

Connectivity and tree structure in infinite graphs and digraphs

Dissertation

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

vorgelegt
im Fachbereich Mathematik

von
Jochen Pascal Gollin

Hamburg

2019

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

Datum der Disputation: 17. Juli 2019

In loving memory of Günter Gollin.

Contents

1. Introduction	1
1.1. Historical Background	1
1.2. Overview of the chapters in this dissertation	3
1.3. Preliminaries	8
1.3.1. Basic notation	8
1.3.2. The compactness principle in combinatorics	12
I. Undirected graphs	14
2. Representations of infinite tree sets	15
2.1. Introduction	15
2.2. Separation systems	17
2.3. Regular tame tree sets and graph-theoretical trees	22
2.3.1. The edge tree set of a tree	22
2.3.2. The tree of a regular tame tree set	23
2.3.3. Regular tame tree sets and trees – A characterisation	25
2.4. Regular tree sets and tree-like spaces	27
2.4.1. Graph-like spaces	27
2.4.2. Tree-like spaces	30
2.4.3. The edge tree set of a tree-like space	32
2.4.4. The tree-like space of a tree set	33
2.4.5. Regular tree sets and tree-like spaces – A characterisation	35
2.5. Proof of Theorem 2.4.4	37
3. Infinite end-devouring sets of rays with prescribed start vertices	46
3.1. Introduction	46
3.2. Theorem	48

3.3. Ends of uncountable degree	52
4. Characterising k-connected sets in infinite graphs	53
4.1. Introduction	53
4.2. Preliminaries	57
4.3. Typical graphs with k -connected sets	63
4.3.1. k -typical graphs	63
4.3.2. Generalised k -typical graphs	70
4.3.3. Statement of the Main Theorem	74
4.4. k -connected sets, minors and topological minors	75
4.5. Structure within ends	76
4.5.1. End defining sequences and combined end degree	76
4.5.2. Constructing uniformly connected rays	81
4.6. Minors for regular cardinalities	82
4.6.1. Complete bipartite minors	82
4.6.2. Minors for regular k -blueprints	85
4.6.3. Characterisation for regular cardinals	86
4.7. Minors for singular cardinalities	87
4.7.1. Cofinal sequence of regular bipartite minors with disjoint cores.	87
4.7.2. Frayed complete bipartite minors	88
4.7.3. Minors for singular k -blueprints	92
4.7.4. Characterisation for singular cardinals	98
4.8. Applications of the minor-characterisation	99
4.9. Nested separation systems	101
II. Directed graphs	105
5. An analogue of Edmonds' branching theorem for infinite di- graphs	106
5.1. Introduction	106
5.2. Preliminaries	110
5.2.1. Topological notions for undirected multigraphs	112
5.2.2. Topological notions for digraphs	114

5.2.3. Basic lemmas	115
5.3. Fundamental statements about topological directed walks in locally finite digraphs	116
5.4. Packing pseudo arborescences	120
5.5. Structure of pseudo-arborescences	122
6. On the infinite Lucchesi-Younger conjecture	128
6.1. Introduction	128
6.2. Basic notions and tools	132
6.2.1. Cuts and dicuts	133
6.2.2. Dijoins and optimal pairs for classes of finite dibonds . .	135
6.2.3. Stars and combs	136
6.2.4. Finitely separable multigraphs	140
6.2.5. Quotients	141
6.2.6. Quotients of rayless digraphs	144
6.2.7. Cost functions and feasible potentials	145
6.2.8. More on dicuts	148
6.3. Comparing Conjecture 6.1.3 with Conjecture 6.1.5	149
6.4. Reductions for the Infinite Lucchesi-Younger Conjecture	152
6.5. Special cases	154
6.5.1. Finite parameters	154
6.5.2. Every edge lies in only finitely many dibonds and reductions to this case	157
6.6. A matching problem about infinite hypergraphs	159
6.7. Extending an algorithmic proof of Frank	162
6.7.1. Tightness and an auxiliary graph	163
6.7.2. Feasible dijoins	167
6.7.3. The auxiliary graph for a feasible dijoin	174
6.7.4. The standard decomposition of F -tight dicuts	175
6.7.5. The auxiliary graph for a healthy feasible dijoin	177
6.7.6. F -tightly corner-closed classes of finite dibonds	177
6.8. Applications of Frank's proof	180
6.8.1. Finite parameters in finite-corner-closed classes	180
6.8.2. Every edge lies on only finitely many dibonds.	180

6.8.3. No infinite dibond in arbitrary corners	181
6.8.4. Dibonds of bounded size	183
6.8.5. Dibonds of minimum size	183
6.8.6. Another class of digraphs	184

Appendix **189**

English summary	189
Deutsche Zusammenfassung	191
Publications related to this dissertation	193
Declaration on my contributions	194
Acknowledgement	196
Bibliography	197
Eidesstattliche Versicherung	201

1. Introduction

1.1. Historical Background

The notion of connectivity is one of the most fundamental concepts in graph theory. Different aspects of connectivity play a role across almost every area of graph theory, going beyond the most basic graph invariants of vertex- or edge-connectivity.

Probably the most fundamental characterisation of connectivity in finite graphs is Menger's Theorem. It states that for any two vertex sets in a finite graph the maximum number of disjoint paths between the two vertex sets equals the minimum size of another vertex set separating the two vertex sets. In fact, there is a structural reformulation due to Erdős of this quantitative description of this dual nature of connectivity: for any two vertex sets in a finite graph there exists a set of disjoint paths between them and another vertex set separating them that consists of precisely one vertex from each of the paths.

While Erdős showed that a simple quantitative generalisation of Menger's Theorem to infinite graphs when just considering cardinalities of these sets is quite easy, Aharoni and Berger [2] proved the conