinal κ and a cofinal sequence $s = (\kappa_i \colon i < \mathrm{cf}(\kappa))$ in κ, a (κ, s)-*star* is a frayed star such that the sequence of degrees of the distributor vertices corresponds to s. Note that if for a frayed star S with κ many leafs the degrees of the distributor vertices form a cofinal sequence for κ, then S contains a (κ, s)-star for any prescribed $\mathrm{cf}(\kappa)$-sequence s of κ as a subgraph. A *frayed comb* is a ray together with infinitely many disjoint stars such that exactly one vertex of every star lies on the ray. The centres of these stars are the *distributor vertices* and the leafs of the stars that do not lie on the ray are the *teeth* of the frayed comb.

Recall that every connected graph on κ many vertices contains a subdivided frayed star or a subdivided frayed comb with κ many leafs or teeth, respectively. To see this we may assume the graph to be a tree T rooted at r. Now, consider the size of every component of $T - r$. If there is one component C of size κ we view C as a tree rooted at the neighbour of r and continue in C. If there is again and again a component of size κ it is straightforward to find a frayed comb with κ many teeth. However, if all components are smaller than κ and their size forms a cofinal sequence of κ it is straightforward to find a subdivided frayed star with κ many leafs. Finally, if all components are smaller than κ and their size does not form a cofinal sequence in κ, there are at least κ many components in which case we even find a star of size κ. See also [41, Corollary 8.1] for a strengthening of the above results.

We now lift these basic facts to ray graphs. For connected graph H, we say that the union of disjoint rays $(R_h \colon h \in V(H))$ in some graph G together with a path family $\mathcal{P}$ in G is a H *graph of rays* in G if the ray graph of $(R_h \colon h \in V(H))$ and $\mathcal{P}$ is H. If $X \subseteq V(H)$ is a set of vertices we say that the rays $(R_x \colon x \in X)$ are the X *rays* in G. If $H = S$ is a star with centre $c \in V(S)$ and leafs L, then we say that R_c is the centre ray and $(R_\ell \colon \ell \in L)$ are the leaf rays.

Lemma 5.3.1. *Let ε be any end of a graph G of degree κ, and suppose that* $\mathrm{HC}(\kappa)$ *holds.*

(1) *If κ is regular and uncountable, then G contains a κ-star of rays all belonging to ε.*
(2) *If κ is singular and s is any* $\mathrm{cf}(\kappa)$*-sequence of cardinals with supremum κ, then G contains either a κ-star of rays all belonging to ε or a (κ, s)-star of rays all belonging to ε.*

Proof. As κ is regular in assertion (1), any ray graph witnessing that $\mathrm{HC}(\kappa)$ holds is connected, and so has a vertex of degree κ. Clearly, any such vertex gives rise to a κ-star of rays in G.

For (2) we find either a frayed star or a frayed comb with κ many leafs or teeth, respectively, in the connected ray graph. First, suppose that we obtain a subdivided frayed star with κ many leafs. This gives rise to a subdivided frayed star of rays in G. However, this can easily be turned into a κ-star of rays or a (κ, s)-star of rays, respectively. Indeed, any ray which has degree two in the ray graph might be abandoned to find a path family between its two neighbour rays (internally disjoint from all other rays and paths). Second, suppose that we obtain a frayed comb of size κ in the ray graph. Then there is a frayed comb of rays in G and by a similar argument as above we may assume that the frayed comb has no vertices of degree two. Hence, there are only countably many vertices in the frayed comb that are not a tooth; denote by U the set of all these vertices. Now, apply Lemma 5.2.1 to the set $U' = \{V(R_u) \colon u \in U\}$ and the ε–U' combs provided by the teeth rays to obtain a κ-star of rays. □

5.4. On (λ, κ)-graphs I

A (λ, κ)-*graph* is a bipartite graph (A, B) with $|A| = \lambda$ and $|B| = \kappa$, where $\lambda < \kappa$ are infinite cardinals and every vertex $b \in B$ has infinitely many neighbours in A.

Next to their role in this chapter, (λ, κ)-graphs and in particular (λ, λ^+)-graphs also occur in the characterisation of the Erdős-Hajnal colouring number [3] as well as in the characterisation of graphs with normal spanning trees [59]. For a structural classification of $(\aleph_0, \aleph_1)$-graphs see [6].

In the following we investigate the question which (λ, κ)-graphs contain an $(\aleph_0, \kappa)$-subgraph. Whenever we speak of a (λ', κ)-subgraph of another (λ, κ)-graph in this chapter, we silently assume that the bipartition classes are respected. This is no significant restriction, since whenever a (λ, κ)-graph (A, B) has a (λ', κ)-subgraph (C, D), then it also has such a subgraph that respects the bipartition classes (simply restrict to the vertices in $D \setminus A$ and their neighbours in $A \cap C$).

Lemma 5.4.1. *Any (λ, κ)-graph with $\mathrm{cf}(\lambda) > \omega$ and $\mathrm{cf}\,(\kappa) \neq \mathrm{cf}\,(\lambda)$ contains a (λ', κ)-subgraph for some $\lambda' < \lambda$.*

Proof. Let (A, B) be a (λ, κ)-graph as in the lemma. We may assume that $N(b) \subseteq A$ is countable for all $b \in B$. Write $A = \bigcup \{A_i \colon i < \mathrm{cf}\,(\lambda)\}$ as an increasing union of subsets A_i with $|A_i| < \lambda$. Let $B_i := \{b \in B \colon N(b) \subseteq A_i\}$. Since $\mathrm{cf}(\lambda) > \omega$, we have $B = \bigcup \{B_i \colon i < \mathrm{cf}\,(\lambda)\}$. Recall that if a non-decreasing γ-sequence in a limit ordinal α is cofinal in α, then $\mathrm{cf}(\gamma) = \mathrm{cf}(\alpha)$, see e.g. [49, Lemma 3.7(ii)]. As $\mathrm{cf}\,(\kappa) \neq \mathrm{cf}\,(\lambda)$, we therefore have $|B_i| = |B|$ for some $i < \mathrm{cf}\,(\lambda)$. Then (A_i, B_i) is the desired (λ', κ)-subgraph. □

Corollary 5.4.2. *Any (λ, κ)-graph with $\kappa \leq \aleph_\omega$ contains an $(\aleph_0, \kappa)$-subgraph.* □

However, this pattern breaks down at the cardinal $\aleph_{\omega+1}$. In fact, we will show below that for every singular cardinal λ of countable cofinality, there exist (λ, λ^+)-graphs without $(\aleph_0, \lambda^+)$-subgraph, Theorem 5.7.1.

In order to prove the affirmative results of item (3) in our main result Theorem 5.1.2, we now show that under GCH[1], these are more or less the only exceptions:

Lemma 5.4.3. *Under* GCH*, every (λ, κ)-graph with $\mathrm{cf}\,(\kappa) \notin \{\mu^+ \colon \mathrm{cf}\,(\mu) = \omega\}$ and $\mathrm{cf}\,(\kappa) > \lambda > \aleph_0$ contains a (λ', κ)-subgraph for some $\lambda' < \lambda$.*

Proof. Let (A, B) be a (λ, κ)-graph as in the lemma. We may assume that $N(b) \subseteq A$ is countable for all $b \in B$. If $\mathrm{cf}\,(\lambda) > \omega$, the assertion follows from Lemma 5.4.1.

If $\mathrm{cf}\,(\lambda) = \omega$, GCH implies that $|[A]^{\aleph_0}| = \lambda^+$, see [49, Theorem 5.15(ii)]. By assumption, $\mathrm{cf}\,(\kappa) > \lambda$ and $\mathrm{cf}\,(\kappa) \neq \lambda^+$, so $\mathrm{cf}\,(\kappa) > \lambda^+$. Hence, for some countable subset $A' \subseteq A$ we

[1] In fact, a close inspection shows that our results only require the consequence of GCH that $\lambda^{\aleph_0} = \lambda^+$ for all singular cardinals λ of countable cofinality.

find a κ-sized $B' \subseteq B$ such that $N(b) = A'$ for all $b \in B'$. Then (A', B') is the desired subgraph. $\square$

Corollary 5.4.4. *Under* GCH, *every* (λ,κ)*-graph with* $\mathrm{cf}\,(\kappa) \notin \{\mu^+ : \mathrm{cf}\,(\mu) = \omega\}$ *and* $\mathrm{cf}\,(\kappa) > \lambda$ *contains an* $(\aleph_0,\kappa)$*-subgraph.* $\square$

Corollary 5.4.5. *Under* GCH, *every* (λ,κ)*-graph for regular* $\kappa \notin \{\mu^+ : \mathrm{cf}\,(\mu) = \omega\}$ *contains an* $(\aleph_0,\kappa)$*-subgraph.* $\square$

Lemma 5.4.6. *Under* GCH, *every* (λ,κ)*-graph with* $\mathrm{cf}\,(\kappa) \notin \{\mu^+ : \mathrm{cf}\,(\mu) = \omega\}$ *contains an* $(\aleph_0,\kappa)$*-subgraph or otherwise a collection of disjoint* $(\aleph_0,\kappa_i)$*-subgraphs for* $\{\kappa_i : i < \mathrm{cf}\,(\kappa)\}$ *cofinal in* κ.

Proof. Let (A, B) be some (λ,κ)-graph as in the statement of the lemma, but without an $(\aleph_0,\kappa)$-subgraph. By Corollary 5.4.5, κ is singular and we may write $\kappa = \sup\{\kappa_i : i < \mathrm{cf}\,(\kappa)\}$ for regular $\kappa_i > \lambda$ with $\kappa_i \notin \{\mu^+ : \mathrm{cf}\,(\mu) = \omega\}$. Recursively choose disjoint $(\aleph_0,\kappa_i)$-subgraphs (A_i, B_i) of (A, B) for $i < \mathrm{cf}\,(\kappa)$ as follows: Given some fixed $i < \mathrm{cf}\,(\kappa)$ with pairwise disjoint (A_j, B_j) already selected for all $j < i$, let

$$A'_i := A \setminus \bigcup\{A_j : j < i\} \quad\text{and}\quad B'_i := \{b \in B \setminus \bigcup\{B_j : j < i\} : |N(b) \cap A'_i| = \infty\}.$$

If $|B'_i| < \kappa$, then since $|\bigcup\{A_j : j < i\}| < \mathrm{cf}\,(\kappa)$, the subgraph

$$(\bigcup\{A_j : j < i\}, B \setminus B'_i)$$

of (A, B) would contain an $(\aleph_0,\kappa)$-subgraph by Corollary 5.4.4, contradicting our initial assumption. Therefore, we may choose $B''_i \subseteq B'_i$ of size κ_i, and apply Corollary 5.4.5 to the graph (A'_i, B''_i) and obtain an $(\aleph_0,\kappa_i)$-subgraph (A_i, B_i) of (A, B) disjoint from all (A_j, B_j) for $j < i$ as desired. $\square$

Corollary 5.4.7. *Under* GCH, *every* (λ,κ)*-graph with* $\mathrm{cf}(\kappa) = \omega$ *contains an* $(\aleph_0,\kappa)$*-subgraph.*

Proof. By Lemma 5.4.6, any (λ,κ)-graph with $\mathrm{cf}\,(\kappa) = \omega$ contains either an $(\aleph_0,\kappa)$-subgraph, in which case we are done, or otherwise a collection of disjoint $(\aleph_0,\kappa_i)$-subgraphs (A_i, B_i) for $\{\kappa_i : i < \omega\}$ cofinal in κ. But then $(\bigcup_{i<\omega} A_i, \bigcup_{i<\omega} B_i)$ is an $(\aleph_0,\kappa)$-subgraph. $\square$

5.5. Affirmative cases in Halin's conjecture

5.5.1. Regular cardinals

Our affirmative results for regular cardinals are now a straightforward consequence of Lemmas 5.2.1 and 5.2.2, together with the results from the previous section.

Proposition 5.5.1. *Let* $\mathcal{R}$ *be any* κ*-sized collection of disjoint equivalent rays in a graph* G *for an uncountable regular cardinal* κ*. If* $\kappa = \aleph_n$ *for some* $n \in \mathbb{N}$ *with* $n \geq 2$*, then* G *contains a* κ*-star of rays with leaf rays in* $\mathcal{R}$*.*

Assuming GCH, *there exists a* κ*-star of rays in* G *with leaf rays in* $\mathcal{R}$ *unless* $\kappa \in \{\mu^+ : \mathrm{cf}\,(\mu) = \omega\}$.

Proof. Let $\mathcal{R}$ be any collection of disjoint equivalent rays of regular cardinality $\kappa > \aleph_1$ belonging to some end ε of G. By Lemma 5.2.2, there exist a set of vertices U in G with $\lambda := |U| < \kappa$ and a κ-sized collection $\mathcal{C}$ of internally disjoint ε–U combs in G with spines in $\mathcal{R}$.

Consider the (λ, κ)-minor $H = (U, \mathcal{C})$ of G where we contract the interior of every ε–U comb $C \in \mathcal{C}$. Applying Corollary 5.4.2 in the case where $\kappa = \aleph_n$ for some $n \in \mathbb{N}$ with $n \geq 2$, and Corollary 5.4.4 otherwise, we obtain that H contains an $(\aleph_0, \kappa)$-subgraph H'. After uncontracting the combs on the κ-side of H', applying Lemma 5.2.1 finishes the proof. □

Corollary 5.5.2. HC$(\aleph_n)$ *holds for all* $n \in \mathbb{N}$ *with* $n \geq 2$. *Moreover, under* GCH, HC(κ) *holds for all regular cardinals* $\kappa \notin \{\mu^+ \colon \mathrm{cf}\,(\mu) = \omega\}$. □

5.5.2. Singular cardinals

We now extend these affirmative results to singular cardinals. First, to singular cardinals of countable cofinality, and then to all singular cardinals whose cofinality is not a successor of a regular cardinal. By case (2) of our main Theorem 5.1.2, this is best possible.

Proposition 5.5.3. HC$(\aleph_\omega)$ *holds. Moreover, under* GCH, HC(κ) *holds whenever* $\mathrm{cf}\,(\kappa) = \omega$.

Proof. Suppose $\kappa > \mathrm{cf}\,(\kappa) = \omega$, and let $(\kappa_n \colon n \in \omega)$ be a strictly increasing sequence of infinite cardinals with supremum κ, and each of the form $\kappa_n = \lambda_n^{++}$. Let G be any graph and ε an end of G of degree κ. For each $n \in \omega$ we use Proposition 5.5.1 to find a κ_n-star of rays S_n in G with all rays in ε. We write R_n for the centre ray of S_n and consider the countable vertex set $U := \bigcup_{n \in \omega} V(R_n)$. For each $n \in \omega$, some κ_n many of the components of S_n are disjoint from both $\bigcup_{i<n} S_i$ and U. Thus, the union of all stars S_n contains κ many internally disjoint ε–U combs. Therefore, we may apply Lemma 5.2.1 to ε and U to find the desired κ-star of rays. □

Theorem 5.5.4. *Under* GCH, HC(κ) *holds for all* κ *with* $\mathrm{cf}\,(\kappa) \notin \{\mu^+ \colon \mathrm{cf}\,(\mu) = \omega\}$.

Proof. Let ε be an end of G with $\deg(\varepsilon) = \kappa$. By the previous results, we may assume that κ is singular and $\mathrm{cf}\,(\kappa) > \aleph_1$. Hence, by the Greedy Lemma 5.2.2, there are some set of vertices $U \subseteq V(G)$ with $|U| < \kappa$ and a κ-sized family $\mathcal{C}$ of internally disjoint ε–U combs. Consider the $(|U|, \kappa)$-minor $H = (U, \mathcal{C})$ of G where we contract the interior of every ε–U comb $C \in \mathcal{C}$. By Lemma 5.4.6, H contains either an $(\aleph_0, \kappa)$-subgraph (in which case we are done by Lemma 5.2.1), or a collection of disjoint $(\aleph_0, \kappa_i)$-subgraphs for $\{\kappa_i \colon i < \mathrm{cf}\,(\kappa)\}$ cofinal in κ with all $\kappa_i > \max\{\aleph_{\omega+1}, \mathrm{cf}\,(\kappa)\}$ regular.

Consider one $(\aleph_0, \kappa_i)$-subgraph H_i. As GCH implies in particular that $2^{\aleph_0} < \aleph_{\omega+1}$, it follows from $\aleph_{\omega+1} < \kappa_i$ and the regularity of κ_i that there is a complete $(\aleph_0, \kappa_i)$-subgraph $H_i' \subseteq H_i$. Uncontracting the κ_i-side of H_i' to combs and applying Lemma 5.2.1 inside the resulting subgraph of G (in which by construction all spines of the combs are still equivalent) gives a star of rays S_i of size κ_i. By construction, any two such stars are disjoint.

Now, we apply Proposition 5.5.1 to the collection $\mathcal{R}$ of all center rays of the S_i to obtain a star of rays S of size $\mathrm{cf}(\kappa)$ with leaf rays in $\mathcal{R}$. Keeping only those S_i whose centre ray is a leaf ray of S, we may assume that S has precisely $\mathcal{R}$ as set of leaf rays. Since $|S| < |S_i|$ for all i, we may assume that each S_i meets S only in the former's centre ray. Then $S \cup \bigcup \{S_i : i < \mathrm{cf}(\kappa)\}$ yields a connected ray graph of size κ. $\square$

5.6. The first counterexample to Halin's conjecture

5.6.1. Order trees, T-graphs and ray inflations

A partially ordered set $(T, \leq)$ is called an *order tree* if it has a unique minimal element (called the *root*) and all subsets of the form $\lceil t \rceil = \lceil t \rceil_T := \{t' \in T : t' \leq t\}$ are well-ordered. Write $\lfloor t \rfloor = \lfloor t \rfloor_T := \{t' \in T : t \leq t'\}$.

A maximal chain in T is called a *branch* of T; note that every branch inherits a well-ordering from T. The *height* of T is the supremum of the order types of its branches. The *height* of a point $t \in T$ is the order type of $\lceil \mathring{t} \rceil := \lceil t \rceil \setminus \{t\}$. The set T^i of all points at height i is the ith *level* of T, and we write $T^{<i} := \bigcup \{T^j : j < i\}$ as well as $T^{\leq i} := \bigcup \{T^j : j \leq i\}$. If $t < t'$, we use the usual interval notation $(t, t') = \{s : t < s < t'\}$ for nodes between t and t'. If there is no point between t and t', we call t' a *successor* of t and t the *predecessor* of t'; if t is not a successor of any point it is called a *limit*.

An order tree T is *normal* in a graph G, if $V(G) = T$ and the two endvertices of any edge of G are comparable in T. We call G a *T-graph* if T is normal in G and the set of lower neighbours of any point t is cofinal in $\lceil \mathring{t} \rceil$. For detailed information on normal tree orders, see also [7].

Given T an order tree, a T-graph G is *sparse* if the down-neighbourhood of any node t is of order type $\mathrm{cf}\left(\lceil \mathring{t} \rceil\right)$. For trees T of height at most ω_1, this means that if t is a successor, it has a unique down-neighbour namely its predecessor, and if t is a limit, its down-neighbours form a cofinal ω-sequence in $\lceil \mathring{t} \rceil$. (For T an ordinal, this corresponds to a *ladder system* in Todorčević's terminology [69]).

An *Aronszajn tree* is an order tree of size $\aleph_1$ where all branches and levels are countable.

We now introduce a concept that is crucial for all our counterexamples to Halin's conjecture.

Definition 5.6.1. *Let G be a sparse T-graph for an order tree T of height at most ω_1. The* ray-inflation $G \sharp \mathbb{N}$ *of G is the graph with vertex set $T \times \mathbb{N}$, and the following edges (cf. Figure 5.6.1):*

1. *For every $t \in T$ we add all the edges $(t, n)(t, n + 1)$ with $n \in \mathbb{N}$ so that $R_t := G[\{t\} \times \mathbb{N}]$ is a* horizontal ray*.*
2. *If $t \in T$ is a successor with predecessor t', we add all edges $(t, n)(t', n)$ for all $n \in \mathbb{N}$.*
3. *If $t \in T$ is a limit with down-neighbours $t_0 <_T t_1 <_T t_2 <_T \dots$ in G we add the edges $(t, n)(t_n, n)$ for all $n \in \mathbb{N}$.*

Lemma 5.6.2. *If G is a sparse T-graph for T an order tree of height at most ω_1, then all the pairwise disjoint horizontal rays R_t in the ray inflation $H = G \sharp \mathbb{N}$ belong to the same sole end ε of $G \sharp \mathbb{N}$; in particular,* $\deg(\varepsilon) = |T|$.

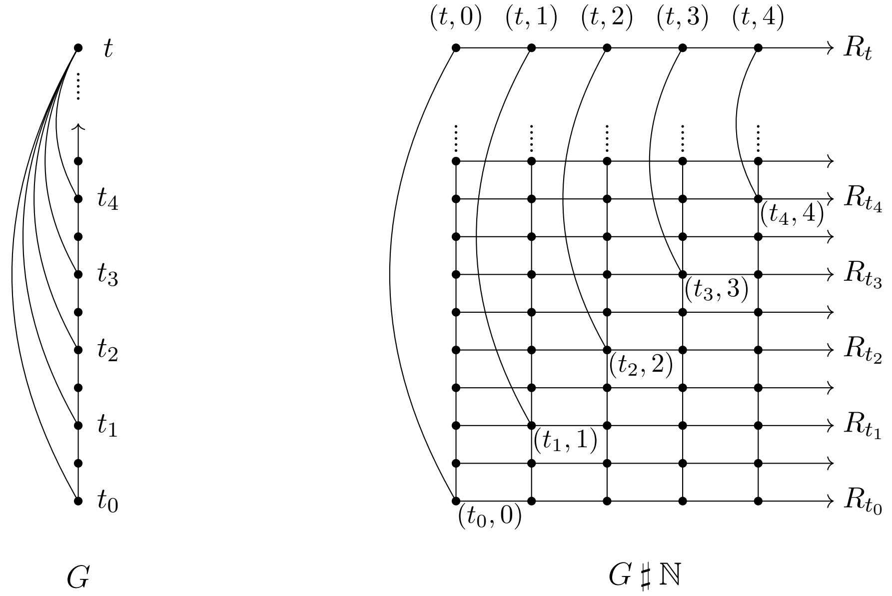

Figure 5.6.1.: The ray inflation of an $(\omega+1)$-graph.

Proof. We show by an ordinal induction on $i \geq 0$ that the rays in the set $\mathcal{R}^{\leq i} := \{R_t \colon t \in T^{\leq i}\}$ belong to the same sole end of $H^{\leq i} := H[T^{\leq i} \times \mathbb{N}]$. For i equal to the height of T, this implies the statement of the lemma. The induction starts because $\mathcal{R}^{\leq 0}$ consists of the ray $H^{\leq 0} = R_r$ for r the root of T. Now suppose that $i > 0$. We have to show that every horizontal ray R_t with $t \in T^i$ is equivalent in $H^{\leq i}$ to some other horizontal ray R_s with $s \in T^{<i}$. For this, let any horizontal ray R_t with $t \in T^i$ be given. Since the vertex set of $H^{<i}$ is partitioned by $\mathcal{R}^{<i}$ into vertex sets of rays that belong to the same sole end of $H^{<i}$ by the induction hypothesis, it suffices to show that R_t can be extended to a comb in $H^{\leq i}$ attached to $H^{<i}$. If t is a successor in T with predecessor t', then such a comb arises from R_t by adding the edges $(t,n)(t',n)$ for all $n \in \mathbb{N}$. And if t is a limit in T with down-neighbours $t_0 <_T t_1 <_T t_2 <_T \dots$ in G, then such a comb arises from R_t by adding the edges $(t,n)(t_n,n)$ for all $n \in \mathbb{N}$. □

Lemma 5.6.3. *Let G be a sparse T-graph for T an order tree of height at most ω_1, let $0 \leq i < \omega_1$ and let H be the ray inflation of G. Then the map $T^i \ni t \mapsto H[\lfloor t \rfloor \times \mathbb{N}]$ is a bijection between T^i and the components of $H - (T^{<i} \times \mathbb{N})$.*

Proof. The T-graph G is the contraction minor of H that arises by contracting every horizontal ray R_t to a single vertex. For T-graphs such as G it is well known, and straightforward to show, that the map $T^i \ni t \mapsto G[\lfloor t \rfloor]$ is a bijection between T^i and the components of $G - T^{<i}$. Since every component of $H - (T^{<i} \times \mathbb{N})$ arises from a component of $G - T^{<i}$ by uncontracting every vertex of that component to a horizontal ray, the claim follows. □

5.6.2. An Aronszajn tree of rays

Our counterexample to $\mathrm{HC}(\aleph_1)$ is based on the ray inflation of an Aronszajn tree. Diestel, Leader and Todorčević showed in [30] that:

Proposition 5.6.4. *There exist an Aronszajn tree T and a sparse T-graph G with the following property:*

For every $t \in T$ there is a finite $S_t \subseteq \lceil \mathring{t} \rceil$ such that every $t' > t$ has all its down-neighbours below t inside S_t. $(\star)$

Proof. We sketch the construction from [30, Theorem 6.2] for convenience of the reader. Let T be an Aronszajn tree with an antichain partition $\{U_n : n \in \mathbb{N}\}$ (the standard Aronszajn tree constructions yield such an antichain partition).

Given a limit $t \in T$, choose its down-neighbours $t_0 < t_1 < t_2 < \dots$ inductively, starting with t_0 the root of T. If $t_{n-1} < t$ has already been defined, consider the least $i \in \mathbb{N}$ such that the antichain U_i meets the interval (t_{n-1}, t), and let t_n be the (unique) point in $U_i \cap (t_{n-1}, t)$. It is easy to check that the T-graph G is as desired. □

Lemma 5.6.5. *If T is an Aronszajn tree and G is any T-graph with property $(\star)$, then the ray inflation $H = G \sharp \mathbb{N}$ of G has the following property:*

For every $t \in T$ there is a finite $S_t \subseteq \lceil \mathring{t} \rceil$ such that every $(t', n) \in H$ with $t' >_T t$ and $n \in \mathbb{N}$ satisfies $N_H((t', n)) \cap (\lceil \mathring{t} \rceil \times \mathbb{N}) \subseteq S_t \times |S_t|$. $(\star\star)$

Proof. Given $t \in T$, and given S_t by property $(\star)$, we show that every $(t', n) \in H$ with $t' >_T t$ and $n \in \mathbb{N}$ satisfies $N := N_H((t', n)) \cap (\lceil \mathring{t} \rceil \times \mathbb{N}) \subseteq S_t \times |S_t|$. If t' is a successor, then N is empty, so $N \subseteq S_t \times |S_t|$. Otherwise t' is a limit and has down-neighbours $t'_0 <_T t'_1 <_T t'_2 \dots$ in G. Then $\{t'_n : n \in \mathbb{N}\} \cap \lceil \mathring{t} \rceil = \{t'_0, \dots, t'_k\} \subseteq S_t$ for $k < |S_t|$, so $N = \{(t'_0, 0), \dots, (t'_k, k)\} \subseteq S_t \times |S_t|$. □

Theorem 5.6.6. *Let T be an Aronszajn tree and let G be a T-graph with property $(\star)$ as in Proposition 5.6.4. Then the ray inflation $G \sharp \mathbb{N}$ of G witnesses that $\mathrm{HC}(\aleph_1)$ fails.*

Proof. Suppose for a contradiction that for every end ε of $H = G \sharp \mathbb{N}$ there is a set $\mathcal{R} \subseteq \varepsilon$ of disjoint rays with $|\mathcal{R}| = \deg(\varepsilon)$ such that some ray graph of $\mathcal{R}$ in G is connected. By Lemma 5.6.2, all the horizontal rays R_t $(t \in T)$ of $G \sharp \mathbb{N}$ belong to the same sole end ε with $\deg(\varepsilon) = |T| = \aleph_1$. Hence, by our assumption and Lemma 5.3.1 we find an $\aleph_1$-star of rays in $G \sharp \mathbb{N}$ which we denote by S. We denote the centre ray of S by R, and we denote the leaf rays of S by R_i $(i < \omega_1)$.

Now, let $\sigma < \omega_1$ be minimal such that $V(R) \subseteq T^{<\sigma} \times \mathbb{N}$. Since $T^{<\sigma} \times \mathbb{N}$ is countable, we may assume without loss of generality that S meets $T^{<\sigma} \times \mathbb{N}$ precisely in R. Then every component of $S - R$ is contained in a component of $H - (T^{<\sigma} \times \mathbb{N})$. By Lemma 5.6.3, the map $T^\sigma \ni t \mapsto H[\lfloor t \rfloor \times \mathbb{N}]$ is a bijection between T^σ and the components of $H - (T^{<\sigma} \times \mathbb{N})$. Thus, we obtain a map $\omega_1 \to T^\sigma$, $i \mapsto t_i$ such that the component D_i of $S - R$ with $D_i \supseteq R_i$ is contained in the component $C_i := H[\lfloor t_i \rfloor \times \mathbb{N}]$ of $H - (T^{<\sigma} \times \mathbb{N})$. Since the level T^σ of the Aronszajn tree T is countable, some $t \in T^\sigma$ is the image t_i of uncountably many indices $i < \omega_1$, and we abbreviate $C_i =: C$ for these indices. Then for uncountably many indices $i < \omega_1$ of these uncountably many indices their components D_i are contained in $C - (\{t\} \times \mathbb{N})$ entirely. But then all these components D_i have infinite neighbourhood in $\lceil \mathring{t} \rceil \times \mathbb{N}$, contradicting that H has property ($\star\star$) by Lemma 5.6.5. $\square$

We remark that the ray inflation $G \sharp \mathbb{N}$ of any (sparse) ω_1-graph G does contain an $\aleph_1$-star of rays. Indeed, by [30, Proposition 3.5] applied to the minor G of $G \sharp \mathbb{N}$, the ray inflation $G \sharp \mathbb{N}$ contains a $K^{\aleph_1}$ minor, and so in fact a subdivision of $K^{\aleph_1}$ by a well-known result of Jung [51]. In particular, the ray inflations giving counterexamples to HC($\aleph_1$) must be based on Aronszajn trees.

We remark that, based on this example, one can construct counterexamples to Halin's conjecture for all κ with cf $(\kappa) = \aleph_1$, see Theorem 5.9.1.

5.7. On (λ, κ)-graphs II – regular trees with tops

Earlier, in Section 5.4 we investigated conditions under which (λ, κ)-graphs possess $(\aleph_0, \kappa)$-subgraphs, and used these results to prove a number of positive instances of Halin's conjecture in Section 5.5. However, there exist (λ, κ)-graphs without $(\aleph_0, \kappa)$-subgraphs, and the question arises whether these can be turned into counterexamples for HC(κ). And indeed, our main result in this section, Theorem 5.7.1, states that there is such a class of (λ, κ)-graphs, which we call (λ, κ)-graphs of type T_λ, that achieve precisely this, see Theorem 5.8.1 below.

Given a cardinal $\lambda \geq 2$, either finite or of countable cofinality, let $(T_\lambda, \leq)$ be the order tree where the nodes of T_λ are all sequences of elements of λ of length $\leq \omega$ including the empty sequence, and $t \leq t'$ if t is an initial segment of t'. Then T_λ is an order tree of height $\omega + 1$ in which every point of finite height has exactly λ successors and above every branch of $T_\lambda^{<\omega}$ there is exactly one point in T_λ^{ω}, represented by a countable sequence of ordinals in λ. Since $\lambda \geq 2$ is finite or of countable cofinality, it follows from König's Theorem [49, Theorem 5.10] that T_λ has strictly more than λ many branches, that is to say we have $|T_\lambda^{\omega}| \geq \max\{\lambda^+, 2^{\aleph_0}\}$.

The down-closure of any κ-sized subset X of T_λ^{ω} for $\kappa > \lambda$ is a λ-*tree with* κ *tops*, the tops themselves being the points in X. Now if T is a λ-tree with κ tops, then any T-graph clearly contains a (λ, κ)-graph with bipartition classes $T^{<\omega}$ and T^{ω}; we shall call any such graph a (λ, κ)-graph of *type* T_λ.

We remark that λ-trees with κ tops form a generalisation of the so-called *binary trees with tops*, studied in more detail in [6, 30], where it was shown that, consistently, every $(\aleph_0, \aleph_1)$-graph contains an $(\aleph_0, \aleph_1)$-subgraph of type T_2 [6, Theorem 1.1], a statement which does not hold under CH [30, Proposition 8.2].

Our main result in this section, which forms the basis for (2) in Theorem 5.1.2, is the following:

Theorem 5.7.1. *For any singular cardinal λ of countable cofinality there is a (λ, λ^+)-graph of type T_λ that does not have an $(\aleph_0, \lambda^+)$-subgraph.*

Before we proceed to the proof, we also state two consistent results about the number of branches of a certain pair of such trees. For this, recall that GCH implies that every λ-tree T_λ from above for λ of countable cofinality has precisely λ^+ branches [49, Theorem 5.15(ii)], in which case every (λ, κ)-graph of type T_λ satisfies $\kappa = \lambda^+$.

Of course, one way to increase the number of branches of T_λ is to work in a model of ZFC where $2^{\aleph_0}$ is large.

It turns out, however, that for at least two trees, namely $T_{\aleph_\omega}$ and T_μ where $\mu = \aleph_\mu$ denotes the least $\aleph$ fixed point (which is a singular cardinal of countably cofinality, namely $\mu = \sup\{\aleph_0, \aleph_{\aleph_0}, \aleph_{\aleph_{\aleph_0}}, \dots\}$), it is known that the number of branches can consistently be much larger despite the continuum being small, i.e. $2^{\aleph_0} = \aleph_1$. This gives rise to our second result in this section, forming the basis for (4) in Theorem 5.1.2:

Theorem 5.7.2. (1) *For any countable ordinal α with $\omega < \alpha < \omega_1$, it is consistent that we have* CH *and there is an $(\aleph_\omega, \aleph_\alpha)$-graph of type $T_{\aleph_\omega}$ that does not have an $(\aleph_0, \aleph_\alpha)$-subgraph.*

(2) *Let $\mu = \aleph_\mu$ denote the first fixed point of the $\aleph$-function. Then for every $\kappa > \mu$ it is consistent that we have* CH *and there is a (μ, κ)-graph of type T_μ that contains no $(\aleph_0, \kappa)$-subgraph.*

Given the result that the trees $T_{\aleph_\omega}$ and T_μ can consistently have as many branches as needed, the proof of Theorem 5.7.2 is remarkably simple.

Proof. (1) Assuming large cardinals, Gitik and Magidor [40] showed that $T_{\aleph_\omega}$ can consistently have any number of branches of the form $\aleph_{\omega+\alpha+1}$ where $\alpha < \omega_1$, while simultaneously having GCH below $\aleph_\omega$ (i.e. $2^{\aleph_n} = \aleph_{n+1}$ for all $n < \omega$). Now let $\omega < \alpha < \omega_1$ and consider any $(\aleph_\omega, \aleph_\alpha)$-graph of type $T_{\aleph_\omega}$ in such a model, and suppose for a contradiction that it contains an $(\aleph_0, \aleph_\alpha)$-subgraph (A, B). Without loss of generality, B is a set of tops of $T_{\aleph_\omega}$, and $A \subseteq T_{\aleph_\omega}$ is a down-closed subtree. But a countable tree contains at most $2^{\aleph_0} = \aleph_1$ many branches, contradicting that every top in B has infinitely many neighbours in A.

(2) Assuming even larger cardinals, Shelah showed that T_μ can consistently have arbitrarily many branches, while simultaneously having GCH below μ. See again [40]. Now given $\kappa > \mu$, consider any (μ, κ)-graph of type T_μ in such a suitable model in which $|T_\mu^\omega| \geq \kappa$ is sufficiently large. As before, such a graph cannot contain an $(\aleph_0, \kappa)$-subgraph. □

Note that by a similar argument, one can obtain a simple proof of Theorem 5.7.1 under GCH. Perhaps remarkably, however, this assertion holds already in ZFC. Our examples rely on the notion of a *scale* (see also [49, Chapter 24]) that have been developed for Shelah's pcf-theory [66]. Recall that an *ideal* on the natural numbers $\mathbb{N}$ is a proper subset of $\mathcal{P}(\mathbb{N})$ that contains the empty set, is closed under finite unions and is closed under taking subsets of its elements. Thus, it is the dual notion of a filter [49, Chapter 7].

Given an ideal I on $\mathbb{N}$ and two sequences $f, g \colon \mathbb{N} \to \lambda$ of ordinals, we write $f <_I g$ if

$$\{n \in \mathbb{N} \colon f(n) \geq g(n)\} \in I.$$

Definition 5.7.3 (Scales). *Let λ be a singular cardinal of countable cofinality and $\kappa > \lambda$ regular. A κ-*scale *for T_λ is a well-ordered collection $X = (f_\alpha)_{\alpha<\kappa}$ of tops of T_λ for which there are*

- *a strictly increasing sequence $(\lambda_n)_{n\in\mathbb{N}}$ of uncountable regular cardinals with supremum λ satisfying $f_\alpha(n) < \lambda_n$ for all $\alpha < \kappa$, and*
- *an ideal I on $\mathbb{N}$ containing all finite sets,*

such that

(1) *for all $\alpha, \beta < \kappa$ with $\alpha < \beta$ we have $f_\alpha <_I f_\beta$ and*
(2) *for all $g \in \prod_{n\in\mathbb{N}} \lambda_n$ there is $\alpha < \kappa$ such that $g <_I f_\alpha$.*

Proposition 5.7.4. *Let λ be a singular cardinal of countable cofinality. Given any κ-scale X, the corresponding tree T_λ with tops X gives rise to a (λ, κ)-graph that has no $(\aleph_0, \kappa)$-subgraph.*

Proof. Let $X = (f_\alpha)_{\alpha<\kappa}$ be a κ-scale on T_λ. Let T be the corresponding λ-tree with tops X.

Suppose $S \subseteq T$ is a countable subtree of T. Then there is a function $g \in \prod_{n\in\mathbb{N}} \lambda_n$ such that for all $t \in S$ and all n in the domain of t we have $t(n) < g(n)$. Let $\alpha < \kappa$ be such that $g <_I f_\alpha$. Let $\beta < \kappa$ be such that $\alpha < \beta$. Then $g <_I f_\beta$. Consider the set

$$A = \{n \in \mathbb{N} \colon g(n) \geq f_\beta(n)\} \in I.$$

Since I is an ideal, we have $A \subsetneq \mathbb{N}$. Hence, for some $n \in \mathbb{N}$ and all $t \in S$ with n in the domain of t, $f_\beta(n) > t(n)$. It follows that f_β is not a branch of S. This shows that S has at most $|\alpha| < \kappa$ branches in X.

Now let G' be any T-graph and $G \subseteq G'$ the corresponding (λ, κ)-graph with bipartition classes $T^{<\omega}$ and X. If H is an $(\aleph_0, \kappa)$-subgraph (C, D) of G, then without loss of generality $D \subseteq X$ and $C \subseteq T^{<\omega}$, and we can choose a countable subtree S of T such that $C \subseteq S$. Let g and α be as in the argument above for the countable tree S.

Since H is an $(\aleph_0, \kappa)$-graph, there is $\beta < \kappa$ such that $f_\beta \in D$ and $\beta > \alpha$. As before, for some $n \in \mathbb{N}$ and all $t \in S$, whenever n is in the domain of t, then $f_\beta(n) > t(n)$. It follows that for no $m \in \mathbb{N}$ with $m > n$, we have $f_\beta \restriction m \in S$. Hence f_β only has finitely many neighbours in C, contradicting the assumption that H is an $(\aleph_0, \kappa)$-graph. $\square$

Proof of Theorem 5.7.1. The assertion follows from Proposition 5.7.4 together with the well-known result by Shelah that for any singular cardinal λ of countable cofinality, there is a λ^+-scale on T_λ, see [49, Theorem 24.8]. $\square$

5.8. More counterexamples to Halin's conjecture

We are now ready to turn the (λ, κ)-graphs of type T_λ from the previous section into counterexamples for $\mathrm{HC}(\kappa)$.

Theorem 5.8.1. *Let T be a λ-tree with κ tops X, and G any sparse T-graph such that the corresponding (λ, κ)-graph on $(T^{<\omega}, X)$ has no $(\aleph_0, \kappa)$-subgraph. Then the ray inflation $G \sharp \mathbb{N}$ of G witnesses that $\mathrm{HC}(\kappa)$ fails.*

Proof. Suppose for a contradiction that HC(κ) holds. By Lemma 5.6.2, the ray inflation $H = G \sharp \mathbb{N}$ has only one end ε, and $\deg(\varepsilon) = |T| = \kappa$. Then by Lemma 5.3.1 we find a κ'-star S of rays in H with $\lambda < \kappa' \leq \kappa$ where κ' is regular. For ease of notation, let us assume $\kappa = \kappa'$. Denote the centre ray of S by R, and the leaf rays of S by R_i ($i < \kappa$).

Since $T^{<\omega} \times \mathbb{N}$ has size $\lambda < \kappa$, we may assume without loss of generality that each leaf ray R_i is a tail of a horizontal ray $R_{t(i)} \subseteq H$ for a top $t(i) \in T^\omega$. Then the map $i \mapsto t(i)$ is injective. Next, we consider the 'down-closed projection' of the center ray R to T, namely

$$\bar{R} := \lceil \{ t \in T \mid \text{the horizontal ray } R_t \text{ meets } R \} \rceil,$$

a countable subtree of T. We claim that the λ-set $X := (T^{<\omega} \setminus \bar{R}) \times \mathbb{N}$ contains $\kappa > \lambda$ many internal vertices of paths in the path system of S, causing a contradiction. By the choice of T, fewer than κ many tops $t \in T^\omega$ satisfy $N_G(t) \subseteq \bar{R}$. Hence we may assume without loss of generality that each leaf ray R_i is a tail of a horizontal ray $R_{t(i)}$ with $N_G(t(i)) \nsubseteq \bar{R}$ and, in particular, $N_G(t(i)) \cap \bar{R}$ finite. But then for every $i < \kappa$ all but finitely many vertices of the neighbourhood $N_H(R_{t(i)})$ are contained in X. Thus, all paths of the path system of S must have an internal vertex in X, as desired. □

Corollary 5.8.2.

(1) HC(κ) *fails for all* $\kappa \in \{\mu^+ \colon \mathrm{cf}\,(\mu) = \omega\}$.
(2) *For every* $\kappa \in \{\aleph_\alpha \colon \omega < \alpha < \omega_1\}$ *it is consistent that* HC(κ) *fails.*
(3) *Let* μ *denote the first fixed point of the* $\aleph$*-function. Then for every* $\kappa > \mu$ *it is consistent that* HC(κ) *fails.*

Proof. Assertion (1) for $\kappa = \aleph_1$ is Theorem 5.6.6. For the remaining cardinals in (1), it follows from Theorem 5.8.1 together with Theorem 5.7.1. Assertions (2) and (3) follow from Theorem 5.8.1 in combination with Theorem 5.7.2. □

5.9. Lifting counterexamples to singular cardinals

Theorem 5.9.1. *If* HC(κ) *fails for some regular cardinal* κ*, then* HC(λ) *fails for all cardinals* λ *with* $\mathrm{cf}(\lambda) = \kappa$.

Our proof strategy is roughly as follows. Given $\lambda > \kappa$ we consider any graph G with an end ε witnessing that HC(κ) fails. Then we obtain a counterexample $\hat{G}$ for HC(λ) from G as follows. We select any κ many disjoint rays X_i ($i < \kappa$) in ε and consider any sequence $s = (\lambda_i \colon i < \kappa)$ of ordinals $\lambda_i < \lambda$ with supremum λ. Then we obtain $\hat{G}$ from G by adding κ many disjoint λ_i-stars of rays all meeting G precisely in their centre ray X_i. Next, in order to verify that $\hat{G}$ is a counterexample, we assume for a contradiction that HC(λ) holds. Using HC(λ) and Lemma 5.3.1 in $\hat{G}$ we find either a λ-star of rays or a (λ, s)-star of rays, with all rays belonging to the end $\hat{\varepsilon}$ that includes ε. A short argument shows that we cannot get a λ-star of rays. Our aim then is to use the (λ, s)-star of rays in $\hat{G}$ to find a κ-star of rays in G with all rays belonging to ε. An obvious candidate is the κ-star of rays formed by the centre ray and the distributor rays of the (λ, s)-star of rays.

However, the

(1) distributor rays,
(2) paths from the distributor rays to the centre ray, and
(3) the centre ray

need not be included in G. In the remainder of the proof, we adjust our candidate in three steps $i = 1, 2, 3$ so that (i) is included in G at the end of step i.

Proof of Theorem 5.9.1. Suppose that $\mathrm{HC}(\kappa)$ fails for some regular cardinal κ and let $\lambda > \kappa$ be any other cardinal with $\mathrm{cf}(\lambda) = \kappa$. Then $\kappa > \aleph_0$ and there are a graph G with $|G| = \kappa$ and an end ε of G of degree κ such that no degree-witnessing collection of disjoint rays in ε admits a connected ray graph. Let $(X_i : i < \kappa)$ be any family of κ disjoint rays in ε, and let $s = (\lambda_i : i < \kappa)$ be any κ-sequence of ordinals $\lambda_i < \lambda$ with supremum λ. Then we let $\hat{G}$ be the graph obtained from G by adjoining λ many new rays, as follows. For each $i < \kappa$ and $\ell < \lambda_i$ we disjointly add a new ray $X_{(i,\ell)}$ and join the nth vertex of $X_{(i,\ell)}$ to the nth vertex of X_i by an edge for all $n \in \mathbb{N}$. Then the end ε is included in a unique end $\hat{\varepsilon}$ of $\hat{G}$, and $\hat{\varepsilon}$ contains all new rays $X_{(i,\ell)}$ so that it has degree λ. For each $i < \kappa$ we let $\hat{X}_i := \hat{G}[X_i \cup \bigcup_{\ell < \lambda_i} X_{(i,\ell)}]$ be the λ_i-star of rays X_i and $X_{(i,\ell)}$. We claim that $\hat{G}$ and $\hat{\varepsilon}$ witness that $\mathrm{HC}(\lambda)$ fails.

Assume for a contradiction that $\mathrm{HC}(\lambda)$ holds. Then, by Lemma 5.3.1, we find either a λ-star of rays in $\hat{G}$ or a (λ, s)-star of rays in $\hat{G}$, with all rays belonging to $\hat{\varepsilon}$. Let $S \subseteq \hat{G}$ be one of the two possibilities with path system $\mathcal{P}$. We denote the centre ray of S by R, and we denote the leaf rays by R_j $(j < \lambda)$. Since $|G| = \kappa < \lambda$, we may assume without loss of generality that every leaf ray R_j is a tail of a ray $X_{(i,\ell)}$. Since R is countable while $\mathrm{cf}(\lambda) = \kappa$ is uncountable, we may assume without loss of generality that every $\hat{X}_i$ that meets R avoids all the leaf rays of S. Furthermore, we may assume without loss of generality that each $\hat{X}_i$ either avoids all the leaf rays of S or includes uncountably many of them. In summary, for all $i < \kappa$ the λ_i-star $\hat{X}_i$ of rays either avoids all the leaf rays of S, or avoids R while uncountably many leaf rays of S are tails of the rays $X_{(i,\ell)}$. Since every ray X_i is countable, this means that S cannot be a λ-star of rays.

Now that we know that S must be a (λ, s)-star of rays, we denote its distributor rays by R_i $(i < \kappa)$, and we revise our notation for its leaf rays which we now denote by $R_{(i,\ell)}$ $(i < \kappa$ and $\ell < \lambda_i)$ so that R_i is neighboured precisely to the leaf rays $R_{(i,\ell)}$ with $\ell < \lambda_i$ and the centre ray R. Our goal is to find a κ-star of rays in G with all rays belonging to ε. For this, we consider the subset $I \subseteq \kappa$ of all $i < \kappa$ for which $\hat{X}_i$ avoids R and uncountably many leaf rays of S are tails of the rays $X_{(i,\ell)}$. Note that I is a κ-set. For each $i \in I$, the uncountably many relevant leaf rays are neighboured to at most countably many distributor rays: otherwise there exist a leaf ray $R_{(i',\ell')} \subseteq X_{(i,\ell)}$ and a distributor ray $R_j \subseteq \hat{G} - X_i$, implying $R_j \subseteq \hat{G} - X_i - X_{(i,\ell)}$, such that the infinitely many $R_{(i',\ell')}$–R_j paths in $\mathcal{P}$ avoid X_i, contradicting the fact that X_i is equal to the neighbourhood of $X_{(i,\ell)}$ in $\hat{G}$. Thus, we may apply the pigeonhole principle to find, for each $i \in I$, a distributor ray R_j with $j =: j(i)$ that is neighboured to uncountably many leaf rays of S which are tails of rays $X_{(i,\ell)}$ $(\ell < \lambda_i)$. Since each X_i is countable and $R_{j(i)}$ is the centre ray of an $\aleph_1$-star of rays with leaf rays in $\hat{X}_i - X_i$, we deduce that each $R_{j(i)}$ meets X_i infinitely.

By deleting combs from S and deleting paths from $\mathcal{P}$ we may assume without loss of generality that S is a κ-star of rays with centre ray R, leaf rays $R_{j(i)}$ $(i \in I)$ and path system $\mathcal{P}$.

Every component of $S - R$ is a comb whose spine meets G infinitely. Since these combs are disjoint for distinct elements of I and $\hat{X}_i \cap G = X_i$ is countable for all $i < \kappa$, each $\hat{X}_i$ $(i < \kappa)$ meets at most countably many of them. Conversely, each comb meets at most countably many $\hat{X}_i$. Therefore, by deleting elements of I we achieve that

$$\textit{each } \hat{X}_i \ (i < \kappa) \textit{ meets at most one of } R \textit{ and the components of } S - R \tag{5.9.1}$$

while maintaining that I is a κ-set.

For every $i \in I$ we let $\mathcal{P}_i$ be the set of all $R_{j(i)}$–R paths in $\mathcal{P}$. As each $R_{j(i)}$ meets X_i infinitely, X_i is equivalent to R in the subgraph $X_i \cup R_{j(i)} \cup \bigcup \mathcal{P}_i \cup R$, and we find a system $\mathcal{P}'_i$ of infinitely many disjoint X_i–R paths in this subgraph. Then the paths in $\mathcal{P}' := \bigcup_{i \in I} \mathcal{P}'_i$ are independent by the choice of the subgraphs in which we found them and because each X_i meets S only in the component of $S - R$ that contains $R_{j(i)}$ by (5.9.1). Hence, $\mathcal{P}'$ ensures that R is the centre ray of a κ-star S' of rays with leaf rays $X_i \subseteq G$ $(i \in I)$. Then (5.9.1) translates to

$$\textit{each } \hat{X}_i \ (i < \kappa) \textit{ meets at most one of } R \textit{ and the components of } S' - R. \tag{5.9.2}$$

It remains to modify S' so that $S' \subseteq G$. By (5.9.2), the neighbourhood of each component of $S' - R$ in S' is included in $R \cap G$, so every path in $\mathcal{P}'$ ends in R in a vertex of G.

If $P \subseteq \hat{G}$ is a path with endvertices in G, then we write $\pi(P)$ for the subgraph of G that arises from P by replacing each G-path $Q \subseteq P$ with the subpath uX_iv between the endvertices u and v of Q on the ray X_i that contains u and v.

For every $i \in I$, each path $P' \in \mathcal{P}'_i$ starts in a vertex $x \in X_i \subseteq G$ and ends in a vertex $y \in R \cap G$ while avoiding $\hat{X}_i - X_i$ and all $\hat{X}_j - X_j$ $(j < \kappa)$ that are met by R, so $\pi(P')$ contains an X_i–R path P'' that starts in x and ends in y. We claim that, for every $i \in I$ and every path $P'_1 \in \mathcal{P}'_i$, the subgraph $\pi(P'_1)$ meets only finitely many other subgraphs $\pi(P'_2)$ with $P'_2 \in \mathcal{P}'_i$. Indeed, given $i \in I$ let us assume for a contradiction that there are infinitely many paths P' and $P'_0, P'_1, \dots$ in $\mathcal{P}'_i$ such that every subgraph $\pi(P'_n)$ meets $\pi(P')$. Then there is a vertex $v \in \pi(P')$ that lies on infinitely many of the other subgraphs, say all of them. Clearly, v must lie on some ray X_j with $j \neq i$. Since the paths P'_n are disjoint, at most one can contain v, say none. Then each path P'_n contains a G-path Q_n with endvertices x_n and y_n in X_j such that $v \in \mathring{x}_n X_j \mathring{y}_n$. But the initial segment $X_j v$ of the ray X_j up to v is finite and cannot contain the infinitely many distinct vertices x_n, a contradiction. Therefore, we find for every $i < \kappa$ an infinite subset $\mathcal{P}''_i \subseteq \{P'' \colon P' \in \mathcal{P}'_i\}$ of disjoint X_i–R paths all contained in G. By (5.9.2) and the choice of $\mathcal{P}'$, the paths in $\mathcal{P}'' := \bigcup_{i \in I} \mathcal{P}''_i$ are independent. Hence $\mathcal{P}''$ ensures that R is the centre ray of a κ-star S'' of rays with leaf rays X_i $(i \in I)$ all belonging to ε and $S'' - R \subseteq G$.

Now, $S'' - R$ is included in G, but R need not be included in G. However, the endvertices in R of paths in $\mathcal{P}''$ all lie in G, because they inherit this property from $\mathcal{P}'$. Hence, applying Lemma 5.2.1 with $\mathcal{C}$ the collection of combs $X_i \cup \bigcup \mathcal{P}''_i$ $(i \in I)$ and the countable set $U := V(R) \cap V(G)$ in G finishes the proof. □

Corollary 5.9.2.

(1) $\mathrm{HC}(\kappa)$ *fails for all* κ *with* $\mathrm{cf}\,(\kappa) \in \{\mu^+ \colon \mathrm{cf}\,(\mu) = \omega\}$.
(2) *For every* κ *with* $\mathrm{cf}\,(\kappa) \in \{\aleph_\alpha \colon \omega < \alpha < \omega_1\}$ *it is consistent that* $\mathrm{HC}(\kappa)$ *fails.*

Proof. This follows from Theorem 5.9.1 together with Corollary 5.8.2. □

Part III.

Ends of digraphs

6. Ends of digraphs I: basic theory

6.1. Introduction

Ends of graphs are one of the most important concepts in infinite graph theory. They can be thought of as points at infinity to which its rays converge. Formally, an *end* of a graph G is an equivalence class of its rays, where two rays are equivalent if for every finite vertex set $X \subseteq V(G)$ they have a tail in the same component of $G - X$. For example, infinite complete graphs or grids have one end, while the binary tree has continuum many ends, one for every rooted ray [25]. The concept of ends was introduced in 1931 by Freudenthal [35], who defined ends for certain topological spaces. In 1964, Halin [45] introduced ends for infinite undirected graphs, taking his cue directly from Carathéodory's *Primenden* of regions in the complex plane [17].

There is a natural topology on the set of ends of a graph G, which makes it into the *end space* $\Omega(G)$. Polat [60, 61] studied the topological properties of this space. Diestel and Kühn [29] extended this topological space to the space $|G|$ formed by the graph G together with its ends. Many well known theorems of finite graph theory extend to this space $|G|$, while they do not generalise verbatim to infinite graphs. Examples include Nash-William's tree-packing theorem [22], Fleischner's Hamiltonicity theorem [37], and Whitney's planarity criterion [8]. In the formulation of these theorems, topological arcs and circles take the role of paths and cycles, respectively.

For directed graphs, a similarly useful notion and theory of ends has never been found. There have been a few attempts, most notably by Zuther [70], but not with very encouraging results. In this part we propose a new notion of ends of digraphs and develop a corresponding theory of their end spaces. In the first chapter of this part we lay the foundation for the whole part by extending to digraphs a number of techniques that are important in the study of ends of graphs. In order to state the main results of the first chapter of this part more formally, we need a few definitions.

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

For example, the digraph D in Figure 6.1.1 has two ends, which are shown as small dots on the right. Both the upper ray R and the lower ray R' are solid in D because the vertex set of any tail of R or R' is strongly connected in D. Deleting finitely many vertices of D always results in precisely two infinite strong components (and finitely many finite strong components) spanned by the vertex sets of tails of R or R'.

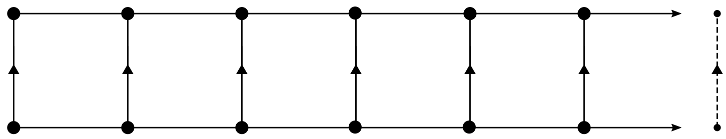

Figure 6.1.1.: A digraph with two ends (depicted as small dots) linked by a limit edge (depicted as a dashed line). Every undirected edge in the figure represents a pair of inversely directed edges.

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

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

For undirected graphs G one often needs to know whether an end ω is in the closure of a given vertex set U, i.e., whether U meets $C(X, \omega)$ for every finite vertex set $X \subseteq V(G)$. This is equivalent to G containing a comb with all its teeth in U. Recall that a *comb* is the union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R. The last vertices of those paths are the *teeth* of this comb. A standard tool in this context is the star-comb lemma [25, Lemma 8.2.2] which states that a connected graph contains for a given set U of vertices either a comb with all its teeth in U or an infinite subdivided star with all its leaves in U. In this first chapter we will prove a directed version of the star-comb lemma.

Call two statements A and B *complementary* if the negation of A is equivalent to B. For a graph G, the statement that G has an end in the closure of $U \subseteq V(G)$ is complementary to the statement that G has a U-rank, see [12]. For $U = V(G)$, the U-rank is known as Schmidt's ranking of rayless graphs [25, 65]. It is a standard technique to prove statements about rayless graphs by transfinite induction on Schmidt's rank. For example Bruhn, Diestel, Georgakopoulos, and Sprüssel [9] employed this technique to prove the unfriendly partition conjecture for countable rayless graphs.

The directed analogue of a comb with all its teeth in U will be a 'necklace' attached to U. The *symmetric ray* is the digraph obtained from an undirected ray by replacing each of its edges by its two orientations as separate directed edges. A *necklace* is an inflated symmetric ray with finite branch sets. (An inflated H is obtained from a digraph H by subdividing some edges of H finitely often and then replacing the 'old' vertices by strongly connected digraphs. The *branch sets* of the inflated H are these strongly connected digraphs. See Section 6.2 for the formal definition of inflated, and of branch sets.) Figure 6.1.2 shows an example of a necklace.

Given a set U of vertices in a digraph D, a necklace $N \subseteq D$ is *attached to* U if infinitely many of the branch sets of N contain a vertex from U. We will see that the statement

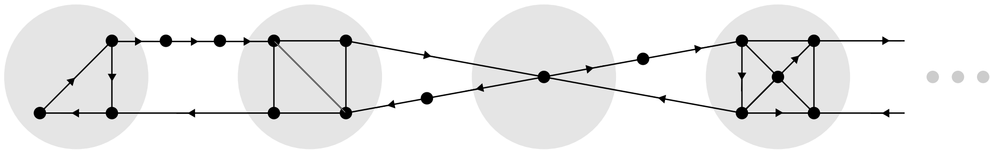

Figure 6.1.2.: A necklace up to the fourth branch set. Every undirected edge in the figure represents a pair of inversely directed edges.

that D has an end in the closure of U is equivalent to the statement that D contains a necklace attached to U as a subdigraph.

We extend Schmidt's result that a graph is rayless if and only if it has a rank. See Section 6.3 for the definition of 'U-rank' in digraphs.

Lemma 6.3.5 (Necklace Lemma). *Let D be any digraph and U any set of vertices in D. Then the following statements are complementary:*

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

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

Theorem 6.1. *Let D be any infinite digraph. The map $\omega \mapsto f_\omega$ with domain $\Omega(D)$ is a bijection between the ends and the vertex-directions of D.*

While most of the concepts that we investigate have undirected counterparts, there is one important exception: limit edges. If ω and η are distinct ends of a digraph, there exists a finite vertex set $X \in \mathcal{X}(D)$ such that ω and η live in distinct strong components of $D - X$. Let us say that such a vertex set X *separates* ω and η. For two distinct ends $\omega, \eta \in \Omega(D)$ we call the pair (ω, η) a *limit edge* from ω to η if D has an edge from $C(X, \omega)$ to $C(X, \eta)$ for every finite vertex set X that separates ω and η.

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

The digraph in Figure 6.1.1 has a limit edge from the lower end to the upper end, and the digraph in Figure 6.1.3 has a limit edge from the lower vertex to the unique end. Let us enumerate from left to right the vertical edges $e_0, e_1, \ldots$ of the digraph D in Figure 6.1.1. We may think of the e_n as converging towards the unique limit edge. This will be made precise in the second chapter of this part.

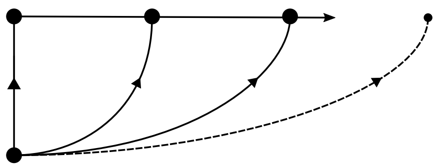

Figure 6.1.3.: A digraph with one end (depicted as a small dot) and a limit edge (depicted as a dashed line) from the lower vertex to the end. Every undirected edge in the figure represents a pair of inversely directed edges.

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

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

Theorem 6.2. *Let D be any infinite digraph. The map $\lambda \mapsto f_\lambda$ with domain $\Lambda(D)$ is a bijection between the limit edges and the edge-directions of D.*

This chapter is organised as follows. In Section 6.2 we provide the basic terminology that we use throughout this chapter. In Section 6.3 we prove the necklace lemma and discuss some basic properties of ends of digraphs. In Section 6.4 we prove Theorem 6.1. Finally, in Section 6.5 we investigate limit edges and prove Theorem 6.2.

6.2. Preliminaries

Any graph-theoretic notation not explained here can be found in Diestel's textbook [25]. For the sake of readability, we sometimes omit curly brackets of singletons, i.e., we write x instead of $\{x\}$ for a set x. Furthermore, we omit the word 'directed'—for example in 'directed path'—if there is no danger of confusion.

Throughout this chapter D is an infinite digraph without multi-edges and without loops, but which may have inversely directed edges between distinct vertices. For a digraph D, we write $V(D)$ for the vertex set of D, we write $E(D)$ for the edge set of D

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

and $\mathcal{X}(D)$ for the set of finite vertex sets of D. We write edges as ordered pairs (v, w) of vertices $v, w \in V(D)$, and we usually write (v, w) simply as vw. The *reverse* of an edge vw is the edge wv. More generally, the *reverse* of a digraph D is the digraph on $V(D)$ where we replace every edge of D by its reverse, i.e., the reverse of D has the edge set $\{ vw \mid wv \in E(D) \}$. A *symmetric* path is a digraph obtained from an undirected path by replacing each of its edges by its two orientations as separate directed edges. Similarly, a *symmetric ray* is a digraph obtained from an undirected ray by replacing each of its edges by its two orientations as separate directed edges. Hence the reverse of any symmetric path or symmetric ray is a symmetric path or symmetric ray, respectively.

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

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

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

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

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

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

An inflated H is any digraph that arises from a subdivision H' of H as follows. Replace every branch vertex v of H' by a strongly connected digraph H_v so that the H_v are disjoint and do not meet any subdividing vertex; here replacing means that we first delete v from H' and then add $V(H_v)$ to the vertex set and $E(H_v)$ to the edge set. Then replace every subdividing path P_{vw} that starts in v and ends in w by an H_v–H_w path that coincides with P_{vw} on inner vertices. We call the vertex sets $V(H_v)$ the *branch sets* of the inflated H. A *necklace* is an inflated symmetric ray with finite branch sets; the branch sets of a necklace are its *beads*. (See Figure 6.1.2 for an example of a necklace.)

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

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

is its supremum and infimum, respectively.

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

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

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

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

6.3. Necklace Lemma

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

The star-comb lemma [25] for undirected graphs is a standard tool in infinite graph theory and reads as follows:

Lemma 6.3.1 (Star-Comb Lemma)**.** *Let U be an infinite set of vertices in a connected undirected graph G. Then G contains a comb with all its teeth in U or a subdivided infinite star with all its leaves in U.*

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

Lemma 6.3.2 (Directed Star-Comb Lemma)**.** *Let D be any strongly connected digraph and let $U \subseteq V(D)$ be infinite. Then D contains a star or comb attached to U and a reverse star or reverse comb attached to U sharing their attachment sets.*

Proof. Since D is strongly connected we find a spanning arborescence T. Applying the star-comb lemma in the undirected tree underlying T yields a comb or a star attached to U (without loss of generality the spine starts in the root of T). Let U' be the attachment set in either case.

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

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

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

The directed star-comb lemma however does not manage the task of finding an end of a digraph in the closure of an infinite set of vertices. Consider for example the digraph D that is obtained from the digraph in Figure 6.1.1 by subdividing each vertical edge once. We write U for the set of subdividing vertices. As D contains neither an infinite star nor an infinite reverse star, the directed star-comb lemma applied to U returns a comb attached to U and a reverse comb attached to U sharing their attachment sets.

Therefore we would expect that the ends that are represented by the spines are contained in the closure of U. But U does not have any end in its closure because the subdividing vertices all lie in singleton strong components of D.

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

We remark that our notion of ranking extends the notion of Schmidt's *ranking of rayless graphs*, in that the rank of a given undirected graph G is precisely the rank of the digraph obtained from G by replacing every edge by its two orientations as separate directed edges, see [65] or Chapter 8.5 of [25] for Schmidt's rank. More generally, for a set U, our U-rank of digraphs extends the notion of the U-rank of graphs, in that an undirected graph G has a U-rank if and only if the digraph that is obtained from G by replacing every edge by its two orientations as separate directed edges has a U-rank; see [12] for the definition of the U-rank of an undirected graph.

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

Lemma 6.3.3. *Let D be a digraph and let $\mathcal{U}$ be a finite set. If D has $\mathcal{U}$-rank α and $H \subseteq D$, then H has some $\mathcal{U}$-rank $\leq \alpha$.*

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

Lemma 6.3.4. *Let D be any digraph and let $\mathcal{U}$ be a finite set. If D has a $\mathcal{U}$-rank $\alpha > 0$ and $X \subseteq V(D)$ is a finite vertex set such that every strong component of $D - X$ has a $\mathcal{U}$-rank $< \alpha$, then infinitely many strong components of $D - X$ meet every set in $\mathcal{U}$.*

Proof. Suppose for a contradiction that the set $\mathcal{C}$ of strong components of $D - X$ that meet every set in $\mathcal{U}$ is finite. We find for every $C \in \mathcal{C}$ a finite vertex set $X_C \subseteq V(C)$ witnessing that C has a $\mathcal{U}$-rank $< \alpha$. Let Y be the union of X with all the finite vertex sets X_C. Then Y witnesses that D has a $\mathcal{U}$-rank $< \alpha$ contradicting our assumption that D has $\mathcal{U}$-rank α. □

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

Lemma 6.3.5 (Necklace Lemma). *Let D be any digraph and $\mathcal{U}$ a finite set of vertex sets of D. Then the following statements are complementary:*

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

Proof. Let us start by showing that not both statements hold at the same time. Suppose for a contradiction there is a digraph D that has a $\mathcal{U}$-rank and contains a necklace attached to $\mathcal{U}$ as a subdigraph. Then, by Lemma 6.3.3, every necklace $N \subseteq D$ has a $\mathcal{U}$-rank. But deleting finitely many vertices from any necklace attached to $\mathcal{U}$ leaves a strong component that is a necklace attached to $\mathcal{U}$ by its own. Hence choosing a necklace $N \subseteq D$ attached to $\mathcal{U}$ with minimal $\mathcal{U}$-rank results in a contradiction.

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

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

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

As our first application of the necklace lemma we describe the connection between Zuther's notion of ends from [70], which we call pre-ends, with our notion of ends. Two rays or reverse rays $R_1, R_2 \subseteq D$ are *equivalent*, if there are infinitely many disjoint paths from R_1 to R_2 and infinitely many disjoint paths from R_2 to R_1. We call the equivalence classes of this relation the *pre-ends* of D.

Lemma 6.3.6. *Let D be any digraph and γ a pre-end of D. Then γ includes an end ω of D if and only if γ is represented both by a ray and a reverse ray. Moreover, ω is the unique end of D included in γ.*

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

So suppose for a contradiction that there is no such necklace. Then by the necklace lemma applied to U in D, the digraph D has a U-rank, say α. Let $X \subseteq V(D)$ be a finite vertex set that witnesses that the U-rank of D is α. As $U \subseteq V(D)$ is infinite, we have

$\alpha > 0$. Now, it follows by Lemma 6.3.4 that the ray R meets infinitely many strong components of $D - X$. We conclude that R has no tail in any strong component of $D - X$ contradicting that R is solid in D.

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

The subdigraph H of D that consists of R, R' and all the paths in $\mathcal{P}$ has exactly one infinite strong component and finitely many finite strong components (possibly none). Moreover, deleting finitely many vertices from H results again in exactly one infinite strong component and finitely many finite strong components. Consequently, R has a tail that is contained in a strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$.

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

Our second application of the necklace lemma demonstrates how the rank can be used to prove statements about digraphs that have no end. A set of vertices of a digraph D is *acyclic* in D if its induced subdigraph does not contain a directed cycle. The *dichromatic number* [58] of a digraph D is the smallest cardinal κ so that D admits a vertex partition into κ partition classes that are acyclic in D. As a consequence of the necklace lemma we obtain a sufficient condition for D to have a countable dichromatic number:

Theorem 6.3.7. *If D is a digraph that contains no necklace as a subdigraph, then the dichromatic number of D is countable.*

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

6.4. Directions

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

We call a direction f of D a *vertex-direction* if $f(X)$ is a strong component of $D - X$ for every $X \in \mathcal{X}(D)$.

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

Theorem 6.1. *Let D be any infinite digraph. The map $\omega \mapsto f_\omega$ with domain $\Omega(D)$ is a bijection between the ends and the vertex-directions of D.*

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

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

Lemma 6.4.1. *Let D be any digraph, and let $\mathcal{U}$ be a finite set of vertex sets of D. Then the following assertions are equivalent:*

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

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

(ii)→(iii): Suppose that f is a vertex-direction in the closure of $\mathcal{U}$. We need to find a necklace attached to $\mathcal{U}$. By the necklace lemma we may equivalently show that D has no $\mathcal{U}$-rank. Suppose for a contradiction that D has a $\mathcal{U}$-rank α. By Lemma 6.3.3 subdigraphs of digraphs that have a $\mathcal{U}$-rank have a $\mathcal{U}$-rank and thus we may choose X' such that $f(X')$ has the smallest $\mathcal{U}$-rank among all $f(X)$ with $X \in \mathcal{X}(D)$. Note that $f(X)$ has $\mathcal{U}$-rank ≥ 1 for every $X \in \mathcal{X}(D)$. Indeed, if $f(X) \cap U$ is finite for some $U \in \mathcal{U}$, then

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

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

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

Let D be any digraph and let f be any vertex-direction of D. We think of a separation (A, B) of D as pointing towards its side B. Now, if (A, B) is a finite order separation of

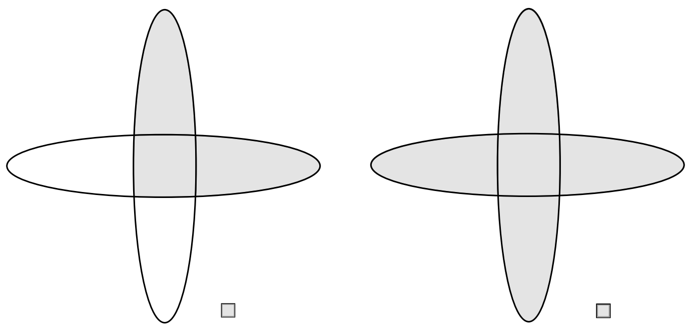

Figure 6.4.1.: The separations (A, B) and (A, B') from the proof of Lemma 6.4.2.

D, then $f(A \cap B)$ is either included in $B \setminus A$ or $A \setminus B$. In the first case we say that (A, B) *points towards* f and in the second case we say that (A, B) *points away from* f. Note that the supremum or infimum of two finite order separations is again a finite order separation. If two separations point towards or away from f, then the same is true for their supremum or infimum, respectively:

Lemma 6.4.2. *Let D be any digraph and let f be a vertex-direction of D. Suppose that (A_1, B_1) and (A_2, B_2) are finite order separations of D.*

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

Proof. (i) We have to show that for $(A, B) := (A_1 \cup A_2, B_1 \cap B_2)$ and $X = A \cap B$ the strong component $f(X)$ is included in $B \setminus A$. For this let us consider the auxiliary separation $(A, B') := (A, X' \cup B)$, where $X' := \bigcup_{i=1,2} A_i \cap B_i$ (cf. Figure 6.4.1). Recall that the separator of a separation separates its two sides. Hence the vertex set $B \setminus A$ is partitioned into the strong components of $D - X$ that it meets.

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

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

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

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

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

(ii) Apply (i) to the reverse of D. $\square$

Recall, that for a given undirected graph G a vertex v is said to *dominate* an end ω of G if there is an infinite v–R fan in G for some (equivalently every) ray R that represents ω. Equivalently v dominates ω if v is contained in $C(X, \omega)$ for every finite vertex set $X \subseteq V(G) \smallsetminus \{v\}$. An end $\omega \in \Omega(G)$ is *dominated* if some vertex of G dominates it. Ends not dominated by any vertex of G are *undominated*, see [25]. The main case distinction in the proof of Diestel and Kühn's theorem [28, Theorem 2.2], which states that the ends of an undirected graph correspond bijectively to its directions, essentially distinguishes between directions that correspond to dominated ends and those that correspond to undominated ends. Our plan is to make a similar case distinction for which we need a concept of domination for ends of digraphs.

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

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

Proposition 6.4.3. *Let D be any digraph and ω an end of D. A vertex (reverse) dominates ω if and only if it (reverse) dominates f_ω.*

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

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

Lemma 6.4.4. *Let D be any strongly connected digraph and let f be any vertex-direction of D. Then the following assertions are complementary:*

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

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

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

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

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

Now, assume that (A_n, B_n) has already been defined. Since no vertex dominates f we find for every $x \in A_n \cap B_n$ a separation (A_x, B_x) pointing away from f such that $x \in B_x \setminus A_x$. Letting (A_{n+1}, B_{n+1}) be the infimum of all the (A_x, B_x) and (A_n, B_n) completes the construction. Indeed, (A_{n+1}, B_{n+1}) points away from f by Lemma 6.4.2 and its separator is disjoint from all the previous ones as $A_{n+1} \cap B_{n+1} \subsetneq A_n \setminus B_n$.

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

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

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

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

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

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

Proof of Theorem 6.1. It is straightforward to show that the map $\omega \mapsto f_\omega$ with domain $\Omega(D)$ and codomain the set of vertex-directions of D is injective; we prove that it is onto. So given a vertex-direction f of D we need to find an end $\omega \in \Omega(D)$ such that $f_\omega = f$. Let S_1^* be the set of all the vertices that dominate f and S_2^* the set of all the vertices that reverse dominate f. We split the proof into three cases:

First, assume that both $S_1^* \cap f(X)$ and $S_2^* \cap f(X)$ are non-empty for every $X \in \mathcal{X}(D)$. Then f is a vertex-direction in the closure of $\mathcal{U}$ for $\mathcal{U} := \{S_1^*, S_2^*\}$. By Lemma 6.4.1 we find an end ω in the closure of $\mathcal{U}$ and we claim that $f_\omega = f$. Indeed, given $X \in \mathcal{X}(D)$ we need to show that $C(X,\omega) = f(X)$. We may view the strong components of $D - X$ partially ordered by $C_1 \leq C_2$ if there is a path in $D - X$ from C_1 to C_2. By the order-extension-principle we choose a linear extension of $\leq$. Let

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

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

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

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

Second, suppose that $S_1^* \cap f(X)$ is empty for some $X \in \mathcal{X}(D)$. If even S_1^* is empty and D strongly connected, then Lemma 6.4.4 and Lemma 6.4.1 do the rest: with U as in the 'moreover' part of Lemma 6.4.4, we have that f is the unique vertex-direction in the closure of U and Lemma 6.4.1 yields an end ω in the closure of U; by uniqueness $f_\omega = f$.

In the following we will argue that we may assume S_1^* to be empty and D to be strongly connected. Fix $X' \in \mathcal{X}(D)$ with $S_1^* \cap f(X') = \emptyset$. Let $D' = f(X')$ and let f' be the vertex-direction of D' induced by f, i.e., f' sends a finite vertex set $X \subseteq V(D')$ to $f(X \cup X')$. Then the set of all the vertices that dominate f' is empty: If (A, B) is a finite order separation of D that points away from f, then $(A \cap V(D'), B \cap V(D'))$ is a finite order separation of D' that points away from f'. As a consequence, any vertex from D' that dominates f' also dominates f, which means there is none.

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

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

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

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

6.5. Limit edges and edge-directions

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

We begin this section with two propositions (Proposition 6.5.1 and Proposition 6.5.2) saying that limit edges are witnessed by subdigraphs that are essentially the digraphs in Figure 6.1.1 or Figure 6.1.3. Subsequently we prove Theorem 6.2.

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

Proposition 6.5.1. *For a digraph D and two distinct ends ω and η of D the following assertions are equivalent:*

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

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

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

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

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

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

There is a natural partial order on the set of ends, where $\omega \leq \eta$ if for every two rays R_ω and R_η that represent ω and η, respectively, there are infinitely many pairwise disjoint paths from R_ω to R_η. By Proposition 6.5.1 we have that $\omega \leq \eta$, whenever $\omega\eta$ is a limit edge for ends ω and η. The converse of this is in general false, for example in the digraph that is obtained from the digraph in Figure 6.1.1 by subdividing every vertical edge once.

Proposition 6.5.2. *For a digraph D, a vertex v and an end ω of D the following assertions are equivalent:*

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

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

For the forward implication (i)→(ii) a similar recursive construction as in the proof of Proposition 6.5.1 yields a necklace N as desired.

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

As a consequence of this proposition, every vertex $v \in V(D)$ for which D has a limit edge from v to an end $\omega \in \Omega(D)$ dominates ω. The converse of this is in general false, for example in the digraph that is obtained from the digraph in Figure 6.1.3 by subdividing every edge once. Similarly, if ωv is a limit edge between an end ω and a vertex v, then v reverse dominates ω; the converse is again false in general.

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

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

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

Theorem 6.2. *Let D be any infinite digraph. The map $\lambda \mapsto f_\lambda$ with domain $\Lambda(D)$ is a bijection between the limit edges and the edge-directions of D.*

Proof. It is straightforward to show that the map given in (ii) is injective; we prove onto. So let f be any edge-direction of D. First suppose that $f(X)$ is always a strong component or a bundle between strong components for every $X \in \mathcal{X}(D)$. Then f defines two vertex-directions f_1 and f_2 as follows. If $f(X) = E(C_1, C_2)$ is a bundle then let $f_1(X) = C_1$ and $f_2(X) = C_2$. Otherwise, $f(X)$ is a strong component and we put $f_1(X) = f_2(X) = f(X)$. Now, the inverse of the function from Theorem 6.1 returns ends ω and η for f_1 and f_2, respectively. We conclude that $\omega\eta$ is a limit edge and that $f = f_{\omega\eta}$.

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

Let us define a vertex-direction f' of D. First, for every $X \in \mathcal{X}(D)$ with $v \in X$ we have that $f(X)$ is a bundle of the form $E(v, C)$ for a strong component C of $D - X$ and we put $f'(X) = C$. Second, if $v \notin X$ for some $X \in \mathcal{X}(D)$ we have that $f(X)$ is either a strong component C' of $D - X$ or a bundle $E(C, C')$ with $v \in C$. We then put $f'(X) = C'$. It is straightforward to check that f' is indeed a vertex-direction. Finally, the inverse of the map from Theorem 6.1 applied to f' returns an end ω. By the definition of f' we have that $v\omega$ is a limit edge of D and a close look to the definitions involved points out that $f = f_{v\omega}$. □

7. Ends of digraphs II: the topological point of view

7.1. Introduction

In 2004, Diestel and Kühn [29] introduced a topological framework for infinite graphs which makes it possible to extend theorems about finite graphs to infinite graphs that do not generalise verbatim. The main point is to consider not only the graph itself but the graph together with its ends, and to equip both together with a suitable topology. For locally finite graphs G, this space $|G|$ coincides with the Freudenthal compactification of G [25, 26].

Diestel and Kühn's approach has become standard and has lead to several results found by various authors. Examples include Nash-William's tree-packing theorem [22], Fleischner's Hamiltonicity theorem [37], and Whitney's planarity criterion [8]. In the formulation of these theorems, topological arcs and circles take the role of paths and cycles, respectively.

To illustrate this, consider Euler's theorem that a connected finite graph contains an Euler tour if and only if every vertex has even degree. This statement fails for infinite graphs, since a closed walk in a connected infinite graph cannot traverse all its infinitely many edges. Diestel and Kühn [29] extended Euler's Theorem to the space $|G|$ for locally finite graphs G, as follows. A *topological Euler tour* of $|G|$ is a continuous map $\sigma\colon S^1 \to |G|$ such that every inner point of an edge of G is the image of exactly one point of S^1. Hence a topological Euler tour 'traverses' every edge exactly once. Diestel and Kühn showed that for a connected locally finite graph G the space $|G|$ admits a topological Euler tour if and only if every finite cut in G is even. Note that in a finite graph every vertex has even degree if and only if all finite cuts are even, but even for locally finite infinite graphs the latter statement is stronger. Theorem 7.3 below is a directed analogue of the Diestel-Kühn theorem about topological Euler tours.

In the first chapter of this part we introduced the concept of *ends* and *limit edges* of a digraph. A *directed ray* is an infinite directed path that has a first vertex (but no last vertex). The directed subrays of a directed ray are its *tails*. For the sake of readability we shall omit the word 'directed' in 'directed path' and 'directed ray' if there is no danger of confusion. We call a ray in a digraph *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. We call two solid rays in a digraph D *equivalent* if for every finite vertex set $X \subseteq V(D)$ they have a tail in the same strong component of $D - X$. The classes of this equivalence relation are the *ends* of D. The ends of a digraph can be thought of as points at infinity to which its solid rays converge.

For limit edges the situation is similar. Informally, they are additional edges that naturally arise between distinct ends of a digraph, as follows. Unlike graphs, digraphs may have two rays R and R' that represent distinct ends and yet there may be a set E of infinitely many independent edges from R to R'. In this situation there will be a limit edge from the end represented by R to the end represented by R', and the edges in E can be thought of as converging towards this limit edge. We will recall the precise definition of limit edges in Section 7.2.

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

For graphs G, a necessary condition for $|G|$ to be compact is that $G - X$ has only finitely many components for every finite vertex set $X \subseteq V(G)$: if $G - X$ has infinitely many components, then these components together with all ends living in them will form a disjoint family of open sets, and combining this family with a suitable cover of the finite graph $G[X]$ yields an open cover of $|G|$ that has no finite subcover. In [21], Diestel proves that this necessary condition is also sufficient. In analogy to this, let us call a digraph D *solid* if $D - X$ has only finitely many strong components for every finite vertex set $X \subseteq V(D)$.

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

Theorem 7.1. *The space $|D|$ is compact if and only if D is solid.*

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

A common way to generalise statements about finite graphs to infinite graphs is to use so-called compactness arguments. These can be phrased in terms of inverse limits. For example, the ray given by Kőnig's infinity lemma [25, Lemma 8.1.2] exists because the inverse limit of compact discrete spaces is non-empty. In order to make this technique applicable for the space $|D|$ we provide the following:

Theorem 7.2. *For a solid digraph D the space $|D|$ is the inverse limit of finite contraction minors of D.*

For the precise statement of this theorem see Section 7.4.

Recall that our motivation for introducing a topology on a digraph D, together with its ends and limit edges, was to extend to the space $|D|$ theorems about finite digraphs that would be either false, or trivial, or undefined for D itself. As a proof of concept, we prove two such applications for $|D|$. Two further applications can be found in the fourth chapter of this part, see also [57].

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

As in the case of undirected graphs, however, there is a natural topological notion of Euler tours of $|D|$. Call a continuous map $\alpha\colon [0,1] \to |D|$ that respects the direction of the edges of $|D|$ a *topological path in* $|D|$, which is *closed* if $\alpha(0) = \alpha(1)$. Call a closed topological path an *Euler tour* if it traverses every edge exactly once, and call $|D|$ *Eulerian* if it admits an Euler tour. See Section 7.5 for precise definitions.

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

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

Any unbalanced finite cut in a digraph D is an obstruction that prevents $|D|$ from being Eulerian: by the directed jumping arc lemma (Lemma 7.5.1), any Euler tour enters a side of a finite cut as often as it leaves it. A second obstruction is a vertex of infinite in- or out-degree, as an Euler tour that traverses a vertex infinitely often forces the tour to converge to that vertex (see the proof of Theorem 7.3 for details).

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

Theorem 7.3. *For a digraph D with a connected underlying graph the following assertions are equivalent:*

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

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

Theorem 7.4. *For a countable solid digraph D the following assertions are equivalent:*

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

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

This chapter is organised as follows. In Section 7.2 we collect together the results that we need from the first chapter of this part [13], or from general topology. In Section 7.3 we formally define the topological space $|D|$ for a given digraph D and prove Theorem 7.1. In Section 7.4 we define an inverse system for a given digraph D and show that the inverse limit of this system coincides with $|D|$ if D is solid (Theorem 7.2). Finally, in Section 7.5 we prove our two applications of our framework, Theorem 7.3 and Theorem 7.4.

7.2. Tools and terminology

In this section we provide the tools and terminology that we use throughout this chapter. Any graph-theoretic notation not explained here can be found in Diestel's textbook [25]. For the sake of readability, we sometimes omit curly brackets of singletons, i.e., we write x instead of $\{x\}$ for a set x. Furthermore, we omit the word 'directed'—for example in 'directed path'—if there is no danger of confusion.

First, we list some tools and terminology that we adopt from the first chapter of this part. If not stated otherwise we consider digraphs without parallel edges and without loops. For a digraph D we write $V(D)$ for the vertex set of D, we write $E(D)$ for the edge set of D and $\mathcal{X} = \mathcal{X}(D)$ for the set of finite vertex sets of D. We write edges as order pairs (v, w) of vertices $v, w \in V(D)$, and usually we write (v, w) simply as vw; except if D is a multi-digraph in which case we write edges of D as triples (e, v, w). The vertex v is the *tail* of vw and the vertex w its *head*.

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

The *reverse* of an edge vw is the edge wv. More generally, the *reverse* of a digraph D is the digraph on $V(D)$ where we replace every edge of D by its reverse, i.e., the reverse of D has the edge set $\{\, vw \mid wv \in E(D)\,\}$. A *symmetric ray* is a digraph obtained from an undirected ray by replacing each of its edges by its two orientations as separate directed edges. Hence the reverse of a symmetric ray is a symmetric ray.

A *directed ray* is an infinite directed path that has a first vertex (but no last vertex). The directed subrays of a ray are its *tails*. Call a ray *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. Two solid rays in D are *equivalent*, if they have a tail in the same strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. We call the equivalence classes of this relation the *ends* of D and we write $\Omega(D)$ for the set of ends of D.

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

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

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

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

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

Theorem 7.2.1 ([Theorem 6.1]). *Let D be any infinite digraph. The map $\omega \mapsto f_\omega$ with domain $\Omega(D)$ is a bijection between the ends and the vertex-directions of D.*

For two distinct ends ω and η of a digraph D, we call the pair (ω, η) a *limit edge from ω to η*, if D has an edge from $C(X,\omega)$ to $C(X,\eta)$ for every finite vertex set $X \subseteq V(D)$ that separates ω and η.

For a vertex $v \in V(D)$ and an end $\omega \in \Omega(D)$ call the pair (v, ω) a *limit edge from v to ω* if D has an edge from v to $C(X,\omega)$ for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X,\omega)$. Similarly, we call the pair (ω, v) a *limit edge from ω to v* if D has an edge from $C(X,\omega)$ to v for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X,\omega)$. We denote by $\Lambda(D)$ the set of all the limit edges of D and we use the usual definitions for edges accordingly; for example we will speak of the head and the tail of a limit edge. Every limit edge λ defines an edge-direction as follows. We say that a limit edge $\lambda = \omega\eta$ *lives* in the bundle defined by $E(X,\lambda)$ if $X \in \mathcal{X}(D)$ separates ω and η. If $X \in \mathcal{X}(D)$ does not separate ω and η, we say that $\lambda = \omega\eta$ *lives* in the strong component $C(X,\omega) = C(X,\eta)$ of $D - X$. We use similar notations for limit edges of the form $\lambda = v\omega$ or $\lambda = \omega v$ with $v \in V(D)$ and $\omega \in \Omega(D)$: We say that the limit edge λ *lives in* the bundle $E(X,\lambda)$ if $x \notin C(X,\omega)$ and we say that λ *lives in* the strong component $C(X,\omega)$ of $D - X$, if $v \in C(X,\omega)$. The edge-direction f_λ defined by λ is the edge-direction that maps every finite vertex set $X \in \mathcal{X}(D)$ to the bundle or strong component of $D - X$ in which λ lives. The limit edges of any digraph correspond bijectively to its edge-directions:

Theorem 7.2.2 ([Theorem 6.2]). *Let D be any infinite digraph. The map $\lambda \mapsto f_\lambda$ with domain $\Lambda(D)$ is a bijection between the limit edges and the edge-directions of D.*

Let H be any fixed digraph. A *subdivision* of H is any digraph that is obtained from H by replacing every edge vw of H by a path P_{vw} with first vertex v and last vertex w

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

so that the paths P_{vw} are internally disjoint and do not meet $V(H) \smallsetminus \{v, w\}$. We call the paths P_{vw} *subdividing* paths. If D is a subdivision of H, then the original vertices of H are the *branch vertices* of D and the new vertices its *subdividing vertices*.

An inflated H is any digraph that arises from a subdivision H' of H as follows. Replace every branch vertex v of H' by a strongly connected digraph H_v so that the H_v are disjoint and do not meet any subdividing vertex; here replacing means that we first delete v from H' and then add $V(H_v)$ to the vertex set and $E(H_v)$ to the edge set. Then replace every subdividing path P_{vw} that starts in v and ends in w by an H_v–H_w path that coincides with P_{vw} on inner vertices. We call the vertex sets $V(H_v)$ the *branch sets* of the inflated H. A *necklace* is an inflated symmetric ray with finite branch sets; the branch sets of a necklace are its *beads*. Note that if a digraph D contains a necklace N, then every ray in N is solid in D and any two such rays are equivalent. Hence all rays in N represent a unique end of D. With a slight abuse of notation, we say that a necklace $N \subseteq D$ *represents* an end ω of D if one (equivalently every) ray in N represents ω. For limit edges we have the following:

Proposition 7.2.3 ([Proposition 6.5.1])**.** *For a digraph D and two distinct ends ω and η of D the following assertions are equivalent:*

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

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

Proposition 7.2.4 ([Proposition 6.5.2])**.** *For a digraph D, a vertex v and an end ω of D the following assertions are equivalent:*

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

For a digraph D and a set $\mathcal{U}$ we say that a necklace $N \subseteq D$ is *attached to* $\mathcal{U}$ if infinitely many beads of N meet every set of $\mathcal{U}$. In the first chapter of this part we introduced an ordinal rank function that can be used to find out whether a digraph D contains for a given set $\mathcal{U}$ a necklace attached to $\mathcal{U}$. For this, consider a finite set $\mathcal{U}$ and think of $\mathcal{U}$ as consisting of infinite sets of vertices. We define in a transfinite recursion the class of digraphs that have a *$\mathcal{U}$-rank*. A digraph D has *$\mathcal{U}$-rank* 0 if there is a set $U \in \mathcal{U}$ such that $U \cap V(D)$ is finite. It has *$\mathcal{U}$-rank* α if it has no $\mathcal{U}$-rank $< \alpha$ and there is some $X \in \mathcal{X}(D)$ such that every strong component of $D - X$ has a $\mathcal{U}$-rank $< \alpha$. In the case $U = V(D)$ we call the U-rank of D the *rank of D* (provided that D has a U-rank). Note that more generally if $U \supseteq V(D)$ for a digraph, then its U-rank equals its rank.

Lemma 7.2.5 (Necklace Lemma 6.3.5)**.** *Let D be any digraph and $\mathcal{U}$ a finite set of vertex sets of D. Then exactly one of the statements is true:*

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

An *arborescence* is a rooted oriented tree T that contains for every vertex $v \in V(T)$ a directed path from the root to v. A *directed star* is an arborescence whose underlying tree is an undirected star that is centred in the root of the arborescence. A *directed comb* is the union of a ray with infinitely many finite disjoint paths (possibly trivial) that have precisely their first vertex on R. Hence the underlying graph of a directed comb is an undirected comb. The *teeth* of a directed comb or reverse directed comb are the teeth of the underlying comb. The ray from the definition of a comb is the *spine* of the comb. Given a set U of vertices in a digraph, a comb *attached* to U is a comb with all its teeth in U and a star *attached* to U is a subdivided infinite star with all its leaves in U. The set of teeth is the *attachment set* of the comb and the set of leaves is the *attachment set* of the star. We adapt the notions of 'attached to' and 'attachment sets' to reverse combs or reverse stars, respectively. We need Lemma 6.3.2 from the first chapter of this part:

Lemma 7.2.6 (Directed Star-Comb Lemma)**.** *Let D be any strongly connected digraph and let $U \subseteq V(D)$ be infinite. Then D contains a star or comb attached to U and a reverse star or reverse comb attached to U sharing their attachment sets.*

In the second part of this section, we list the tools and terminology about inverse limits that we need. Here we follow the textbook of Zalesskii and Ribes [63].

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

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

In this situation, we say that the map Φ is *induced* by the maps φ_i. For a topological space Y together compatible maps $\varphi_i \colon Y \to X_i$, for $i \in I$, the collection of maps $\{ \varphi_i \mid i \in I \}$ is called *eventually injective* if for every two distinct $y, y' \in Y$ there is some $i \in I$ with $\varphi_i(y) \neq \varphi_i(y')$, see [25, Lemma 8.8.4].

Lemma 7.2.7 (Lifting Lemma). *Let $\{X_i, f_{ij}, I\}$ be any inverse system and let Y be a topological space together with eventually injective compatible maps $\varphi_i\colon Y \to X_i$, for $i \in I$. Then the unique continuous map $\Phi\colon Y \to \varprojlim(X_i)_{i\in I}$ given by the universal property of the inverse limit is injective.*

We need the following two results [63, Lemma 1.1.7] and [63, Lemma 1.1.9]:

Lemma 7.2.8. *Let $\{X_i, f_{ij}, I\}$ be an inverse system of topological spaces over a directed set I, and let $\varphi_i\colon X \to X_i$ be compatible surjections from the space X onto the spaces X_i $(i \in I)$. Then either $\varprojlim(X_i)_{i\in I} = \emptyset$ or the induced mapping $\Phi\colon X \to \varprojlim(X_i)_{i\in I}$ maps X onto a dense subset of $\varprojlim(X_i)_{i\in I}$.*

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

Lemma 7.2.9. *Let $\{X_i, f_{ij}, I\}$ be an inverse system of compact topological spaces over a directed poset I and assume that I' is a cofinal subset of I. Then*

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

Finally, we need the following theorem [33, Theorem 3.1.13] from basic topology:

Lemma 7.2.10. *Every continuous injective mapping of a compact space onto a Hausdorff space is a homeomorphism.*

7.3. A topology for digraphs

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

In this topological space, topological arcs and circles take the role of paths and cycles, respectively. This makes it possible to extend to the space $|D|$ statements about finite digraphs. As an important cornerstone we characterise, in this section, those digraphs D for which $|D|$ is compact, see Theorem 7.1.

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

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

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

Neighbourhoods $\hat{C}_\varepsilon(X, \omega)$ of an end ω are of the following form: Given $X \in \mathcal{X}(D)$ let $\hat{C}_\varepsilon(X, \omega)$ be the union of (see Figure 7.3.1)

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

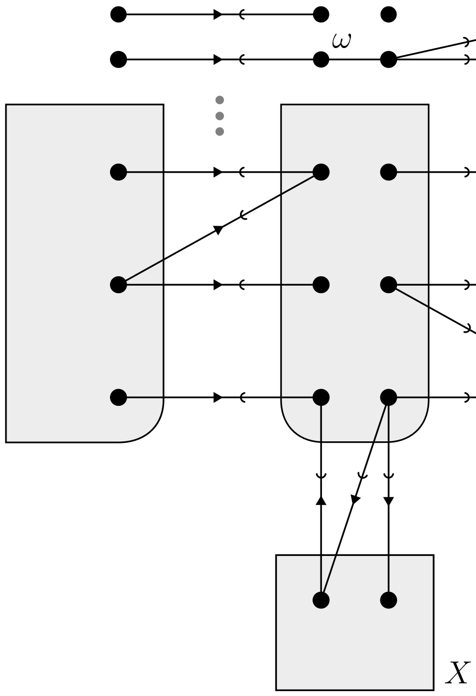

Figure 7.3.1.: A basic open neighbourhood of the form $\hat{C}_\varepsilon(X, \omega)$.

Neighbourhoods $\hat{E}_{\varepsilon,z}(X, \omega\eta)$ of an inner point z of a limit edge $\omega\eta$ between ends are of the following form: Given $X \in \mathcal{X}(D)$ that separates ω and η let $\hat{E}_{\varepsilon,z}(X, \omega\eta)$ be the union of (see Figure 7.3.2)

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

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

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

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

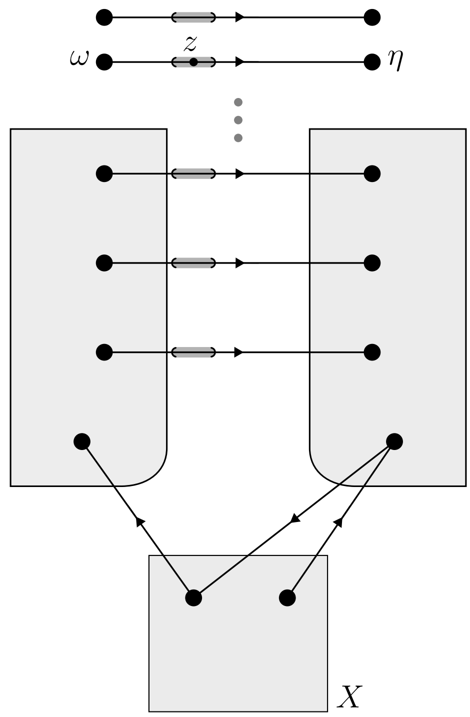

Figure 7.3.2.: A basic open neighbourhood of the form $\hat{E}_{\varepsilon,z}(X, \omega\eta)$.

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

One of the key definitions in the first chapter of this part, was that an end ω of D is said to be in the closure of $\mathcal{U}$, for a set of vertex sets $\mathcal{U}$, if for all $X \in \mathcal{X}(D)$ every $U \in \mathcal{U}$ has a vertex in $C(X, \omega)$. Now that DTop is at hand this is tantamount to $\omega \in \overline{U}$ for every $U \in \mathcal{U}$. We therefore obtain an extension of Lemma 6.4.1:

Lemma 7.3.1. *Let D be any digraph, and let $\mathcal{U}$ be a finite set of vertex sets of D. Then the following assertions are equivalent:*

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

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

Theorem 7.1. *The space $|D|$ is compact if and only if D is solid.*

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

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

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

So let us assume for a contradiction that there is a bad set for every $X \in \mathcal{X}(D)$. We will find a bad set for every $X \in \mathcal{X}(D)$ in a consistent way, i.e., for every two vertex sets $X, Y \in \mathcal{X}(D)$ the bad set of X contains that of Y whenever $X \subseteq Y$. In other words the bad sets will give rise to a direction f and we will then conclude that $f(X)$ is covered by finitely many sets in $\mathcal{O}$ for some $X \in \mathcal{X}(D)$, contradicting that $f(X)$ is bad.

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

In order to arrive at a contradiction we consider three cases. First, if f is a vertex-direction, then by Theorem 7.2.1, we have that f corresponds to an end ω which is covered by some cover set $O \in \mathcal{O}$. As O is a basic open set it is of the form $\hat{C}_\varepsilon(X, \omega)$ for some $X \in \mathcal{X}(D)$. This contradicts that $f(X)$ is bad.

Second, suppose that f is an edge-direction and that f corresponds to a limit edge $\omega\eta$ between ends in the sense of Theorem 7.2.2. This limit edge $\omega\eta$ is covered by a finite subset $\mathcal{O}' \subseteq \mathcal{O}$, as it is homeomorphic to the unit interval. Since each cover set $O \in \mathcal{O}'$ is basic open it comes by its definition together with a finite vertex set $X_O \in \mathcal{X}(D)$. Let $\mathcal{X}' := \{ X_O \mid O \in \mathcal{O}' \}$ and let X be large enough so that it contains $\bigcup \mathcal{X}'$ and so that it separates ω and η. To get a contradiction, we show that $\widehat{f(X)}$ is covered by $\mathcal{O}'$. Consider a point $z \in \widehat{f(X)}$ and let z' be its corresponding point on $\omega\eta$. Then z' is covered by some $O \in \mathcal{O}'$. Since $X_O \subseteq X$ we have $f(X) \subseteq f(X')$ and therefore O' also contains z.

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

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

Corollary 7.3.2. *The following statements are equivalent for any digraph D:*

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

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

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

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

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

(v)→(i) Theorem 7.1. □

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

In this section we show that the space $|D|$ for a solid digraph D can be obtained as an inverse limit of finite contraction minors of D, Theorem 7.2. We begin by defining an inverse system of finite digraphs for any digraph. Then, we show that every digraph embeds in the inverse limit of its inverse system. This gives a compactification for arbitrary digraphs, Theorem 7.4.1.

Let us introduce an inverse system for a given digraph D. For this, we define a directed partially ordered set $(\mathcal{P}, \leq)$ as follows. We call a finite partition P of $V(D)$ *admissible* if any two partition classes of P can be separated in D by a finite vertex set. We denote by $\mathcal{P} := \mathcal{P}(D)$ the set of all the admissible partitions of D. For any two partitions P_1 and P_2 of the vertex set of D we write $P_1 \leq P_2$ and say that P_2 is *finer* than P_1 if every partition class of P_2 is a subset of a partition class of P_1.

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

Let us proceed by defining the topological spaces associated with the admissible partitions of D. Every admissible partition P of D gives rise to a finite (multi-) digraph D/P by contracting each partition class and replacing all the edges between two partition classes by a single edge whenever there are infinitely many. Formally, declare P to be the vertex set of D/P. Given distinct partition classes $p_1, p_2 \in P$ we define an edge (e, p_1, p_2) of D/P for every edge e in D from p_1 to p_2 if there are finitely many such edges. And if there are infinitely many edges from p_1 to p_2 we just define a single edge (p_1p_2, p_1, p_2). We call the latter type of edges *quotient edges*. Endowing D/P with the 1-complex topology turns it into a compact Hausdorff space, i.e., basic open sets are uniform ε stars around vertices and open subintervals of edges. In other words, D/P is defined as our topological space from the previous section (for finite D) with the only difference that multi-edges are taken into account. We will usually not distinguish between the finite (multi-) digraphs D/P and the topological space $|D/P|$. Now, let us turn to the final ingredient of our inverse system for D: bonding maps. We define for every two distinct admissible partitions $P \leq P'$ of D a bonding map $f_{P'P} \colon D/P' \to D/P$ as follows. Vertices $p' \in P'$ of D/P' get mapped to the unique vertex $p \in P$ of D/P with $p' \subseteq p$. Edges get mapped according to their endvertices: For edges (e', p_1', p_2') of D/P' we consider two cases: First, if $p_1', p_2' \subseteq p$ for a partition class $p \in P$, then (e', p_1', p_2') gets mapped to the vertex p of D/P. Second, if $p_1' \subseteq p_1$ and $p_2' \subseteq p_2$ for two distinct partition classes $p_1, p_2 \in P$, then there is at least one edge from p_1 to p_2 in D/P. If (e', p_1', p_2') is a quotient edge in D/P', then also (p_1p_2, p_1, p_2) is a quotient edge in D/P and we map (e', p_1', p_2') to (p_1p_2, p_1, p_2). If (e', p_1', p_2') is not a quotient edge in D/P' and (p_1p_2, p_1, p_2) is a quotient edge in D/P, we map (e', p_1', p_2') to (p_1p_2, p_1, p_2). Finally, if (e', p_1', p_2') is not a quotient edge and there is no quotient edge between p_1 and p_2, then (e', p_1, p_2) is an edge in D/P and we map (e', p_1', p_2') to (e', p_1, p_2).

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

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

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

Theorem 7.4.1. *For every digraph D the space $\varprojlim(D/P)_{P\in\mathcal{P}}$ is a Hausdorff compactification of D, in particular, the map*

$$\Phi \colon D \to \varprojlim(D/P)_{P\in\mathcal{P}}$$

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

Proof. We have to show that $\varprojlim(D/P)_{P\in\mathcal{P}}$ is compact and Hausdorff, that the image of Φ is dense in $\varprojlim(D/P)_{P\in\mathcal{P}}$ and that Φ is an embedding i.e., it is a homeomorphism onto its image. The inverse limit $\varprojlim(D/P)_{P\in\mathcal{P}}$ is compact and Hausdorff because all the topological spaces D/P are compact and Hausdorff. As every φ_P is surjective the image of Φ is dense in $\varprojlim(D/P)_{P\in\mathcal{P}}$, by Lemma 7.2.8. In order to show that Φ is a homeomorphism onto its image, note first that the collection of maps $\{\, \varphi_P \mid P \in \mathcal{P} \,\}$ is eventually injective. Hence Φ is injective by the lifting lemma.

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

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

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

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

Now, consider an inner point z of a limit edge λ. Note that we have already defined the images of the two endpoints of λ. If these images coincide, then map z to the unique image of the endpoints of λ. Otherwise, Proposition 7.2.3 or Proposition 7.2.4 gives rise to a quotient edge λ' between the partition classes of the endpoints of λ. In this case we map z to the corresponding point on λ'. This completes the definition of $\hat{\varphi}_P$.

Lemma 7.4.2. *The map $\hat{\varphi}_P\colon |D| \to D/P$ is continuous for every $P \in \mathcal{P}$.*

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

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

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

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

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

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

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

Theorem 7.2. *Let D be a solid digraph. The map induced by the $\hat{\varphi}_P \colon |D| \to D/P$*

$$\hat{\Phi} \colon |D| \to \varprojlim (D/P)_{P \in \mathcal{P}}$$

is a homeomorphism.

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

By the lifting lemma and Lemma 7.4.2 the $\hat{\varphi}_P$ induce a continuous injective map $\hat{\Phi} \colon |D| \to \varprojlim (D/P)_{P \in \mathcal{P}}$. By Lemma 7.2.8 we have that the image of the map $\hat{\Phi}$ is dense in $\varprojlim (D/P)_{P \in \mathcal{P}}$. Moreover, as D is solid, we have that $|D|$ is compact by Theorem 7.1 so the image of $\hat{\Phi}$ is closed; hence it is all of $\varprojlim (D/P)_{P \in \mathcal{P}}$. The statement now follows from Lemma 7.2.10. □

In the proof of Theorem 7.2 we used those admissible partitions that arise by deleting a finite vertex set from a solid digraph to ensure that the map Φ that is induced by the $\hat{\varphi}_P$ is injective. Next, we show that these admissible partitions capture the whole inverse system for a solid digraph.

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

Now, $\{D/P_X, f_{P_X P_{X'}}, \mathcal{P}_X\}$ is an inverse system by itself and by Lemma 7.2.9 we have that

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

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

Corollary 7.4.3. *Let D be a countable solid digraph and let X_n consist of the first n vertices of D with regard to a any fixed enumeration of $V(D)$. Then $|D| \cong \varprojlim(D/P_{X_n})_{n\in\mathbb{N}}$.* □

7.5. Applications

In this last section we prove two statements about finite digraphs that naturally generalise to the space $|D|$, but do not generalise verbatim to infinite digraphs, Theorem 7.3 and Theorem 7.4. We begin this section by introducing all the definitions needed. We then provide an important tool that describes how (topological) paths in $|D|$ can pass through cuts in D, the directed jumping arc lemma. Finally, we prove our two main results of this section, Theorem 7.3 and Theorem 7.4.

A continuous function $\alpha\colon [0,1] \to |D|$ is called a *local homeomorphism on the edges of* $|D|$ if for every $x \in (0,1)$ that is mapped to an inner point of an edge $e \in E \cup \Lambda$, there is a neighbourhood (a,b) of x such that α restricts on (a,b) to a homeomorphism to the interior of e, i.e., $\alpha \restriction (a,b) \cong \mathring{e}$. Note that by the continuity of α any such homeomorphism $\alpha \restriction (a,b)$ extends to a homeomorphism $\alpha \restriction [a,b] \cong e$. If in addition α *respects the orientation of the edges* in $|D|$, that is if $[a,b] \subseteq [0,1]$ is mapped to an edge $e \in E \cup \Lambda$ we have that $x \leq y$ implies $\alpha(x) \leq \alpha(y)$ for all the $x, y \in [a,b]$, then we call α a *directed topological path* in $|D|$. (Here $\alpha(x) \leq \alpha(y)$ refers to $\leq$ in $[0,1]_e$.)

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

Next, let us gain some understanding of how directed topological paths in $|D|$ can pass through cuts of D, see also the jumping arc lemma [25, Lemma 8.6.3].

Lemma 7.5.1 (Directed Jumping Arc Lemma)**.** *Let D be any digraph and let $\{V_1, V_2\}$ be any bipartition of $V(D)$.*

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

Proof. (i) Suppose that $\overline{V_1} \cap \overline{V_2}$ is empty. Then every end of D is either contained in $\overline{V_1}$ or $\overline{V_2}$. First, we show that every edge of $|D|$ that has both of its endpoints in $\overline{D[V_i]}$ is contained in $\overline{D[V_i]}$, for $i = 1, 2$. For edges of D this is trivial. So consider a limit edge λ with both endpoints in $\overline{D[V_i]}$. All but finitely many vertices of a subdigraph obtained by Proposition 7.2.3 or Proposition 7.2.4 applied to λ are contained in $D[V_i]$, otherwise this gives an end in $\overline{V_1} \cap \overline{V_2}$. Consequently, for every inner point $z \in \lambda$ there is a sequence of inner points of edges in $D[V_i]$ that converge to z, giving $z \in \overline{D[V_i]}$.

From this first observation, we now know that $|D| \setminus (\overline{D[V_1]} \cup \overline{D[V_2]})$ consists only of inner points of edges (possibly limit edges) between $\overline{V_1}$ and $\overline{V_2}$. Now, consider a directed topological path α that connects a point in V_1 to a point in V_2. As $[0,1]$ is connected and α is continuous there is a point $x \in [0,1]$ with $\alpha(x) \in |D| \setminus (\overline{D[V_1]} \cup \overline{D[V_2]})$. Hence the preimage of $|D| \setminus (\overline{D[V_1]} \cup \overline{D[V_2]})$ is non-empty and a union of pairwise disjoint intervals (a, b) each of which is mapped homeomorphically to an open edge between $\overline{V_1}$ and $\overline{V_2}$. The usual relation $\leq$ on the reals defines a linear order on these intervals. Among these intervals, choose (a, b) minimal. That this is possible can be seen as follows: If not, we find a strictly decreasing sequence $(a_0, b_0) \geq (a_1, b_1) \geq \ldots$ of intervals with $\alpha \restriction [a_i, b_i] \cong e_i$ for some edges e_i between $\overline{V_1}$ and $\overline{V_2}$. Then $(a_i)_{i\in\mathbb{N}}$ and $(b_i)_{i\in\mathbb{N}}$ converge to some $c \in [0,1]$ and using that α is continuous, we get $\alpha(c) \in \overline{V_1} \cap \overline{V_2}$, a contradiction.

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

(ii) First note that no inner point of a limit edge is a limit point of a set of vertices. Hence D has at least one end that is contained in both the closure of V_1 and V_2. By Lemma 7.3.1 we find a necklace $N \subseteq D$ attached to $\{V_1, V_2\}$. Let ω be the end that is represented by N. Apply Corollary 7.3.2 to the solid digraph $D[V_1]$ and the infinite set $U_1 := V_1 \cap V(N)$ in order to obtain a necklace N_1 attached to U_1. Let U_2 consist of all the vertices in V_2 that are contained in those beads of N that intersect a bead of N_1. Apply Corollary 7.3.2 to the solid digraph $D[V_2]$ and the infinite set U_2 in order to obtain a necklace N_2 attached to U_2. Note that both necklaces N_1 and N_2 represent ω. A ray in N_1 together with a reverse ray in N_2 defines a directed topological path that is as desired. □

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

A closed directed topological path α that traverses every edge of $|D|$ exactly once is called *Euler tour*, i.e., for every edge e of $|D|$ there is exactly one subinterval of $[0, 1]$ that is mapped homeomorphiccally to e via α. If $|D|$ has an Euler tour we call $|D|$ *Eulerian*. There are two obstructions for digraph D to be Eulerian: one is a vertex of infinite degree and the other one is an unbalanced cut. A digraph is *locally finite* if all its vertices have finite in- and out-degree. Theorem 7.3 states that there are no further obstructions. We need one more lemma for its proof:

Lemma 7.5.2. *Let D be a digraph with a connected underlying graph. If D is locally finite and every finite cut of D is balanced, then D is solid.*

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

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

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

Theorem 7.3. *For a digraph D with a connected underlying graph the following assertions are equivalent:*

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

Proof. For the forward implication (i)→(ii) suppose that D has an Euler tour α. Using the directed jumping arc lemma it is straightforward to show that D has only balanced cuts. Let us show that D needs to be locally finite for α to be continuous. Suppose for a contradiction there is a $v \in V(D)$ with infinitely many edges $e_0, e_1, \ldots$ with head v; the case where v is the tail of infinitely many edges is analogue. Let $(a_i, b_i) \subseteq [0, 1]$ the

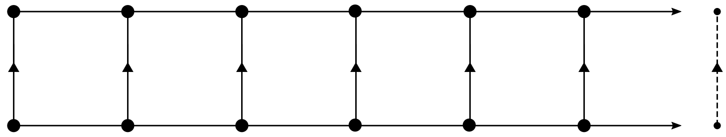

Figure 7.5.1.: A solid digraph (every undirected edge in the figure stands for two directed edges in opposite directions) that is not strongly connected. Adding the one end of the underling undirected graph makes it possible to find a closed directed topological path that contains all the vertices.

subinterval that is mapped homeomorphic by α to e_i. As the unit interval is compact, the sequence of the a_i has a convergent subsequence $(a_{i_n})_{n\in\mathbb{N}}$ and we write x for the limit point of this subsequence. Now, the subsequence $(b_{i_n})_{n\in\mathbb{N}}$ of the b_i forms a convergent subsequence, too, with limit point x. As $\alpha(b_{i_n}) = v$ for all the $n \in \mathbb{N}$ we have $\alpha(x) = v$, by the continuity of α; but $\alpha(a_{i_n})$ is a sequence of neighbours of v which does not converge in $|D|$ to v, a contradiction.

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

Let us now find an Euler tour for $|D|$. By Lemma 7.5.2 the digraph D is solid. As it is locally finite and its underlying graph is connected $V(D)$ is countable. Choose an enumeration of $V(D)$ and let X_n denote the set of the first n vertices. Then D/P_{X_n} contains no quotient edge and every cut of D/P_{X_n} is balanced. As the statement of Theorem 7.3 holds for finite digraphs, we have that D/P_{X_n} is Eulerian. Moreover, as D/P_{X_n} is a finite digraph there are only finitely many (combinatorial) Euler tours of D/P_{X_n}. By Kőnig's infinity lemma there is a consistent choice of one Euler tour for every D/P_{X_n}. Now, take a parametrisation $\alpha_n\colon [0,1] \to D/P_{X_n}$ of the Euler tour chosen for D/P_{X_n} such that the α_n are compatible. Using Theorem 7.2 it is straightforward to check that the universal property of the inverse limit gives an Euler tour for $|D|$. □

It is well know that a finite digraph is strongly connected if and only if it has a directed closed walk that contains all its vertices. Clearly, the statement does not generalise verbatim to infinite digraphs nor does a spanning directed (double)–ray ensure the digraph to be strongly connected. Moreover, the statement does not hold if one adds the ends of the underling undirected graph, see Figure 7.5.1.

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

Theorem 7.4. *For a countable solid digraph D the following assertions are equivalent:*

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

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

Once the W_n are defined, it is not hard to find parametrisations α_n of each W_n such that $f_{X_n X_{n-1}} \circ \alpha_n = \alpha_{n-1}$. Then the universal property of the inverse limit together with Corollary 7.4.3 and Theorem 7.2 gives the desired closed directed topological path in $|D|$.

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

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

8. Ends of digraphs III: normal arborescences

8.1. Introduction

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

For connected finite graphs, the depth-first search trees are precisely the normal spanning trees. Here, a rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T-path in G are comparable in the tree-order of T. (A *T-path* in G is a non-trivial path that meets T exactly in its endvertices.) Normal spanning trees generalise depth-first search trees, since they are also defined for infinite graphs; they are perhaps the single most important structural tool in infinite graph theory [25].

In this third chapter of this part we introduce and study normal spanning arborescences. These are generalisations of depth-first search trees to infinite digraphs that promise to be as powerful for a structural analysis of digraphs as normal spanning trees are for graphs, both from a combinatorial and a topological point of view. For example, they play a crucial role in the next chapter, see also [57], to extend theorems about finite tournaments to the infinite.

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

Consider a finite digraph D together with a spanning depth-first search tree $T \subseteq D$. If vw is an edge of D between $\leq_T$-incomparable vertices of T, then w is visited at an earlier stage of the depth-first search than v.[1] Together with all such edges, T forms an acyclic subdigraph of D [19].[2]

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

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

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

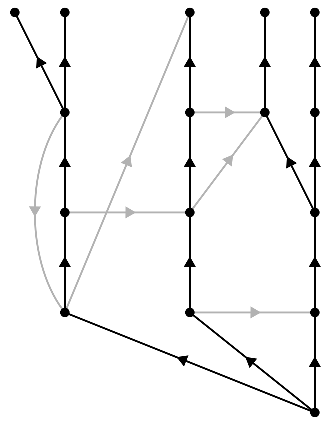

Figure 8.1.1.: A depth-first search arborescence visiting vertices from right to left.

only if T defines a depth-first search tree; see Corollary 8.3.3.

One aspect of why normal spanning trees of infinite undirected graphs are so useful is that they are end-faithful. A spanning tree T of a graph G is *end-faithful* if the map that assigns to every end of T the end of G that contains it as a subset (of rays) is bijective, see [25]. Equivalently T is end-faithful if every end of G is represented by a unique ray in T that starts from a fixed root. Our first main result in this chapter will be an analogue of this for normal arborescences, so let us recall the definition of ends of digraphs from the first chapter of this part that we need.

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

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

Theorem 8.1. *Every normal spanning arborescence of a digraph is end-faithful.*

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

The end space of any normal spanning tree T of an undirected graph G coincides with the end space of G, not only combinatorially but also topologically. Indeed, the map that assigns to every end of T the end of G that contains it as a subset is a homeomorphism between the end space of T and that of G, see [25]. Hence, in order to understand the end space of G one just needs to understand the simple structure of the tree T.

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

For two distinct ends ω and η of D, we call the pair (ω, η) a *limit edge* from ω to η if D has an edge from $C(X, \omega)$ to $C(X, \eta)$ for every finite vertex set X for which ω and η live in distinct strong components of $D - X$. Similarly, for a vertex $v \in V(D)$ and an end ω of D we call the pair (v, ω) a *limit edge from v to ω* if D has an edge from v to $C(X, \omega)$ for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. And we call the pair (ω, v) a *limit edge from ω to v* if D has an edge from $C(X, \omega)$ to v for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. The digraph D, its ends, and its limit edges together form a topological space $|D|$, in which the edges are copies of the real interval $[0, 1]$; see 7.

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

Recall that the digraphs D that are compactified by $|D|$ are precisely the *solid* ones, those such that $D - X$ has only finitely many strong components for every finite vertex set $X \subseteq V(D)$, see [14] or the previous Chapter 7. Let T be a normal spanning arborescence of D, with root r, say. By Theorem 8.1, there exists a well-defined map ψ that sends every end ω of D to the end of the solidification $\overline{T}$ of T represented by the unique ray $R \subseteq T$ starting from r that represents ω in D. This map ψ is clearly injective. If D is solid, every ray in T represents an end of D, so ψ is also surjective. Let ζ denote the map from the set of ends of $\overline{T}$ to that of the solidification $\overline{H}$ of the normal assistant H of T in D that assigns to every end of $\overline{T}$ the end of $\overline{H}$ that contains it as a subset (of rays). This is always bijective, see Lemma 8.5.1. Note that $\overline{H}$, unlike $\overline{T}$, can have limit edges. We say that T *reflects the horizon* of D if the map $\zeta \circ \psi \colon \Omega(D) \to \Omega(\overline{H})$ extends to a homeomorphism from the horizon of D to that of $\overline{H}$.

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

Theorem 8.2. *Every normal spanning arborescence of a solid digraph reflects its horizon.*

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

Jung [50] characterised the connected graphs with a normal spanning tree in terms of dispersed sets. A set $U \subseteq V(G)$ of vertices of a graph G is *dispersed* if there is no comb in G with all its teeth in U. Recall that a *comb* is the union of a ray R with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R. The last vertices of those paths are the *teeth* of this comb, see [25]. Jung proved that a connected graph has a normal spanning tree if and only if its vertex set is a countable union of dispersed sets.

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

Theorem 8.3. *Let D be any digraph and suppose that $r \in V(D)$ can reach all the vertices of D. If $V(D)$ is a countable union of dispersed sets, then D has a normal spanning arborescence rooted in r.*

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

In an undirected graph, the levels of any normal spanning tree are dispersed, so the forward implication in Jung's characterisation is easy. Theorem 8.3 implies the harder backward implication when applied to the digraph obtained from the graph by replacing every edge by its two orientations as separate directed edges.

The easy forward implication in Jung's theorem does not have a directed analogue, since the converse implication in Theorem 8.3 may fail (see Section 8.6). However, the converse of Theorem 8.3 does hold if the digraph D is solid.

This chapter is organised as follows. We provide the tools and terminology that we use throughout this chapter in Section 8.2. Then in Section 8.3 we introduce normal arborescences and provide some basic lemmas that we need for the proofs of our main results. In Section 8.4, we show that normal spanning arborescences are end-faithful, Theorem 8.1. In Section 8.5, we prove that normal spanning arborescences reflect the horizon, Theorem 8.2. Finally, we prove our existence criterion for normal arborescences in digraphs, Theorem 8.3, in Section 8.6.

8.2. Tools and terminology

Any graph-theoretic notation not explained here can be found in Diestel's textbook [25]. For the sake of readability, we sometimes omit curly brackets of singletons, i.e., we write x instead of $\{x\}$ for a set x. Furthermore, we omit the word 'directed'—for example in 'directed path'—if there is no danger of confusion.

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

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

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

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

A vertex set $Y \subseteq V(D)$ *separates* A and B in D with $A, B \subseteq V(D)$ if every A–B path meets Y, or if every B–A path meets Y. For two vertices v and w of D we say that $Y \subseteq V(D) \setminus \{v, w\}$ *separates* v and w in D, if it separates $\{v\}$ and $\{w\}$ in D.

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

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

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

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

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

An inflated H is any digraph that arises from a subdivision H' of H as follows. Replace every branch vertex v of H' by a strongly connected digraph H_v so that the H_v are disjoint and do not meet any subdividing vertex; here replacing means that we first delete v from H' and then add $V(H_v)$ to the vertex set and $E(H_v)$ to the edge set. Then replace every subdividing path P_{vw} that starts in v and ends in w by an H_v–H_w path that coincides with P_{vw} on inner vertices. We call the vertex sets $V(H_v)$ the *branch sets* of the inflated H. A *necklace* is an inflated symmetric ray with finite branch sets; the branch sets of a necklace are its *beads*. With a slight abuse of notation, we say that a necklace $N \subseteq D$ *represents* an end ω of D if one (equivalently every) ray in N represents ω. Given a set U of vertices in a digraph D, a necklace $N \subseteq D$ is *attached to* U if infinitely many of the branch sets of N contain a vertex from U.

For two distinct ends $\omega, \eta \in \Omega(D)$, we call the pair (ω, η) a *limit edge* from ω to η, if D has an edge from $C(X, \omega)$ to $C(X, \eta)$ for every finite vertex set $X \subseteq V(D)$ that separates ω and η. For a vertex $v \in V(D)$ and an end $\omega \in \Omega(D)$ we call the pair (v, ω) a *limit edge from v to ω* if D has an edge from v to $C(X, \omega)$ for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. Similarly, we call the pair (ω, v) a *limit edge from ω to v* if D has an edge from $C(X, \omega)$ to v for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. We write $\Lambda(D)$ for the set of all the limit edges of D. As we do for 'ordinary' edges of a digraph, we will suppress the brackets and the comma in our notation of limit edges. For example we write $\omega\eta$ instead of (ω, η) for a limit edge between ends ω and η. For limit edges we need Proposition 7.2.4 from the first chapter of this part:

Proposition 8.2.1. *For a digraph D, a vertex v and an end ω of D the following assertions are equivalent:*

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

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

8.3. Normal arborescences

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

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

Lemma 8.3.1. *Let D be any digraph and let $T \subseteq D$ be an arborescence. If the normal assistant of T in D contains a cycle, then it also contains a cycle so that consecutive vertices on the cycle are $\leq_T$-incomparable.*

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

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

Lemma 8.3.2. *Let D be any digraph and let $T \subseteq D$ be an arborescence in D. Then T is normal in D if and only if there is a sensitive order on T.*

Proof. For the forward implication assume that T is normal in D. Let us write L_n for the nth level of T and let us write T_n for the arborescence that T induces on $\bigcup\{\, L_m \mid m \le n \,\}$. We recursively construct an ascending sequence of orders $(\preccurlyeq_n)_{n\in\mathbb{N}}$ such that $\preccurlyeq_n$ is a sensitive order on T_n as follows. In the base case, we let $\preccurlyeq_0 := \le_{T_0}$. In the recursive step, suppose that we have defined $\preccurlyeq_n$. Let us write for every $v \in L_n$ the set of up-neighbours (children) of v in T as N_v. For every $v \in L_n$ let $\preccurlyeq_v$ be a linear extension of the restriction of $\trianglelefteq_T$ to N_v. And for every two distinct vertices $v, w \in L_n$ with $v \preccurlyeq_n w$ we define $v' \preccurlyeq_{vw} w'$ whenever $v' \in N_v$ and $w' \in \lceil N_w \rceil_T \smallsetminus \lceil v \rceil_T$. Now, let $\preccurlyeq_{n+1}$ be the transitive closure of

$$\preccurlyeq_n \cup \bigcup\{\, \preccurlyeq_v \mid v \in L_n \,\} \cup \bigcup\{\, \preccurlyeq_{vw} \mid v \neq w \text{ in } L_n \,\}.$$

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

For the backward implication assume that T has a sensitive order $\preccurlyeq$ on T. Suppose for a contradiction that T is not normal in D. Let H be the normal assistant of T in D. Then H contains a cycle C and by Lemma 8.3.1 we may assume that consecutive vertices on C are $\le_T$-incomparable. Let c be the $\preccurlyeq$-largest vertex on C and let c' be its successor on C. Note that $c' \preccurlyeq c$ by the choice of c. The edge cc' of $C \subseteq H$ is witnessed by a T-path P from $\lfloor c \rfloor_T$ to $\lfloor c' \rfloor_T$. Let w be the first vertex and v the last vertex of P. As $\preccurlyeq$ is branch sensitive, we have $v \preccurlyeq w$. But then the two vertices v and w together with P show that $\preccurlyeq$ is not path sensitive contradicting that $\preccurlyeq$ is a sensitive order on T. □

Corollary 8.3.3. *A spanning arborescence of a finite digraph is normal if and only if it defines a depth-first search tree.*

Proof. Let T be a spanning arborescence of a finite digraph D. For the forward implication assume that T is normal in D. By Lemma 8.3.2, we find a sensitive order $\preccurlyeq$ on T. Then T is defined by the traversed edges of the depth-first search that starts in the root of T and always chooses the $\preccurlyeq$-largest up-neighbour (child) in T in each step.

For the backward implication assume that T is a depth-first search tree and suppose for a contradiction that T is not normal in D. Then the normal assistant of T contains a cycle C and by Lemma 8.3.1 we may choose C so that consecutive vertices on C are $\leq_T$-incomparable. Let x be the vertex on C that is visited first in the depth-first search and let y be its successor on C. The edge xy of the normal assistant of T is witnessed by a T-path from $\lfloor x \rfloor_T$ to $\lfloor y \rfloor_T$. As T is spanning this path is just an edge e. Note that all vertices in $\lfloor x \rfloor_T$ are visited earlier than those in $\lfloor y \rfloor_T$ in the depth-first search. Hence the edge e should have been visited by the depth-first search; this is a contraction because e is not an edge of T. □

We think of (countable) normal spanning arborescence $T \subseteq D$ as being drawn in the plane with all the edges between $\leq_T$-incomparable vertices running from left to right; see Figure 8.1.1.

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

Lemma 8.3.4. *Let D be any digraph and let $T \subseteq D$ be a normal arborescence in D. If $v, w \in V(T)$ are $\leq_T$-incomparable vertices of T with $w \not\trianglelefteq_T v$, then every w–v path in D meets $X := \lceil v \rceil_T \cap \lceil w \rceil_T$. In particular, X separates v and w in D.*

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

The *dichromatic number* [58] of a digraph D is the smallest cardinal κ so that D admits a vertex partition into κ many partition classes that are acyclic in D. From the path sensitivity of normal arborescences we obtain the following:

Proposition 8.3.5. *Every digraph that has a normal spanning arborescence does have a countable dichromatic number.*

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

8.4. Arborescences are end-faithful

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

Theorem 8.1. *Let D be any digraph and let $U \subseteq V(D)$ be any vertex set. If T is a normal arborescence containing U, then T is end-faithful for the set of ends in the closure of U.*

We will employ the following star-comb lemma [25, Lemma 8.2.2] in order to prove Theorem 8.1:

Lemma 8.4.1 (Star-comb lemma)**.** *Let W be an infinite set of vertices in a connected undirected graph G. Then G contains a comb with all its teeth in W or a subdivided infinite star with all its leaves in W.*

Proof of Theorem 8.1. First, let R_1 and R_2 be distinct normal rays of T that represent ends of D in the closure of U, say ω_1 and ω_2, respectively. Our goal is to show that ω_1 and ω_2 are distinct ends of D. By Lemma 8.3.4, the rays R_1 and R_2 have tails in distinct strong components of $D - X$ for $X = V(R_1) \cap V(R_2)$. Hence X witnesses that R_1 and R_2 are not equivalent; in particular $\omega_1 \neq \omega_2$.

It remains two show that every end ω in the closure of U is represented by a normal ray of T. We claim that there is a necklace N attached to U in D that represents ω. For this consider the auxiliary digraph D' obtained from D by adding a new vertex v^* and adding new edges v^*u, one for every $u \in U$. Since ω is contained in the closure of U, we have that $v^*\omega$ is a limit edge of D'. Note, that adding v^* does not change the set of ends, in the sense that every end of D' contains a unique end of D as a subset (of rays), and we may identify the ends of D' with the ends of D. Now, Proposition 8.2.1 yields a necklace $N \subseteq D$ that represents ω such that v^* sends an edge to every bead of N. By the definition of D', we conclude that N is attached to U.

Having N at hand, fix a vertex from U of every bead of N and let W be the set of these fixed vertices. Now, apply the star-comb lemma in the undirected tree underlying T to W. We claim that the return is a comb. Indeed, suppose for a contradiction that we get a star and let c be its centre. By Lemma 8.3.4, the finite set $\lceil c \rceil_T$ separates any two leaves of the star, which is impossible because they are all contained in the necklace $N \subseteq D$.

So the return of the star-comb lemma is indeed a comb and we may assume that its spine R, considered as a ray in T, is a normal ray. Our aim is to prove that R represents ω and we may equivalently show that ω is contained in the closure of $V(R)$. So given a finite vertex set $X \subseteq V(D)$, fix teeth u and u' of the comb that are contained in $C(X, \omega)$. These exist because the teeth of the comb are contained in W and the choice of W. By

Lemma 8.3.4, the strong component $C(X,\omega)$ contains a vertex of $\lceil u \rceil_T \cap \lceil u' \rceil_T$. As this intersection is included in R we have verified that $C(X,\omega)$ contains a vertex of R. This completes the proof that ω is contained in the closure of R and with it the proof of this theorem. □

Corollary 8.4.2. *Let D be any digraph and let $U \subseteq V(D)$ be any vertex set. If T is a reverse normal spanning arborescence containing U, then T is end-faithful for the set of ends in the closure of U.*

Proof. Applying Theorem 8.1 to the digraph $\overleftarrow{D}$ and the normal arborescence $\overleftarrow{T} \subseteq \overleftarrow{D}$ shows that $\overleftarrow{T}$ is end-faithful for the ends of $\overleftarrow{D}$ in the closure of U. Hence the statement is a consequence of the fact that the ends of D in the closure of U correspond bijectively to the ends of $\overleftarrow{D}$ in the closure of U, via the map that sends an end ω of D to the end of $\overleftarrow{D}$ that is represented by some (equivalently every) reverse ray of D that represents ω. □

8.5. Arborescences reflect the horizon

One of the most useful facts about normal spanning trees is that the end space of any normal spanning tree T coincides with the end space of the graph G it spans—even topologically, i.e., the map that assigns to every end of T the end of G that contains it as a subset is a homeomorphism, see [27]. Hence, in order to understand the end space of G one just needs to understand the simple structure of the tree T.

In [14] we defined a topological space $|D|$ formed by a digraph D together with its ends and limit edges. The *horizon* of a digraph D is the subspace of $|D|$ formed by the ends of D and all the limit edges between them. In order to understand the results of this section it is not necessary to know the topology on $|D|$, as the subspace topology on the horizon of D is particularly simple. Let us give a brief description of the subspace topology for the horizon of D.

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

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

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

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

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

Now suppose that we have fixed a root r of T and suppose that every ray of D is solid in D. By Theorem 8.1, there exists a well-defined map ψ that sends every end ω of D to the end of the solidification $\overline{T}$ of T represented by the unique ray $R \subseteq T$ starting from r that represents ω in D. This map ψ is clearly injective. Note that the map ψ is also surjective, by our assumption that every ray of D is solid in D. Let ζ denote the map from the set of ends of $\overline{T}$ to that of the solidification $\overline{H}$ of the normal assistant H of T in D that assigns to every end of $\overline{T}$ the end of $\overline{H}$ that contains it as a subset (of rays). This is always bijective:

Lemma 8.5.1. *Let D be a digraph, T a normal spanning arborescence of D and H the normal assistant of T in D. The map $\zeta\colon \Omega(\overline{T}) \to \Omega(\overline{H})$ that assigns to every end of $\overline{T}$ the end of $\overline{H}$ that contains it as a subset is bijective.*

Proof. To see that ζ is injective, let ω_1 and ω_2 be distinct ends of $\overline{T}$ and let R_i be the ray in T starting from the root of T that represents ω_i for $i = 1, 2$. By Lemma 8.3.4 the two rays R_1 and R_2 have a tail in distinct strong components of $\overline{H} - X$ for $X := \lceil R_1 \rceil_T \cap \lceil R_2 \rceil_T$; hence ζ maps the ends ω_1 and ω_2 to distinct ends of $\overline{H}$.

To see that ζ is onto, let ω be an end of $\overline{H}$ and let R be any solid ray in $\overline{H}$ that represents ω. Our goal is to find a solid ray R' in $\overline{T}$ that is equivalent to R in $\overline{H}$: then the end of $\overline{T}$ that is represented by R' is included in ω as a subset of rays. For this apply the star-comb lemma in the undirected tree underlying T to the vertex set of R. If the return is a comb, then the comb's spine defines the desired ray R'. Indeed, the paths between the comb's spine and its teeth, define (in $\overline{H}$) a family of disjoint directed paths from R' to R and from R to R', hence R' and R are equivalent in $\overline{H}$. It now suffices to show that the return of the star-comb lemma is always a comb; so suppose for a contradiction that it is a star with centre c say. Then, by Lemma 8.3.4, the down-closure of c in T separates infinitely many vertices of $V(R)$ in $\overline{H}$, contradicting that R has a tail in a strong component of $\overline{H} - X$ for every finite vertex set X. □

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

The horizon of a normal spanning arborescences might differ from the horizon of the digraph it spans. This is due to the nature of connectivity in digraphs: a digraph might have a normal spanning arborescence and many strong components at the same time. For example consider the digraph D depicted in Figure 8.5.1. On the one hand, every ray in D is solid in D. On the other hand, consider the unique normal spanning arborescence

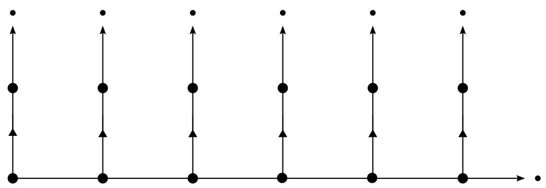

Figure 8.5.1.: A digraph D with a normal spanning arborescence T where the horizon of T differs from that of D. Every undirected edge in the figure represents a pair of inversely directed edges. Every line that ends with an arrow stands for a symmetric ray.

T that is rooted in the leftmost vertex of the bottom ray. Note that T coincides with its normal assistant. Hence T is normal in D and the end ω in the horizon of T is a limit point of the ends ω_i in the horizon of T. In contrast to that, all points in the horizon of D are isolated as every end lives in exactly one strong component of D.

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

Theorem 8.2. *Every normal spanning arborescence of a solid digraph reflects its horizon.*

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

Lemma 8.5.2. *Let D be a solid digraph, let T be a normal spanning arborescence of D and H the normal assistant of T in D. For ends ω of D and ω' of $\overline{H}$, with $\omega' = \zeta(\psi(\omega))$ the following statements hold:*

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

Proof. (i) Let X be any finite vertex set of D. We may assume that X is down-closed with respect to $\leq_T$. We write $R_{\omega'}$ for the unique ray of T that starts from the root of T and represents ω', Theorem 8.1. Note that since D is solid every ray in T is solid in D. We will find a vertex $v \in R_{\omega'}$ such that the up-closure of v in T is contained in $C(X, \omega)$; then $X' := \lceil v \rceil_T \setminus \{v\}$ is as desired by the separation properties of normal arborescences, Lemma 8.3.4.

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

(ii) Let X' be any finite vertex set of $\overline{H}$. By the separation properties of normal arborescences, Lemma 8.3.4, the finite vertex set $X := \lceil X' \rceil_T$ is as desired. □

Lemma 8.5.3. *Let D be a solid digraph, T a normal spanning arborescence of D and H the normal assistant of T in D. Then $\omega\eta$ is a limit edge of D if and only if $\omega'\eta'$ is a limit edge of $\overline{H}$, where ω' and η' is the image under the map $\zeta \circ \psi$ of ω and η, respectively.*

Proof. We write ω' and η' for the image under the map $\zeta \circ \psi$ of ω and η, respectively. Let us first show that $\omega'\eta'$ is a limit edge of $\overline{H}$ if $\omega\eta$ is a limit edge of D. For this let any finite vertex set X' that separates ω' and η' in $\overline{H}$ be given. Our goal is to find an edge in $\overline{H}$ from $C(X', \omega')$ to $C(X', \eta')$. By Theorem 8.1 there are rays R_ω and R_η in T that represent ω and η in D, respectively. As $\omega\eta$ is a limit edge of D, there is an edge in D from $\lfloor v_\omega \rfloor_T$ to $\lfloor v_\eta \rfloor_T$ for any two $\leq_T$-incomparable vertices $v_\omega \in R_\omega$ and $v_\eta \in R_\eta$. Now, choose such vertices v_ω and v_η so that both $\lfloor v_\omega \rfloor_T$ and $\lfloor v_\eta \rfloor_T$ avoid X'. In $\overline{H}$ both $\lfloor v_\omega \rfloor_T$ and $\lfloor v_\eta \rfloor_T$ are strongly connected, by the definition of the solidification. And as R_ω and R_η have a tail in $C(X', \omega')$ and $C(X', \eta')$, respectively, we have that $\lfloor v_\omega \rfloor_T \subseteq C(X', \omega')$ and $\lfloor v_\eta \rfloor_T \subseteq C(X', \eta')$. In particular, $\lfloor v_\omega \rfloor_T \cap \lfloor v_\eta \rfloor_T = \emptyset$. Consequently, any edge in D from $\lfloor v_\omega \rfloor_T$ to $\lfloor v_\eta \rfloor_T$ has $\leq_T$-incomparable endvertices and therefore is an edge in $\overline{H}$ from $C(X', \omega')$ to $C(X', \eta')$.

Now, let $\omega'\eta'$ be a limit edge of $\overline{H}$. We write ω and η for the unique preimage under $\zeta \circ \psi$ of ω' and η', respectively. We show that $\omega\eta$ is a limit edge in D. For this let any finite vertex set X that separates ω and η in D be given. Our goal is to find an edge in D from $C(X, \omega)$ to $C(X, \eta)$. As in the proof of Lemma 8.5.2, there are vertices v_ω and v_η such that $\lfloor v_\omega \rfloor_T \subseteq C(X, \omega)$ and $\lfloor v_\eta \rfloor_T \subseteq C(X, \eta)$. Let $X' = (\lceil v_\omega \rceil_T \cup \lceil v_\eta \rceil_T) \setminus \{v_\omega, v_\eta\}$ and consider $C(X', \omega')$ and $C(X', \eta')$ in $\overline{H}$. Using that T is normal in D it is easy to show that $C(X', \omega') = \lfloor v_\omega \rfloor_T$ and $C(X', \eta') = \lfloor v_\eta \rfloor_T$. As $\omega'\eta'$ is a limit edge of $\overline{H}$ there is an edge e in $\overline{H}$ from $C(X', \omega')$ to $C(X', \eta')$. Furthermore, the endpoints of e are $\leq_T$-incomparable. Now, e was added to T in the definition of H because there is an edge f of D from $\lfloor v_\omega \rfloor_T$ to $\lfloor v_\eta \rfloor_T$ and this edge f is as desired. □

Proof of Theorem 8.2. By Lemma 8.5.1 and its preceding text, the map $\zeta \circ \psi$ is a bijection. We extend this map to a bijection Θ between the horizon of D and that of $\overline{H}$ as follows. Let y be an inner point of a limit edge $\omega\eta$ between ends of D. We write ω' and η' for the image under $\zeta \circ \psi$ of ω and η, respectively. By Lemma 8.5.3 we have that $\omega'\eta'$ is a limit edge of $\overline{H}$. Then we declare $\Theta(y) := y'$ for y' the point that corresponds to y on $\omega'\eta'$. Again, by Lemma 8.5.3, the map Θ is bijective; we claim that Θ is even a homeomorphism. Indeed, using Lemma 8.5.2 (ii) it is straightforward to check that Θ is continuous and using Lemma 8.5.2 (i) it is straightforward to check that the inverse of Θ is continuous. □

8.6. Existence of arborescences

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

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

Theorem 8.3. *Let D be any digraph, $U \subseteq V(D)$ and suppose that $r \in V(D)$ can reach all the vertices in U. If U is a countable union of dispersed sets, then D has a normal arborescence that contains U and is rooted in r.*

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

However, the converse of Theorem 8.3 holds in an important case, namely if the digraph D is solid. Indeed, if D is solid then any arborescence $T \subseteq D$ that is normal in D is locally finite by the separation properties of normal arborescences, Lemma 8.3.4. Hence the levels of T are finite; in particular, dispersed.

An analogue of Theorem 8.3 holds for reverse normal arborescences:

Corollary 8.6.1. *Let D be any digraph and suppose that $U \subseteq V(D)$ is a countable union of dispersed sets in $\overleftarrow{D}$. If $r \in V(D)$ can be reached by all the vertices in U, then D has a reverse normal arborescence that contains U and is rooted in r.*

Proof. Apply Theorem 8.3 to the reverse of D. □

Proof of Theorem 8.3. Suppose that the vertex set U can be written as a countable union $\bigcup\{ U_n \mid n \in \mathbb{N} \}$ of sets that are dispersed in D. Then we can write U as a collection $\{ u_\alpha \mid \alpha < \kappa\}$ for a finite or limit ordinal κ such that every proper initial segment of the collection is dispersed in D as follows: We may assume that the U_n are pairwise disjoint. Choose a well-ordering $\leq_n$ of every U_n. Then write $u \leq u'$ for vertices $u \in U_m$ and $u' \in U_n$ with $m < n$, or with $m = n$ and $u \leq_m u'$. It is straightforward to show that $\leq$ defines a well-ordering of U that is as desired.

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

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

(i) T_α contains $\{ u_\xi \mid \xi \leq \alpha \}$ cofinally[3] with regard to $\leq_{T_\alpha}$;
(ii) if $v, w \in T_\alpha$ with $v \preccurlyeq_\alpha w$ are distinct and have a common $\leq_{T_\alpha}$-predecessor, then $w \in T_\xi$ and $v \notin T_\xi$ for some $\xi < \alpha$;
(iii) there is no infinite strictly ascending sequence of vertices in T_α with regard to $\preccurlyeq_\alpha$.

Once the T_α are defined the arborescence $T := \bigcup\{ T_\alpha \mid \alpha < \kappa \}$ is as desired; indeed, $\bigcup\{ \preccurlyeq_\alpha\mid \alpha < \kappa\}$ is a sensitive order on T and thus T is normal in D by Lemma 8.3.2. Finally, $V(T)$ contains U by condition (i).

Conditions (ii) and (iii) become relevant in the construction of the T_α, which now follows. If $\alpha = 0$, then let T_0 be any r–u_0 path in D and let $\preccurlyeq_0 := \leq_{T_0}$. Otherwise

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

$\beta > 0$. If β is a limit ordinal, then let $T_\beta := \bigcup\{ T_\alpha \mid \alpha < \beta \}$ and $\preccurlyeq_\beta = \bigcup\{ \preccurlyeq_\alpha \mid \alpha < \beta\}$. Then $\preccurlyeq_\beta$ is a sensitive order on T_β as each $\preccurlyeq_\alpha$ with $\alpha < \beta$ is a sensitive order on T_α. Condition (i) for β follows from (i) for $\alpha < \beta$ and our assumption that u_β coincides with u_α for some $\alpha < \beta$. Similarly, condition (ii) for β follows from (ii) for $\alpha < \beta$. Condition (iii) can be seen as follows. Suppose for a contraction that there is an infinite strictly ascending sequence $(w_n)_{n\in\mathbb{N}}$ in T_β with regard to $\preccurlyeq_\beta$. Apply the star-comb lemma to the set $\{ w_n \mid n \in \mathbb{N} \}$ in the undirected tree underlying T_β. The return is an infinite subdivided undirected star since an undirected comb would give rise to a directed comb in T_β with all its teeth in U; here we use that by (i) every tooth has a vertex of U in its up-closure and that every proper initial segment of $\{ u_\alpha \mid \alpha < \kappa\}$ is dispersed. Let Z be the set of $\le_{T_\beta}$-up-neighbours of the centre of the subdivided star that contain a tooth in their $\le_{T_\beta}$-up-closure. Since $\preccurlyeq_\beta$ is branch sensitive we may write Z as a strictly ascending collection $Z = \{ z_n \mid n \in \mathbb{N} \}$ with regard to $\preccurlyeq_\beta$. Choose $z^* \in Z \cap V(T_\alpha)$ so that α is minimal with $Z \cap V(T_\alpha) \neq \emptyset$. By (ii) we have that $z_n \preccurlyeq_\beta z^*$ for every $z_n \neq z^*$ contradicting that the z_n form a strictly ascending sequence with regard to $\preccurlyeq_\beta$.

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

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

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

9. Hamiltonicity in infinite tournaments

9.1. Introduction

A natural aim in infinite graph theory is to extend known theorems about finite graphs to infinite graphs. However the right way to do this is not always to apply the finite statement to infinite graphs verbatim: it often fails for trivial reasons, or becomes trivially true. A fruitful attempt to overcome this issue is to not only consider the graph itself, but the graph together with 'points at infinity': its ends.

Formally, an *end* of a graph is an equivalence class of its rays, where two rays are equivalent if no finite set of vertices separates them. A graph G together with its ends naturally forms a topological space $|G|$. For locally finite G, this is its well-known Freudenthal compactification. The topological properties of the space $|G|$ have been extensively studied [21, 28, 54, 60, 61, 67]. Letting topological arcs and circles in $|G|$ take the role of paths and cycles in G, it often becomes possible to extend theorems about paths and cycles in finite graphs to infinite graphs. Examples include Euler's theorem [2, 36], arboricity and tree-packing [22, 68], Hamiltonicity [11, 20, 25, 34, 37, 38, 46], and various planarity criteria [8, 10, 31].

In the first three chapters of this part we introduced a notion of ends in digraphs for which the fundamental techniques of undirected end space theory naturally generalise to digraphs. Unlike for undirected graphs, some ends of digraphs are joined by *limit edges*. A digraph D together with its ends and limit edges naturally forms a topological space $|D|$. So the scene is set now to attempt, also for digraphs D, to extend finite to infinite theorems by letting the naturally oriented topological paths and circles in $|D|$ take the role of directed paths and cycles in D. The purpose of this chapter is to make a start on this programme, with two well-known Hamiltonicity theorems for digraphs.

Two folklore theorems in finite graph theory, due to Rédei [62] and Camion [16], respectively, say that every finite tournament has a Hamilton path, and every finite strongly connected tournament has a Hamilton cycle. In this chapter we show that these results have natural analogues in the space $|D|$. We shall see that ends and limit edges are both crucial for such extensions to exist: there exists a countable tournament D whose compactification by just the ends of the underlying undirected graph contains no topological Hamilton path. (Similarly, D has no topological Hamilton path in $|D|$ that avoids all its limit edges, and D has no spanning ray or double ray.)

To state our results formally, we need a few definitions. A *ray* is an infinite directed path that has a first vertex (but no last vertex). The subrays of a ray are its *tails*. A ray in a digraph D is *solid* in D if it has a tail in some strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$. Two solid rays in D are *equivalent* if for every finite vertex set $X \subseteq V(D)$ they have a tail in the same strong component of $D - X$. The classes of this equivalence relation are the *ends* of D. For an end ω we write $C(X, \omega)$ for the strong component of $D - X$ in which every ray that represents ω has a tail. For two ends ω and η of D a finite vertex set $X \subseteq V(D)$ is said to *separate* ω and η if $C(X, \omega) \neq C(X, \eta)$. For two distinct ends ω and η of D we call the pair (ω, η) a *limit edge* of D from ω to η if D has an edge from $C(X, \omega)$ to $C(X, \eta)$ for every finite vertex set $X \subseteq V(D)$ that separates ω and η. For a vertex $v \in V(D)$ and an end ω we call the pair (v, ω) a *limit* *edge* of D from v to ω if D has an edge from v to $C(X, \omega)$ for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. Similarly, we call the pair (ω, v) a *limit edge* of D from ω to v if D has an edge from $C(X, \omega)$ to v for every finite vertex set $X \subseteq V(D)$ with $v \notin C(X, \omega)$. For example if R is a ray and every vertex of R sends an edge to a vertex v, then there is a limit (ω, v) from the end ω that is represented by R to v.

The topological space $|D|$ has as its ground set the digraph D, viewed as a 1-complex, together with the ends and limit edges of D. The topology on $|D|$ will be defined formally in Section 9.2.

A *topological path in* $|D|$ is a continuous map $\alpha\colon [0, 1] \to |D|$ that respects the direction of the edges of D when it traverses them. For example, a ray that represents an end ω naturally defines a topological path from its first vertex to ω, and might be extended by a limit edge that starts at ω. A *Hamilton path in* $|D|$ is an injective topological path in $|D|$ that traverses every vertex exactly once. We remark that, as every end of D is a limit point of vertices of D, any topological path that traverses all the vertices of D also traverses all its ends.

There are two trivial obstacles for $|D|$ to containing a Hamilton path. The first is that the cardinality of D may be larger than the cardinality of the unit interval. For this reason we will only consider countable digraphs: these can have continuum many ends, but no more. Another potential obstruction to the existence of a Hamilton path in $|D|$ is that the space $|D|$ may not be compact. As any continuous image of $[0, 1]$ is compact, it is not hard to show that $|D|$ is compact as soon as a topological path traverses all the vertices of D. (We shall disallow parallel edges, as these do not affect hamiltonicity.) For this reason we will only consider those digraphs for which $|D|$ is a compactification of D. These can be described combinatorially, as fallows.

A digraph D is called *solid* if $D - X$ has only finitely many strong components for all finite vertex sets $X \subseteq V(D)$. As shown in the second chapter of this part, see Theorem 7.1, a digraph D is solid if and only if $|D|$ is compact. A vertex v *can reach* a vertex w in D if there is a (finite) path in D from v to w. Our first main theorem of this chapter reads as follows:

Theorem 9.1. *Every countable solid tournament has a Hamilton path. This Hamilton path may be chosen so as to start at any vertex that can reach every other vertex.*

For finite tournaments D there is a standard proof of this result: given a directed path P in D, a quick case distinction shows that any vertex not yet contained in P can be inserted into P. This proof strategy does not adapt easily to infinite tournaments D. It is still possible to insert vertices one after the other to obtain a sequence of longer and longer finite paths which, eventually, contain all the vertices of D. But even if we can show that these paths converge to a topological path in $|D|$ that contains all its vertices, this need not be a Hamilton path by our definition: it might visit some ends multiple times, and thus fail to be injective.

But there is another proof for the finite case, which, as we shall see, can be adapted to infinite digraphs. Every finite tournament D contains a vertex r that can reach every other vertex of D. Let T be a tree obtained by a depth-first search starting at r; note that r can reach every vertex of D even in T. Now T imposes a partial ordering $\leq_T$ on $V := V(D) = V(T)$ defined by letting $v \leq_T w$ if v lies on the path in T from r to w.

As T is a depth-first search tree, this order $\leq_T$ has a linear extension $\leq$ on V in which $u \geq v$ for $\leq_T$-incomparable vertices u, v if our search found u before v. (Note that this differs from when u and v are $\leq_T$-comparable: in that case we have $u \leq_T v$, and hence $u \leq v$, if our search found u before v.)

This is indeed a total order; it is known as the *reverse post-order* and widely used in computer science. Crucially for us one can show that, if v is the predecessor of w in $\leq$, the unique edge of D between them is directed from v to w [19]. Clearly, therefore, this total order on V defines a directed Hamilton path in D. Let us see now how the above proof adapts to infinite digraphs in our topological setting.

In the third chapter of this part, depth-first search trees were adapted to infinite digraphs; these infinite analogues are called *normal* arborescences. (An *arborescence*, in any digraph, is an oriented rooted tree in which the root can reach every vertex.) The notion of normality will be repeated in Section 9.2. For now we only need that normal spanning arborescences exist in every countable solid tournament; they define a tree-order on the vertices as in finite digraphs, and this ordering extends to a total order on its vertices and ends in which (x, y) is an (oriented) edge or limit edge whenever x is the predecessor of y. To prove Theorem 9.1 it then only remains to show that this ordering is continuous at ends, and thus defines a Hamilton path in $|D|$ as before.

To illustrate Theorem 9.1, let us look at an example. Let D be a solid tournament in which the infinite binary tree T is a normal spanning arborescence. Then D has a Hamilton path α which traverses every tree-edge from a vertex v to its right child $v|1$, and every limit edge from an end ω_v represented by a ray $v|1000\ldots$ to the vertex $v|0$; see Figure 9.1.1. This α can be viewed as a limit of the Hamilton paths discussed earlier of the finite subtournaments D_n of D spanned by the subtrees T_n of height n in T, which are depth-first search trees of D_n.

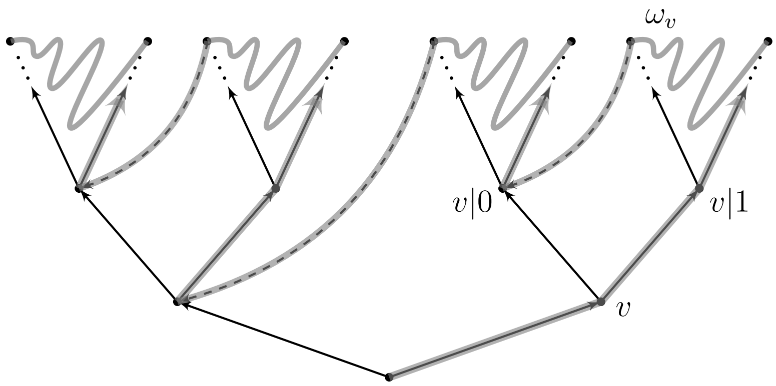

Figure 9.1.1.: A tournament with the infinite binary tree as a normal spanning arborescence. All edges between $\leq_T$-incomparable vertices run from right to left.

A topological path α in $|D|$ is *closed* if $\alpha(0) = \alpha(1)$. A *Hamilton circle* of D is a closed but otherwise injective topological path in $|D|$ that traverses every vertex. As remarked earlier, this implies that it also traverses every end (exactly once).

Theorem 2. *Every countable strongly connected solid tournament has a Hamilton circle.*

As shown in the second chapter of this part, see also [14, Theorem 4, Lemma 5.1], a countable solid digraph D is strongly connected if and only if for any two points $x, y \in |D|$ there is a topological path in $|D|$ from x to y. Again, there is a standard proof of Theorem 9.3 for finite tournaments D: a quick case distinction shows that any vertex not yet contained in a given cycle can be inserted. Again, for infinite D it is possible to insert new vertices, one after the other, into a cycle to obtain at the limit a closed topological path containing all the vertices. But, as earlier, this topological path might traverse ends multiple times.

We will instead use Theorem 9.1 to prove Theorem 2. Ideally, we would like to fix a Hamilton path α in $|D|$ and then use an edge or a limit edge from the endpoint of α to its starting point to obtain a Hamilton circle in $|D|$. This is not always possible, as we are only free to choose the starting point of α. However, an analysis of D and its ends will give us enough control over the endpoint of α to construct the desired Hamilton circle.

This chapter can be read without reading the previous chapters of this part first. We only need a few terms and results from these chapters, which we collect in Section 9.2. We will then prove Theorem 9.2 in Section 9.3, and Theorem 9.3 in Section 9.4.

9.2. Preliminaries

For graph-theoretic terms we follow the terminology in [25]. Throughout this chapter, D is an infinite digraph without infinitely many parallel edges and without loops. We write $V(D)$ for its vertex set, $E(D)$ for its edge set, $\Omega(D)$ for its set of ends and $\Lambda(D)$ for its set of limit edges. For a finite vertex set X and an end ω of D we write $C(X, \omega)$ for the strong component of $D - X$ that contains a tail of every ray that represents ω. We write $\Omega(X, \omega)$ for the set of ends which are represented by a ray in $C(X, \omega)$.

In the following, we give a concise definition of the space $|D|$ and its topology. For a detailed introduction of the space $|D|$ and its topology, see Chapter 7. The ground set of $|D|$ is obtained by taking $V(D) \cup \Omega(D)$ together with a copy of the unit interval $[0, 1]_e$ for every edge or limit edge e of D. Then we identify every vertex or end x with the copy of 0 in $[0, 1]_e$ for which x is the tail of e and with the copy of 1 in $[0, 1]_f$ for which x is the head of f, for all $e, f \in E(D) \cup \Lambda(D)$. Basic open sets of vertices v are uniform stars of radius ε around v, i.e. an ε length from every edge or limit edge that is adjacent to v. Basic open sets of inner points of edges e are open subintervals of $[1, 0]_e$ containing it. For ends ω, basic open sets $\hat{C}_\varepsilon(X, \omega)$ are the union of $C(X, \omega)$ together with $\Omega(X, \omega)$ and every limit edge which has both its endpoints in $C(X, \omega) \cup \Omega(X, \omega)$ and an ε length of every edge or limit edge which has precisely one endpoint in $C(X, \omega) \cup \Omega(X, \omega)$, for finite vertex sets $X \subseteq V(D)$. For inner points z of limit edges (ω, η), basic open sets $\hat{E}_{\varepsilon,z}(X, (\omega, \eta))$ are the union of ε intervals around the copy of z in every edge between $C(X, \omega) \cup \Omega(X, \omega)$ and $C(X, \eta) \cup \Omega(X, \eta)$, for finite vertex sets X which separate ω and η. Similarly, for inner points z of limit edges (v, ω), basic open sets $\hat{E}_{\varepsilon,z}(X, (v, \omega))$ are the union of ε intervals around the copy of z in every edge or limit edge between v and $C(X, \omega) \cup \Omega(X, \omega)$, for finite vertex sets X which contain v. Basic open sets $\hat{E}_{\varepsilon,z}(X, (\omega, v))$ are defined analogously.

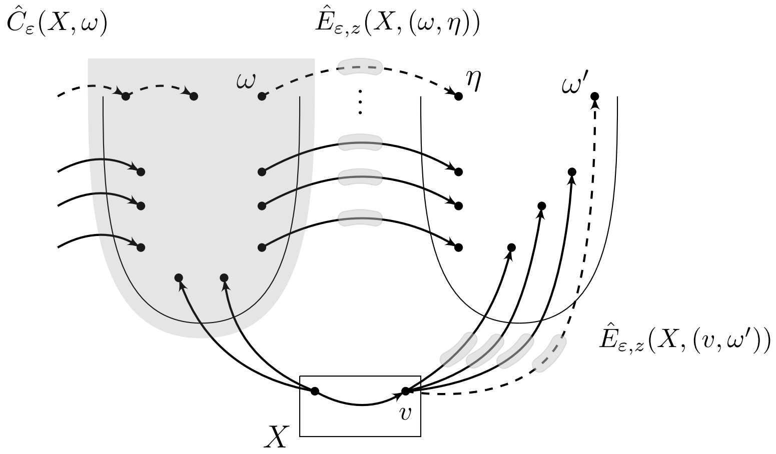

Figure 9.2.1.: Basic open sets for ends and inner points of limit edges.

We will find the desired Hamilton path in Theorem 9.1 via an inverse limit construction. As shown in Section 7.4 for solid D, the space $|D|$ is the inverse limit of its finite contraction minors. In the following, we will give a concise recap of this result for countable digraphs. For the general definition of an inverse system and its inverse limit, see [25, Chapter 8] and for their topological properties, see [63]. Let D be a countable digraph, fix any enumeration of its vertex set and write X_n for the set of the first n vertices. We denote by P_n the partition of $V(D)$ where each vertex in X_n is a singleton partition class and the other partition classes consist of the strong components of $D - X_n$. Every such partition P_n gives rise to a finite (multi-)digraph D/P_n by contracting each partition class and replacing the edges running from a partition class to another by a single edge whenever there are infinitely many. Formally, declare P_n to be the vertex set of D/P_n. Given distinct partition classes $p_1, p_2 \in P_n$, we define an edge (e, p_1, p_2) of D/P_n for every edge $e \in E(D)$ from p_1 to p_2 if there are finitely many such edges. If there are infinitely many edges from p_1 to p_2 we just define a single edge (p_1p_2, p_1, p_2). We call the latter type of edges *quotient edges*. Endowing D/P_n with the 1-complex topology turns it into a compact Hausdorff space, i.e. basic open sets are uniform ε stars around vertices and open subintervals of edges. For $m \leq n$ there is a map $f_{n,m}$ from $V(D/P_n)$ to $V(D/P_m)$, mapping every vertex of D/P_n to the vertex of D/P_m containing it. This map extends naturally to a continuous map from D/P_n to D/P_m by mapping edges of D/P_n to vertices or edges of D/P_m according to the images of its endpoints. This gives an inverse system and if D is solid, its inverse limit coincides with $|D|$:

Corollary 9.2.1 (Corollary 7.4.3)**.** *Let D be a countable solid digraph and let X_n consist of the first n vertices of D with regard to any fixed enumeration of $V(D)$. Then* $|D| \cong \varprojlim (D/P_n)_{n\in\mathbb{N}}$.

Finally, let us recap the notion of normal arborescences from Chapter 8. An *arborescence* is a rooted oriented tree T that contains for every vertex $v \in V(T)$ a directed path from the root to v. The vertices of any arborescence are partially ordered as $v \leq_T w$ if T contains a directed path from v to w. We write $\lfloor v \rfloor_T$ for the up-closure and $\lceil v \rceil_T$ for the down-closure of v in T. Consider a digraph D and a spanning arborescence $T \subseteq D$.

The *normal assistant* of T in D is the auxiliary digraph H that is obtained from T by adding an edge (v, w) for every two $\leq_T$-incomparable vertices $v, w \in V(T)$ for which there is an edge from $\lfloor v \rfloor_T$ to $\lfloor w \rfloor_T$ in D, regardless of whether D contains such an edge. The spanning arborescence T is *normal* in D if the normal assistant of T in D is acyclic; in this case, the transitive closure of the normal assistant defines a partial order $\trianglelefteq_T$ on the vertices of D and we call $\trianglelefteq_T$ the *normal order* of T. We remark that if D is a finite spanning arborescence it is normal in D if and only if it defines a depth-first search tree, see Corollary 8.3.3. One of the most useful properties of normal arborescences is that they capture the separation properties of their host graph, see Lemma 8.3.4:

Lemma 9.2.2. *Let D be any digraph and let $T \subseteq D$ be a normal spanning arborescence in D. If $v, w \in V(T)$ are $\leq_T$-incomparable vertices of T with $w \ntrianglelefteq_T v$, then every path from w to v in D meets $X := \lceil v \rceil_T \cap \lceil w \rceil_T$. In particular, X separates v and w in D.*

Not every digraph has a normal spanning arborescence. However, as a direct consequence of Theorem 8.3, we have that all countable digraphs have one:

Lemma 9.2.3. *Let D be a countable digraph and $r \in V(D)$ a vertex that can reach every other vertex in D. Then D has a normal spanning arborescence rooted in r.*

If T is a normal spanning arborescence of a solid digraph D, then every ray of T is solid in D and therefore represents an end of D. By Lemma 9.2.2, any two distinct rays in T that start at the root represent distinct ends of D. Conversely, we have by Theorem 8.1 that any end of D is represented by a ray in T:

Lemma 9.2.4. *Let D be any digraph and T a normal spanning arborescence of D. Then for every end of D there is exactly one ray in T that represents the end in D and starts at the root of T.*

9.3. Hamilton paths

In this section we prove Theorem 9.1 and give an example which shows that ends and limit edges are crucial for such an extension to exist. As mentioned in the introduction we will find the desired Hamilton path of Theorem 9.1 alongside a normal spanning arborescence of the tournament. This makes it possible to prove a slightly stronger statement and we will need this strengthening in our proof of Theorem 9.3. For a tournament D with a normal spanning arborescence T, we say that an injective topological path α in $|D|$ *respects the normal order* of T if α traverses a vertex t before a vertex t' if and only if t is less than t' in the normal order of T. Note that the normal order of T is a total order, as D is a tournament.

Theorem 9.2. *Let D be a countable solid tournament with a normal spanning arborescence T. Then D has a Hamilton path that respects the normal order of T.*

Note that every solid tournament D has a vertex that can reach every other vertex. Indeed, as D is solid it has only finitely many strong components and one of them sends an edge to every other strong component; every vertex in this component can reach any other vertex in D. Since D is countable, it has a normal spanning arborescence T rooted at any given vertex that can reach every other vertex of D, Lemma 9.2.3. And any Hamilton path that respects the normal order of T starts at the root of T, since it is the smallest element. So the above formulation of Theorem 9.2 implies the formulation in the introduction.

Proof. Our goal is to show that the normal order of T naturally defines a Hamilton path in $|D|$. We will show this by an inverse limit construction. It is straightforward to find a sequence $X_1 \subseteq X_2 \subseteq \dots$ of finite vertex sets of $V(D)$ such that:

(i) the union of all the X_n is $V(D)$ and
(ii) X_n is down-closed in T with regards to the tree-order.

Now, every X_n defines a partition P_n of $V(D)$ and a finite contraction minor D/P_n of D as in Section 9.2. These D/P_n form an inverse system with bonding maps $f_{n,m}$. By the first property (i), the partitions P_n are cofinal in a sequence of partitions that arise by an enumeration of $V(D)$. Hence, the D/P_n form a cofinal (sub-)inverse system of an inverse system that arises by an enumeration of $V(D)$; so the inverse limit of both inverse systems coincides and we have by Corollary 9.2.1 that $|D| \cong \varprojlim(D/P_n)_{n\in\mathbb{N}}$. Next, we will find compatible Hamilton paths in every D/P_n so that the universal property of the inverse limit gives the desired Hamilton path in $|D|$. By the second property (ii) and Lemma 9.2.2, the edges of T in D/P_n form a spanning arborescence T_n of D/P_n. Moreover, as T is normal in D we have that T_n is normal in D/P_n. As D is a tournament we have that the normal order of T_n is a total order on the vertices of D/P_n. Let $v_1 \trianglelefteq_{T_n} \dots \trianglelefteq_{T_n} v_k$ be the sequence of vertices of D/P_n ordered by the normal order of T_n.[1] We claim that $W_n = v_1, \dots, v_k$ is a Hamilton path in D/P_n. Indeed, either v_i and v_{i+1} are $\leq_{T_n}$-comparable in which case (v_i, v_{i+1}) is a tree-edge of T_n or v_i and v_{i+1} are $\leq_{T_n}$-incomparable, in which case there is cross-edge from v_i to v_{i+1} as D is a tournament. These Hamilton paths are compatible in the sense that $f_{n,m}(v_1) \trianglelefteq_{T_m} \dots \trianglelefteq_{T_m} f_{n,m}(v_k)$ is the sequence of vertices of W_m. However, as there might be parallel cross-edges from v_i to v_{i+1}, it might happen that $(f_{n,m}(v_i), f_{n,m}(v_{i+1}))$ does not coincide with the edge from $f_{n,m}(v_i)$ to $f_{n,m}(v_{i+1})$ in W_m. However, for every D/P_n there are only finitely many Hamilton paths with vertex sequence $v_1, \dots, v_k$. So by Kőnig's infinity lemma, we might choose the edges of W_n such that $f_{n,m}(W_n)$ gives W_m.

Finally, fix for every $n \in \mathbb{N}$ a parameterisation $\alpha_n\colon [0,1] \to D/P_n$ of W_n. It is straightforward to choose the α_n in a compatible way, i.e. the projection of α_n coincides with α_{n-1}. Moreover, we may choose the α_n so that they are nowhere constant but on the strong components of $D - X_n$, i.e. on the vertices of D/P_n not in X_n and that the intervals in $[0,1]$ on which α_n is constant have length less or equal to $\frac{1}{n}$. Now, the universal property of the inverse limit gives an injective topological path α. It traverses every vertex of D as the X_n contain every vertex eventually, so α is a Hamilton path in $|D|$. To show that α respects the normal order of T, consider two vertices $t \trianglelefteq_T t'$ and choose $n \in \mathbb{N}$ so that $t, t' \in X_n$. Then $t \trianglelefteq_{T_n} t'$ and α_n traverses t before t'. As the projection of α to D/P_n gives α_n, we have that α traverses t before t'. Conversely, if α traverses t before t' a we have that α_n traverses t before t' in D/P_n for $n \in \mathbb{N}$ with

[1] We remark that this is the reverse post-order of T_n.

$t, t' \in X_n$. Hence, we have $t \trianglelefteq_{T_n} t'$ and as the normal order of T induces the normal order of T_n, we have that $t \trianglelefteq_T t'$. □

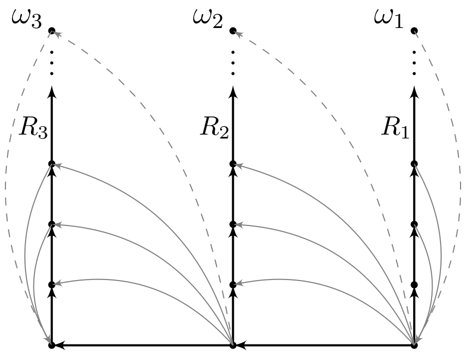

Figure 9.3.1.: A countable solid tournament together with a normal spanning arborescence (black edges). All grey edges along a branch are oriented from top to bottom, and all edges between any two branches are oriented from right to left.

Example 9.3.1. The tournament D in Figure 9.3.1 satisfies the following properties:

(i) There is no spanning ray or double ray in D.
(ii) There is no Hamilton path in the space formed by D and its end compactification of the underlying undirected graph.
(iii) There is no Hamilton path in $|D|$ that avoids all inner points of limit edges of D.

Proof. To (i): There are no edges into $V(R_1)$ and there are no edges out of $V(R_3)$; hence, any ray or double ray that contains a vertex of R_1 and a vertex of R_3 contains only finitely many vertices of R_2.

To (ii): The underlying undirected graph of D is a clique; hence, it has exactly one end ω. Moreover, its end compactification by the ends of the underlying undirected graph coincides with its one-point compactification D^*. So the image of any Hamilton path in D^* defines either a spanning (reverse) ray or the disjoint union of a ray and a reverse ray containing together every vertex of D. A similar argument as in (i) shows that there are no such (reverse) rays in D.

To (iii): Suppose for a contradiction that D has a Hamilton path α that avoids all inner points of limit edges of D. Then one of the ends ω_i is not an endpoint of α. So the image of α contains a ray that represents ω_i and a reverse ray that contains infinitely many vertices of R_i. However, D contains no such ray and reverse ray. □

9.4. Hamilton circles

In this section, we prove Theorem 9.3. For this we need some definitions and two Lemmas. The reverse subrays of a reverse ray are its *tails*. A reverse ray R in a digraph D *represents* an end ω if there is a solid ray R' in D that represents ω such that R and R' have a tail in the same strong component of $D - X$ for every finite vertex set $X \subseteq V(D)$.

It is straightforward to show that a reverse ray R that represents an end ω defines a topological path from ω to the first vertex of R. For a (reverse) ray $R = v_1, v_2, \dots$ the subpaths of the form $v_1, \dots, v_n$ are the *finite initial segments* of R. For a double ray $W = \dots, w_{-1}, w_0, w_1, \dots$ we denote by $W_{n<}$ the ray $w_{n+1}, w_{n+2}, \dots$ and by $W_{<n}$ the reverse ray $\dots, w_{n-2}, w_{n-1}$.

We say that a vertex v of D can be *inserted* into a (reverse) ray $R = v_1, v_2, \dots$ if there is a path P that starts (ends) at v such that $v_1, \dots, v_{i-1}, P, v_i, \dots$ is a (reverse) ray in D, for some $i \in \mathbb{N}$. Similarly, we say that a vertex v of D can be *inserted* into a double ray $W = \dots, w_{-1}, w_0, w_1, \dots$ if there is a path P that contains v such that $\dots, w_{i-1}, P, w_i, \dots$ is a double ray in D, for some $i \in \mathbb{Z}$. A quick case distinction shows:

Lemma 9.4.1. *Let D be a strongly connected tournament. Any given vertex can be inserted into any given (reverse) ray. Any given vertex with finite in- or finite out-degree can be inserted into any given double ray.* □

There is a natural partial order on the set of ends of a digraph D. For two ends $\omega, \eta \in \Omega(D)$ write $\omega \leq_\Omega \eta$ if there are rays R_ω and R_η that represent ω and η respectively, such that there are infinitely many disjoint paths from R_ω to R_η. This gives a (well-defined) partial order on $\Omega(D)$. If (ω, η) is a limit edge of D then clearly $\omega \leq_\Omega \eta$; the converse is false in general. If D is a tournament then any two ends of D are comparable, so $\leq_\Omega$ gives a total order on $\Omega(D)$.

Lemma 9.4.2. *Let D be any countable solid tournament, then $\Omega(D)$ has a greatest and a least element.*

Proof. First note that $\Omega(D)$ is non-empty; it is straightforward to construct a ray in D since the deletion of any finite vertex set leaves only finitely many strong components and every ray in a solid digraph is solid. We show that $\Omega(D)$ has a greatest element; the proof for the least element is analogue.

Fix an enumeration of $V(D)$ and let X_n denote the set of the first n vertices. We have that $|D| \cong \varprojlim (D/P_n)_{n \in \mathbb{N}}$ by Corollary 9.2.1. Now, consider the strong components of $D - X_n$. We may view them as partially ordered by $C_1 \leq C_2$ if there is a path in $D - X_n$ from C_1 to C_2. As D is a tournament, this gives a total order on the strong components of $D - X_n$. Hence, for every X_n there is a greatest strong component C_n of $D - X_n$ with regard to the aforementioned order of strong components. This strong component is a vertex in D/P_n, and this choice of vertices is compatible in the sense that C_n includes C_{n+1} as a subset. So this choice of vertices gives a point in the inverse limit, which in turn corresponds to a point ω in $|D|$. It is straightforward to show that this point ω is an end of D and we claim that it is the greatest element of $\Omega(D)$. Indeed, for any other end $\eta \in \Omega(D)$ there is a X_n that separates ω and η. Now, $C(X_n, \omega)$ and $C(X_n, \eta)$ are distinct strong components of $D - X_n$, and by the choice of ω we have that $C(X_n, \omega)$ is greater than $C(X_n, \eta)$. Consequently, there are only finitely many disjoint paths from a ray representing ω to a ray representing η. □

Theorem 9.3. *Every countable strongly connected solid tournament has a Hamilton circle.*

Proof. First note that for every vertex $v \in V(D)$ and any end $\omega \in \Omega(D)$ the tournament D has a limit edge from v to ω or vice versa. Indeed, for a ray R that represents ω the vertex v sends infinitely many edges to $V(R)$ or receives infinitely many edges from $V(R)$. Furthermore, for a vertex $v \in V(D)$ there is a limit edge from some end of D to v if and only if v has infinite in-degree, and there is a limit edge from v to some end of D if and only if v has infinite out-degree. We split the proof into two main cases:

First case: The tournament D has only one end ω. In this case, first suppose that every vertex of D has infinite in- and out-degree. Then it is straightforward to construct a spanning ray R. The first vertex v of R receives a limit edge from ω, and following first R to ω and then the limit edge (ω, v) yields the desired Hamilton circle in $|D|$. We remark that in this situation it is also possible to construct a spanning reverse ray or a spanning double ray to obtain a Hamilton circle in $|D|$.

Second, suppose that there is a vertex, v say, of D that has finite in- or finite out-degree. We discuss the case where v has finite in-degree, the other case follows by considering the reverse of D. Fix a normal spanning arborescence T of D rooted at v, Lemma 9.2.3, and apply Theorem 9.2 to D and T to obtain a Hamilton path α in $|D|$ that starts at v. As D is solid and one-ended, T has exactly one ray R_ω that starts at v, Lemma 9.2.4. All vertices of R_ω are traversed by α before ω, in particular the vertices that are traversed by α before ω form a ray R, in the order in which they are traversed by α. Conversely, the vertices that are traversed by α after ω form a reverse path or a reverse ray $\overleftarrow{R}$. Furthermore, every vertex v' of $\overleftarrow{R}$ has finite out-degree as there are only finitely many vertices greater than v' in the normal order. If $\overleftarrow{R}$ is a reverse path we are done by Lemma 9.4.1, as we can insert all the vertices of $\overleftarrow{R}$ into R_ω one after another to obtain a spanning ray and then use a limit edge from ω to its start vertex to obtain the desired Hamilton circle in $|D|$. If $\overleftarrow{R}$ is a reverse ray there is an edge from some vertex of $\overleftarrow{R}$ to v, since v has finite out-degree. Consequently, we find a double ray W that contains R and all but finitely many vertices of $\overleftarrow{R}$. By Lemma 9.4.1, any vertex of $\overleftarrow{R}$ not yet contained in W can be inserted into W to obtain a spanning double ray of D. As D has only one end, a spanning double ray naturally defines a Hamilton circle in $|D|$.

Second case: The tournament D has more than one end. For the rest of the proof denote by ω_* the least and by ω^* the greatest end of D, Lemma 9.4.2. Our first goal is to find a double ray W such that its finite subpaths separate every vertex and every other end from ω_* and ω^* respectively, and such that its subrays represent ω_* and its reverse subrays represent ω^*. In order to find such a double ray, fix a normal spanning arborescence T of D. Then there is exactly one ray R_{ω_*} in T that starts at the root of T and represents ω_* in D, Lemma 9.2.4. This ray R_{ω_*} has the property that its finite initial segments separate ω_* from every vertex and every other end eventually, Lemma 9.2.2. Now, consider $\overleftarrow{D}$ the reverse of D and fix a normal spanning arborescence $\overleftarrow{T}$ of $\overleftarrow{D}$. The ends of D and $\overleftarrow{D}$ are in a one-to-one correspondence in that the reverse of every ray in D represents an end in $\overleftarrow{D}$. Let R_{ω^*} be the unique ray in $\overleftarrow{T}$ that starts at the root and represents the least end of $\overleftarrow{D}$, Lemma 9.2.4. Then the reverse $\overleftarrow{R}_{\omega^*}$ of R_{ω^*} is a reverse ray in D that represents ω^* and its finite initial segments separate ω^* from every vertex and every other end eventually, Lemma 9.2.2. As $\omega_* \neq \omega^*$, we have that R_{ω_*} and $\overleftarrow{R}_{\omega^*}$ have only finitely many vertices in common. Consequently, there is a double ray W' that contains a tail of R_{ω_*} and a tail of $\overleftarrow{R}_{\omega^*}$. Let S be the set of all the vertices of R_{ω_*} and $\overleftarrow{R}_{\omega^*}$ not contained in W'. If we can insert all those vertices of S into W' that are not separated from ω_* or ω^* by any finite subpath of W' then we obtain our desired double ray. So let $s \in S$ be a vertex that is not separated from ω_* or ω^* by any finite subpath of W'. If s can be separated from ω_* but not from ω^* by a finite subpath of W', or vice versa, then it is straightforward to insert s into W', Lemma 9.4.1. So we may assume that s cannot be separated from ω_* and from ω^* by any finite subpath of W'. Furthermore, we may assume that for $W' = \dots, w_{-1}, w_0, w_1, \dots$ we have that s receives an edge from w_n for $0 \leq n$ and sends an edge to all the other vertices of W', otherwise s can clearly be inserted into W'. As $\omega_* <_\Omega \omega^*$, there is a finite subpath $W_N = w_{-N}, \dots, w_N$ of W', for $2 \leq N$, such that there is no edge from $W_{<N}$ to $W_{N<}$. Now, there is a non-trivial path P in $C(W_N, \omega_*) = C(W_N, \omega^*)$ from $W_{<N}$ to $W_{N<}$. If P contains s it can be inserted into $W_{N<}$ via a subpath of P (and into $W_{N<}$). If s is not contained in P consider any inner vertex v of P, then s can be inserted into $W_{N<}$ or $W_{<N}$ depending on whether D contains the edge (s, v) or the edge (v, s). Hence, we obtain a double ray W such that its finite subpaths separate every vertex and every other end from ω_* and ω^* respectively and such that its subrays represent ω_* and its reverse subrays represent ω^*.

Our next goal is to show that there is even a double ray with the defining properties of W that contains all vertices which send a limit edge to ω_* and all vertices that receive a limit edge from ω^*. We will show this claim by inserting all these vertices, not yet contained in W, one after the other into W in such a way that the limit is still a double ray.

Denote by S_* all vertices not in W which send a limit edge to ω_*, and by S^* all vertices not in W which receive a limit edge from ω^*. Note that S_* and S^* are disjoint, as there is a finite subpath of W that separates ω_* and ω^*. First consider S_* and choose $N \in \mathbb{N}$ so that W_N separates ω_* and ω^*. It is straightforward to check that all vertices in S_* that are separated from ω_* by W_N can be inserted into W without changing $W_{N<}$ or $W_{<N}$. Now, for any other vertex $s \in S_*$ there is a smallest $n(x) \in \mathbb{N}$ such that $W_{n(s)}$ separates s from ω_*. Again, it is straightforward to check that all vertices in S_* with index $n(s)$ can be inserted into W by a path from $w_{n(s)}$ to $w_{n(s)+1}$. An analogue technique shows that all vertices in S^* can be inserted into W. As we substituted only edges of W by a path at most once, we end up in the limit step with a double ray.

So let us assume that W additionally has the property to contain $S_* \cup S^*$. Our final goal is to find an injective topological path α from ω_* to ω^* that contains precisely the vertices not in W. Having α at hand, the desired Hamilton circle in $|D|$ is obtained by first following α and then following W.

Consider the strong components of $D - W$, for any such strong component there is a finite subpath W_n of W such that C is a strong component of $D - W_n$. Indeed, for every $v \in C$ there is an $n(v) \in \mathbb{N}$ such that $W_{n(v)}$ separates v from ω_* and ω^* and this $n(v)$ has to be the same for any two vertices in C. Moreover, these strong components are totally ordered in that every vertex of C sends an edge to any vertex of C', or vice versa, for any two strong components of $D - W$. For all strong components of $D - W$

fix a Hamilton path α_C in $|C|$ (or in C if it is finite). Now, all these Hamilton paths can be linked up to the desired injective topological path α, see Figure 9.4.1. Indeed, if C is the predecessor of C' in the aforementioned order of strong components of $D-W$, then there is an edge or limit edge from the endpoint of α_C to the starting point of $\alpha_{C'}$. Moreover, if there is a least element, C_* say, of the strong components of $D-W$, then ω_* sends a limit edge to every vertex of C_*. Similarly, if there is a greatest element, C^* say, then every vertex of C^* sends a limit edge to ω^*.

Conversely, if there is no greatest element, then the strong components of $D-X$ converge to ω^* in that traversing the α_C one after the other in their total order yields an injective continuous path that ends at ω^*. Similarly, if there is no least element, then the strong components of $D-X$ converge to ω_* in that traversing the α_C one after the other in their inverted total order yields an injective continuous path that starts at ω_*. Note that the open sets $\hat{C}_\varepsilon(W_n,\omega_*)$ and $\hat{C}_\varepsilon(W_n,\omega^*)$ form a neighbourhood base for ω_* and ω^*, respectively. This topological path traverses all the vertices of $D-W$ as W contains $S_*\cup S^*$. We remark that if there are no strong components of $D-W$, i.e. W is spanning then we obtain a directed topological path from ω_* to ω^* that avoids W by the limit edge (ω_*,ω^*). □

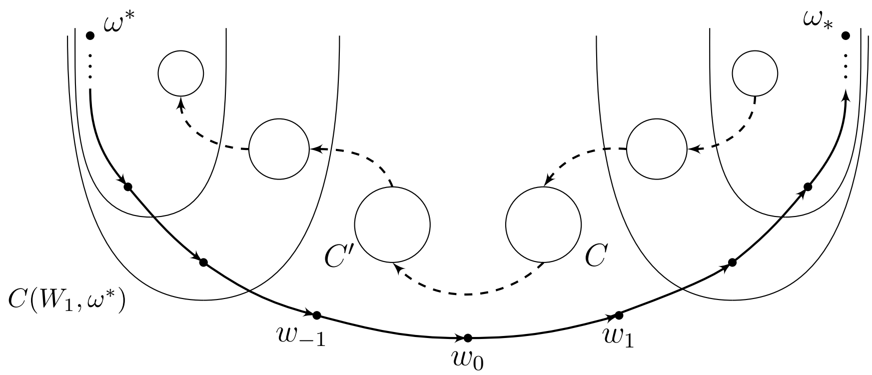

Figure 9.4.1.: Strong components of $D-W$ are indicated as circles. Strong components of the form $C(W_n,\omega_*)$ or $C(W_n,\omega^*)$ are indicated as parabolas, which might contain strong components of $D-W$ not yet separated by W_n.

Appendix

10. English summary

Part I:

Chapter 2. We show that every connected graph can be approximated by a normal tree, up to some arbitrarily small error phrased in terms of neighbourhoods around its ends. The existence of such approximate normal trees has consequences of both combinatorial and topological nature.

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

On the topological side, we show that the end space $\Omega(G)$, as well as the spaces $|G| = G \cup \Omega(G)$ naturally associated with a graph G, are always paracompact. This gives unified and short proofs for a number of results by Diestel, Sprüssel and Polat, and answers an open question about metrizability of end spaces by Polat.

Chapter 3. The *directions* of an infinite graph G are a tangle-like description of its ends: they are choice functions that choose compatibly for all finite vertex sets $X \subseteq V(G)$ a component of $G - X$.

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

Curiously, there exist graphs whose directions are all countably determined but which cannot be distinguished all at once by countably many choices.

We structurally characterise the graphs whose directions can be distinguished all at once by countably many choices, and we structurally characterise the graphs which admit no such countably many choices. Our characterisations are phrased in terms of normal trees and tree-decompositions.

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

Part II:

Chapter 4. We show that for every infinite collection $\mathcal{R}$ of disjoint equivalent rays in a graph G there is a subdivision of the hexagonal half-grid in G such that all its vertical rays belong to $\mathcal{R}$. This result strengthens Halin's grid theorem by giving control over which specific set of rays is used, while its proof is significantly shorter.

Chapter 5. An end of a graph G is an equivalence class of rays, where two rays are equivalent if there are infinitely many vertex-disjoint paths between them in G. The degree of an end is the maximum cardinality of a collection of pairwise disjoint rays in this equivalence class.

Halin conjectured that the end degree can be characterised in terms of certain typical ray configurations, which would generalise his famous *grid theorem*. In particular, every end of regular uncountable degree κ would contain a *star of rays*, i.e. a configuration consisting of a central ray R and κ neighbouring rays $(R_i \colon i < \kappa)$ all disjoint from each other and each R_i sending a family of infinitely many disjoint paths to R so that paths from distinct families only meet in R.

We show that Halin's conjecture fails for end degree $\aleph_1$, holds for $\aleph_2, \aleph_3, \ldots, \aleph_\omega$, fails for $\aleph_{\omega+1}$, and is undecidable (in ZFC) for the next $\aleph_{\omega+n}$ with $n \in \mathbb{N}$, $n \geq 2$. Further results include a complete solution for all cardinals under GCH, complemented by a number of consistency results.

Part III:

In this part we develop an end space theory for directed graphs.

Chapter 6. As for undirected graphs, the *ends* of a digraph are points at infinity to which its rays converge. Unlike for undirected graphs, some ends are joined by *limit edges*; these are crucial for obtaining the end space of a digraph as a natural (inverse) limit of its finite contraction minors.

As our main result in this first chapter of this part, we show that the notion of *directions* of an undirected graph, a tangle-like description of its ends, extends to digraphs: there is a one-to-one correspondence between the 'directions' of a digraph and its ends and limit edges.

In the course of this we extend to digraphs a number of fundamental tools and techniques for the study of ends of graphs, such as the star-comb lemma and Schmidt's ranking of rayless graphs.

Chapter 7. Here in the second chapter of this part, we introduce the topological space $|D|$ formed by a digraph D together with its ends and limit edges. We then characterise those digraphs that are compactified by this space. Furthermore, we show that if $|D|$ is compact, it is the inverse limit of finite contraction minors of D.

To illustrate the use of this, we extend to the space $|D|$ two statements about finite digraphs that do not generalise verbatim to infinite digraphs. The first statement is the characterisation of finite Eulerian digraphs by the condition that the in-degree of every vertex equals its out-degree. The second statement is the characterisation of strongly connected finite digraphs by the existence of a closed Hamilton walk.

Chapter 8. Here in the third chapter of this part, we introduce a concept of depth-first search trees in infinite digraphs, which we call *normal spanning arborescences*.

We show that normal spanning arborescences are end-faithful: every end of the digraph is represented by exactly one ray in the normal spanning arborescence that starts from the root. We further show that this bijection extends to a homeomorphism between the end space of a digraph D, which may include limit edges between ends, and the end space of any normal arborescence with limit edges induced from D. Finally, we prove a Jung-type criterion for the existence of normal spanning arborescences.

Chapter 9. We prove that for all countable tournaments D its compactification $|D|$ by its ends and limit edges contains a topological Hamilton path: a topological arc that contains every vertex. If D is strongly connected, then $|D|$ contains a topological Hamilton circle.

These results extend well-known theorems about finite tournaments, which we show do not extend to the infinite in a purely combinatorial setting.

11. Deutsche Zusammenfassung

Chapter 2.

Wir zeigen, dass jeder zusammenhängende Graph durch normale Bäume approximierbar ist, bis auf einen beliebig kleinen Fehler um seine Enden. Die Existenz solcher Bäume hat Konsequenzen sowohl kombinatorischer als auch topologischer Natur.

Als kombinatorisches Resultat zeigen wir, dass jeder Graph einen normalen Spannbaum hat sobald er lokal einen normalen Spannbaum um jedes seiner Enden hat, d.h. das einzige Hindernis für einen Graphen, einen normalen Spannbaum zu haben, ist ein Ende, dessen sämtliche Nachbarschaften keinen Spannbaum besitzen.

Als topologische Resultate zeigen wir, dass der Endenraum $\Omega(G)$ sowie auch der Raum $|G| = G \cup \Omega(G)$ parakompakt sind. Dies ergibt vereinheitlichte und kurze Beweise von einer Reihe an Resultaten von Diestel, Sprüssel and Polat. Darüber hinaus beantwortet dies eine Frage von Polat über die Metrisierbarkeit von Endenräumen.

Chapter 3.

Die *Richtungen* eines unendlichen Graphen G sind knäuelartige Beschreibung seiner Enden: Sie sind Auswahlfunktionen, die in einer konsistenten Weise zu jeder endlichen Eckenmenge $X \subseteq V(G)$ eine Komponente von $G - X$ auswählen.

Auch wenn jede Richtung von einem Strahl induziert wird gibt es Richtungen von Graphen, die nicht eindeutig durch irgendeine abzählbare Teilmenge ihrer Auswahlen bestimmt sind. Wir charakterisieren diese Richtungen und jene Richtungen, die durch abzählbar viele ihrer Auswahlen bestimmt sind, mithilfe von sternartigen Teilstrukturen und Strahlen des Graphen.

Interessanterweise gibt es Graphen, für die jede Richtung eindeutig durch abzählbar viele ihrer Auswahlen bestimmt ist, aber für die es keine abzählbare Menge an Auswahlen gibt, die alle Richtungen unterscheidet.

Wir geben eine strukturelle Charakterisierung aller Graphen, deren Richtungen sich durch abzählbar viele Auswahlen unterscheiden lassen und wir geben eine strukturelle Charakterisierung aller Graphen, für die es keine solche abzählbaren Auswahlen gibt. Unsere Charakterisierungen sind formuliert mithilfe von normalen Bäumen und Baumzerlegungen.

Unsere vier (sub)strukturellen Charakterisierungen implizieren kombinatorische Charakterisierungen jener vier Klassen von unendlichen Graphen, deren Endenräume das erste und zweite Abzählbarkeitsaxiom erfüllen: die zwei Klassen von Graphen, deren Endenräume das erste oder zweite Abzählbarkeitsaxiom erfüllen und die Komplemente dieser zwei Klassen.

Part II:

Chapter 4.

Wir zeigen, dass es für jede unendliche Familie $\mathcal{R}$ von disjunkten äquivalenten Strahlen in einem Graphen G eine Unterteilung des hexagonalen Halbgitters in G gibt, sodass all seine vertikalen Strahlen zu $\mathcal{R}$ gehören. Dieses Resultat verstärkt Halins Gittersatz, indem es Kontrolle über die zuvor spezifizierten Strahlen gibt. Unser Beweis für dieses stärkere Resultat ist deutlich kürzer als Halins Beweis.

Chapter 5.
Ein Ende eines Graphen G ist eine Äquivalenzklasse von Strahlen, in der zwei Strahlen äquivalent sind, wenn es in G unendlich viele eckendisjunkte Wege zwischen ihnen gibt. Der Grad eines Endes ist die größte Kardinalzahl einer Familie von paarweise disjunkten Strahlen in dem Ende.

Halin vermutete, dass der Grad eines Endes durch typische Konfigurationen von Strahlen charakterisiert werden kann, was seinen berühmten *Gittersatz* verallgemeinern würde. Insbesondere enthielte jedes Ende mit regulärem überabzählbaren Grad κ einen *Stern von Strahlen*, d.h. eine Konfiguration von Strahlen bestehend aus einem zentralen Strahl R und κ vielen benachbarten Strahlen $(R_i : i < \kappa)$, welche alle paarweise disjunkt sind und wo jeder Strahl R_i eine Familie von unendlich vielen disjunkten Wegen zu R schickt, sodass Wege von unterschiedlichen Familien sich höchstens in R treffen.

Wir zeigen, dass Halins Vermutung falsch ist für Enden mit Grad $\aleph_1$, dass sie wahr ist für Enden mit Grad $\aleph_2, \aleph_3, \ldots, \aleph_\omega$, dass sie wieder falsch ist für Enden mit Grad $\aleph_{\omega+1}$ und dass sie unentscheidbar (in ZFC) ist für Enden mit Grad $\aleph_{\omega+n}$ für $n \in \mathbb{N}$ und $n \geq 2$. Unter anderem geben wir eine vollständige Lösung unter GCH an sowie einige Konsistenzresultate.

Part III:
In diesem Teil der Arbeit entwickeln wir eine Endenraum-Theorie für gerichtete Graphen.

Chapter 6.
Wie für ungerichtete Graphen sind *Enden* von gerichteten Graphen Punkte im Unendlichen, zu welchen ihre Strahlen konvergieren. Anders als bei ungerichteten Graphen sind einige Enden durch *Limes-Kanten* verbunden; diese sind wesentlich, um den Endenraum eines gerichteten Graphen auf natürliche Weise als (inversen) Limes seiner endlichen Kontraktionsminoren zu erhalten.

Als unser Hauptresultat im ersten Kapitel von diesem Teil der Arbeit zeigen wir, dass der Begriff der *Richtung* eines ungerichteten Graphen, welcher eine knäuelartige Beschreibung seiner Enden ist, sich auf gerichtete Graphen erweitern lässt: Es gibt eine Eins-zu-eins-Beziehung zwischen den 'Richtungen' eines gerichteten Graphen und seinen Enden und Limes-Kanten.

Im Zuge dessen erweitern wir eine ganze Reihe von Werkzeugen und Techniken für ungerichtete Graphen auf gerichtete Graphen, zum Beispiel das Stern-Kamm-Lemma und Schmidts Rangfunktion für strahlenlose Graphen.

Chapter 7.
Im zweiten Kapitel von diesem Teil der Arbeit stellen wir den topologischen Raum $|D|$ vor, den ein gerichteter Graph D zusammen mit seinen Enden und Limes-Kanten bildet. Wir charakterisieren dann genau die gerichteten Graphen D, welche durch $|D|$ kompaktifiziert werden. Darüber hinaus zeigen wir, dass wenn der Raum $|D|$ kompakt

ist, er der inverse Limes der endlichen Kontraktionsminoren von D ist.

Um zu zeigen wie nützlich dies ist erweitern wir zwei Resultate über endliche gerichtete Graphen auf den Raum $|D|$, welche nicht wörtlich für unendlich gerichtete Graphen gelten. Das erste Resultat ist die Charakterisierung endlicher Euler-Graphen durch die Bedingung, dass der Eingrad einer jeden Ecke ihrem Ausgrad entspricht. Das zweite Resultat ist die Charakterisierung der stark zusammenhängenden endlichen gerichteten Graphen durch die Existenz eines geschlossenen Hamilton-Kantenzuges.

Chapter 8.

Im dritten Kapitel von diesem Teil der Arbeit stellen wir ein Konzept der Tiefensuche-Bäume für unendliche gerichtete Graphen vor, welche wir *normale aufspannende Arboreszenzen* nennen.

Wir zeigen, dass normale aufspannende Arboreszenzen endentreu sind: Jedes Ende eines gerichteten Graphen ist durch genau einen Strahl in der normalen aufspannenden Arboreszenz, welcher bei ihrer Wurzel startet, repräsentiert. Darüber hinaus zeigen wir, dass diese Bijektion einen Homöomorphismus zwischen dem Endenraum eines gerichteten Graphen D, welcher möglicherweise Limes-Kanten enthält, und dem Endenraum der normalen aufspannenden Arboreszenz mit den Limes-Kanten von D definiert. Schießlich zeigen wir ein Jung-artiges Kriterium für die Existenz von normalen aufspannenden Arboreszenzen.

Chapter 9.

Wir zeigen, dass die Kompaktifizierung $|D|$ durch die Enden und Limes-Kanten eines jeden abzählbaren Turniers D einen topologischen Hamiltonweg enthält: eine topologische Kurve, die jede Ecke enthält. Ist D stark zusammenhängend, dann enthält $|D|$ sogar einen topologischen Hamiltonkreis.

Diese Resultate erweitern wohlbekannte Sätze über endliche Turniere, die sich, wie wir zeigen werden, nicht im rein kombinatorischen Sinne auf unendliche Turniere erweitern lassen.

12. Publications related to this dissertation

The following articles are related to this dissertation:

Part I:

(i) Chapter 2 is based on [54].
(ii) Chapter 3 is based on [52].

Part II:

(iii) Chapter 4 is based on [53].
(iv) Chapter 5 is based on [39].

Part III:

(v) Chapter 6 is based on [13].
(vi) Chapter 7 is based on [14].
(vii) Chapter 8 is based on [15].
(viii) Chapter 9 is based on [57].

13. Declaration of my contributions

Part I:

Chapter 2 This chapter is based on the paper [54] that I wrote together with Jan Kurkofka and Max Pitz. We conducted the research in this chapter together. I noticed that we can not only show that end spaces are ultra-paracompact, Corollary 2.3.1, but that our cover also refines the prescribed components, yielding the formulation of our main result in this chapter. I drafted a proof of Corollary 2.3.2. I noticed all the consequences of the approximating result and drafted the Section 2.4.

Chapter 3 This chapter is based on the paper [52] that I wrote together with Jan Kurkofka. We conducted the research in this chapter together. We drafted Sections 3.1 and 3.2 together. I wrote an early draft for Section 3.3 and Jan Kurkofka wrote an early draft for Section 3.4. I wrote an early draft for Section 3.5

Part II:

Chapters 4. This chapter is based on the paper [53] that I wrote together with Jan Kurkofka and Max Pitz. I came up with the presented proof. Max Pitz wrote a draft of the paper. I created figure 4.1.1.

Chapters 5 This chapter is based on the paper [39] that I wrote together with Stefan Geschke, Jan Kurkofka and Max Pitz. We conducted the research in this chapter together. I drafted an early version of Section 5.2 (excluding the second proof of Lemma 5.2.1). I drafted an early version of Section 5.3. I drafted an early version of Section 5.5.

Part III:

Chapter 6-8 This part is based on a series of three papers [13–15] that I wrote together with Carl Bürger. The research was conducted in close collaboration and we share an equal amount of work: both in developing the main ideas and in drafting the papers. The project started with a new approach to ends of digraphs. Carl Bürger told me about his ideas and invited me to join the project. After that I developed the idea of limit edges. Carl Bürger first drafted Chapter 6 and I wrote the first draft of Chapter 7. We drafted the introduction and preliminaries of Chapter 8 together. I first drafted Section 8.5.

Chapter 9 This chapter is based on the paper [57] that I wrote entirely on my own.

Acknowledgement

First of all, I want to thank my supervisor Reinhard Diestel for his support, guidance and honesty. Also, I would like to thank Max Pitz, it is always a pleasure to work with you and you have taught me countless things.

I am grateful for my colleagues and friends Carl Bürger and Jan Kurkofka; together, we found new directions to unknown ends and climbed the highest trees to the largest suns.

Finally, on a more personal note, I want to say thank you to my family and those friends who stand by my side since our early school days, you are like family to me.

Bibliography

[1] J. Bang-Jensen and G. Gutin, *Digraphs: theory, algorithms and applications*, Springer Science & Business Media, 2008. ↑7.1

[2] E. Berger and H. Bruhn, *Eulerian edge sets in locally finite graphs*, Combinatorica **31** (2011), no. 1, 21–38, DOI 10.1007/s00493-011-2572-0. ↑9.1

[3] N. Bowler, J. Carmesin, P. Komjáth and Christian Reiher, *The Colouring Number of Infinite Graphs*, Combinatorica **39** (2019), no. 6, 1225–1235, DOI 10.1007/s00493-019-4045-9. ↑5.4

[4] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, P. Pitz and M. Teegen, *Ubiquity in graphs II: Ubiquity of graphs with non-linear end structure* (2018), available at arXiv:1809.00602. Submitted. ↑4.1

[5] N. Bowler, J. Erde, P. Heinig, F. Lehner and M. Pitz, *Non-reconstructible locally finite graphs*, J. Combin. Theory (Series B) **133** (2018), 122–152, DOI 10.1016/j.jctb.2018.04.007. ↑1

[6] N. Bowler G. Geschke and M. Pitz, *Minimal obstructions for normal spanning trees*, Fund. Math. **241** (2018), 245–263, DOI 10.4064/fm337-10-2017. ↑1.1.1, 2.1, 5.4, 5.7

[7] J-M. Brochet and R. Diestel, *Normal tree orders for infinite graphs*, Transactions of the American Mathematical Society **345** (1994), no. 2, 871–895, DOI 10.1090/S0002-9947-1994-1260198-4. ↑5.6.1

[8] H. Bruhn and R. Diestel, *Duality in infinite graphs*, Comb., Probab. & Comput. **15** (2006), 75–90, DOI 10.1017/S0963548305007261. ↑1.3, 6.1, 7.1, 9.1

[9] H. Bruhn, R. Diestel, A. Georgakopoulos and P. Sprüssel, *Every rayless graph has an unfriendly partition*, Combinatorica **30** (2010), no. 5, 521–532, DOI 10.1007/s00493-010-2590-3. ↑6.1

[10] H. Bruhn and M. Stein, *MacLane's planarity criterion for locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), 225–239, DOI 10.1016/j.jctb.2005.07.005. ↑9.1

[11] H. Bruhn and X. Yu, *Hamilton circles in planar locally finite graphs*, SIAM J. Discrete Math. **22** (2008), 1381–1392, DOI 10.1137/050631458. ↑9.1

[12] C. Bürger and J. Kurkofka, *Duality theorems for stars and combs I: Arbitrary stars and combs* (2020), available at arXiv:2004.00594. Submitted. ↑1.1.2, 3.1, 3.2.2, 3.2.3, 3.2.4, 3.4, 6.1, 6.3

[13] C. Bürger and R. Melcher, *Ends of digraphs I: Basic theory* (2020), available at arXiv:2009.03295. Submitted. ↑7.1, (v), 13

[14] ———, *Ends of digraphs II: The topological point of view* (2020), available at arXiv:2009.03293. Submitted. ↑8.1, 8.5, 9.1, (vi), 13

[15] ———, *Ends of digraphs III: Normal arborescences* (2020), available at arXiv:2009.03292. Submitted. ↑(vii), 13

[16] P. Camion, *Chemins et circuits hamiltoniens des graphes complets*, Comptes Rendus de l'Académie des Sciences de Paris **249** (1959), 2151–2152. MR0122735 ↑1.3.4, 9.1

[17] C. Carathéodory, *Über die Begrenzung einfach zusammenhängender Gebiete*, Math. Annalen **73** (1913), no. 3, 323–370, DOI 10.1007/BF01456699. ↑6.1

[18] J. Carmesin, M. Hamann and B. Miraftab, *Canonical trees of tree-decompositions* (2020), available at arXiv:2002.12030. Submitted. ↑3.2.3

[19] T. H. Cormen, C. E. Leiserson, R. L. Rivest and C. Stein, *Introduction to algorithms*, MIT press, 2009. MR2572804 ↑8.1, 9.1

[20] Q. Cui, J. Wang and X. Yu, *Hamilton circles in infinite planar graphs*, J. Combin. Theory (Series B) **99** (2009), 110–138, DOI 10.1016/j.jctb.2008.05.002. ↑9.1

[21] R. Diestel, *End spaces and spanning trees*, J. Combin. Theory (Series B) **96** (2006), no. 6, 846–854, DOI 10.1016/j.jctb.2006.02.010. ↑1.1.2, 2.1, 2.4, 2.4, 3.1, 3.1, 3.4, 7.1, 9.1

[22] ———, *Graph Theory*, 3th, Springer, 2005. MR2159259 ↑1.3, 6.1, 7.1, 9.1

[23] ———, *Ends and Tangles*, Abh. Math. Sem. Univ. Hamburg **87** (2017), no. 2, 223–244, DOI 10.1007/s12188-016-0163-0, available at arXiv:1510.04050v3. Special issue in memory of Rudolf Halin. ↑1.1.2, 3.1

[24] ———, *A short proof of Halin's grid theorem*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **74** (2004), no. 1, 237–242, DOI 10.1007/BF02941538. ↑4.1, 5.1.1

[25] ———, *Graph Theory*, 5th, Springer, 2016. MR3822066 ↑1, 1.1.1, 1.3.3, 1.4, 2.1, 2.2, 3.1, 3.1, 3.2, 3.2.1, 3.2.2, 3.2.3, 4.1, 5.1.1, 5.2, 5.2, 6.1, 6.1, 6.2, 6.3, 6.3, 6.4, 7.1, 7.1, 7.2, 7.2, 7.5, 8.1, 8.1, 8.1, 8.1, 8.2, 8.4, 9.1, 9.2, 9.2

[26] ———, *Locally finite graphs with ends: a topological approach*, Discrete Math. **310–312** (2010), 2750-2765 (310); 1423-1447 (311); 21–29 (312), available at arXiv:0912.4213v3. ↑7.1

[27] ———, *The end structure of a graph: recent results and open problems*, Disc. Math. **100** (1992), no. 1, 313–327, DOI 10.1016/0012-365X(92)90650-5. ↑2.1, 2.4, 2.5, 8.5

[28] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, J. Combin. Theory (Series B) **87** (2003), 197–206, DOI 10.1016/S0095-8956(02)00034-5. ↑1.1.2, 2.1, 2.2.1, 3.1, 3.1, 6.4, 9.1

[29] ———, *On Infinite Cycles I*, Combinatorica **24** (2004), 68-89, DOI 10.1007/s00493-004-0005-z. ↑6.1, 7.1

[30] R. Diestel and I. Leader, *Normal spanning trees, Aronszajn trees and excluded minors*, J. London Math. Soc. **63** (2001), 16–32, DOI 10.1112/S0024610700001708. ↑1.1.1, 2.1, 5.1.3, 5.6.2, 5.6.2, 5.6.2, 5.7

[31] R. Diestel and J. Pott, *Dual trees must share their ends*, J. Combin. Theory (Series B) **123** (2017), 32–53, DOI 10.1016/j.jctb.2016.11.005. ↑9.1

[32] J.A. Dieudonné, *Une généralisation des espaces compacts*, J. Math. Pures. Appl. **23** (1944), 65–76. MR0013297 ↑2.5

[33] R. Engelking, *General Topology*, 2nd ed., Sigma Series in Pure Mathematics, vol. 6, Heldermann Verlag, Berlin, 1989. MR1039321 ↑1.4, 2.2, 2.4, 2.4, 2.5, 3.2, 3.4, 7.2

[34] J. Erde, F. Lehner and M. Pitz, *Hamilton decompositions of one-ended Cayley graphs*, J. Combin. Theory (Series B) **140** (2020), 171–191, DOI 10.1016/j.jctb.2019.05.005. ↑9.1

[35] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, Mathematische Zeitschrift **33** (1931), no. 1, 692–713, DOI 10.1007/BF01174375. ↑6.1

[36] P. Gartside and M. Pitz, *Eulerian spaces* (2019), available at arXiv:1904.02645. Submitted. ↑1, 9.1

[37] A. Georgakopoulos, *Infinite Hamilton cycles in squares of locally finite graphs*, Advances in Mathematics **220** (2009), no. 3, 670-705, DOI 10.1016/j.aim.2008.09.014. ↑1.3, 6.1, 7.1, 9.1

[38] ———, *Topological circles and Euler tours in locally finite graphs*, Electronic J. Comb. **16** (2009), no. R40, DOI 10.37236/129. ↑9.1

[39] S. Geschke, J. Kurkofka, R. Melcher, and M. Pitz, *Halin's end degree conjecture*, available at arXiv:2010.10394. Submitted. ↑4.1, (iv), 13

[40] M. Gitik and M. Magidor, *The singular cardinal hypothesis revisited.*, Set theory of the continuum. Papers presented at the Mathematical Sciences Research Institute (MSRI) workshop, Berkeley, California, USA, October 16-20, 1989, 1992, pp. 243–279. MR1233822 ↑5.7

[41] J.P. Gollin and K. Heuer, *Characterising k-connected sets in infinite graphs* (2018), available at arXiv:1811.06411. Submitted. ↑5.3

[42] R. Halin, *Miscellaneous Problems on Infinite Graphs*, J. Graph Theory **35** (2000), 128–151, DOI 10.1002/1097-0118(200010)35:2<128::AID-JGT6>3.0.CO;2-6. ↑5.1.1

[43] ———, *A problem in infinite graph-theory*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, 1975, pp. 79–84, DOI 10.1007/BF02995937. ↑4.1

[44] ———, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Math. Nachr. **30** (1965), no. 1–2, 63–85, DOI 10.1002/mana.19650300106. ↑1.2.1, 4.1, 4.1, 4.1, 5.1.1, 5.1.1

[45] ———, *Über unendliche Wege in Graphen*, Math. Annalen **157** (1964), 125–137, DOI 10.1007/BF01362670. ↑1.1.2, 3.1, 6.1

[46] K. Heuer, *Hamiltonicity in locally finite graphs: two extensions and a counterexample*, Electronic J. Comb. **25** (2018), no. P3.13, DOI 10.37236/6773. ↑9.1

[47] ———, *Excluding a full grid minor*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, 2017, pp. 265–274, DOI 10.1007/s12188-016-0165-y. ↑4.1

[48] B. Hughes, *Trees and ultrametric spaces: a categorical equivalence*, Advances in Mathematics **189** (2004), no. 1, 148–191, DOI 10.1016/j.aim.2003.11.008. ↑2.4

[49] T. Jech, *Set theory, The Third Millennium Edition*, Springer Monographs in Mathematics, 2013. MR1940513 ↑1.4, 5.2, 5.4, 5.4, 5.7, 5.7, 5.7, 5.7

[50] H.A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, Math. Nachr. **41** (1969), 1–22, DOI 10.1002/mana.19690410102. ↑2.2.2, 3.2.4, 5.2, 8.1

[51] ———, *Zusammenzüge und Unterteilungen von Graphen*, Mathematische Nachrichten **35** (1967), no. 5-6, 241–267, DOI 10.1002/mana.19670350503. ↑5.6.2

[52] J. Kurkofka and R. Melcher, *Countably determined ends and graphs* (2020), available at arXiv:2101.07247. Submitted. ↑(ii), 13

[53] J. Kurkofka, R. Melcher and M. Pitz, *A strengthening of Halin's grid theorem* (2021), available at arXiv:2104.10672. Submitted. ↑5.1.1, (iii), 13

[54] ———, *Approximating infinite graphs by normal trees*, J. Combin. Theory (Series B) **148** (2021), 173–183, DOI 10.1016/j.jctb.2020.12.007. ↑1.1.2, 3.1, 9.1, (i), 13

[55] J. Kurkofka and M. Pitz, *Ends, tangles and critical vertex sets*, Math. Nachr. **292** (2019), no. 9, 2072–2091, DOI 10.1002/mana.201800174, available at arXiv:1804.00588. ↑1.1.2, 3.1

[56] ———, *Tangles and the Stone-Čech compactification of infinite graphs*, J. Combin. Theory (Series B) **146** (2021), 34–60, DOI 10.1016/j.jctb.2020.07.004, available at arXiv:1806.00220. ↑1.1.2, 3.1

[57] R. Melcher, *Hamiltonicity in infinite tournaments* (2021), available at arXiv:2101.05264. Submitted. ↑7.1, 8.1, (viii), 13

[58] V. Neumann-Lara, *The dichromatic number of a digraph*, Journal of Combinatorial Theory, Series B **33** (1982), no. 3, 265–270, DOI 10.1016/0095-8956(82)90046-6. ↑6.3, 8.3

[59] M. Pitz, *Proof of Halin's normal spanning tree conjecture* (2020), available at arXiv:2005.02833. to appear in Israel J. Math. ↑5.4

[60] N. Polat, *Ends and multi-endings I*, J. Combin. Theory (Series B) **67** (1996), 86–110, DOI 10.1006/jctb.1996.0035. ↑2.1, 2.2, 2.3, 2.4, 3.1, 3.5, 6.1, 9.1

[61] ———, *Ends and multi-endings II*, J. Combin. Theory (Series B) **1** (1996), 56–86, DOI 10.1006/jctb.1996.0057. ↑2.1, 2.4, 3.1, 6.1, 9.1

[62] L. Rédei, *Ein kombinatorischer Satz*, Acta Scientiarum Mathematicarum **7** (1934), 39–43. ↑1.3.4, 9.1

[63] L. Ribes and P. Zalesskii, *Profinite Groups*, Springer, 2010. MR2599132 ↑7.2, 7.2, 9.2

[64] N. Robertson, P. Seymour and R. Thomas, *Excluding infinite clique minors*, Vol. 566, American Mathematical Society, 1995. MR1303095 ↑4.1

[65] R. Schmidt, *Ein Ordnungsbegriff für Graphen ohne unendliche Wege mit einer Anwendung auf n-fach zusammenhängende Graphen*, Arch. Math. **40** (1983), no. 1, 283–288, DOI 10.1007/BF01192782. ↑6.1, 6.3

[66] S. Shelah, *Cardinal Arithmetic*, Oxford University Press, 1994. MR1318912 ↑5.7

[67] P. Sprüssel, *End spaces of graphs are normal*, J. Combin. Theory (Series B) **98** (2008), 798–804, DOI 10.1016/j.jctb.2007.10.006. ↑2.1, 2.3, 2.5, 3.1, 9.1

[68] M. Stein, *Arboricity and tree-packing in locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), 302–312, DOI 10.1016/j.jctb.2005.08.003. ↑9.1

[69] S. Todorčević, *Walks on ordinals and their characteristics*, Vol. 263, Springer Science & Business Media, 2007. MR2355670 ↑5.6.1

[70] J. Zuther, *Planar strips and an end concept for digraphs*, Technische Universität Berlin, 1996. ↑1.3, 6.1, 6.3

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.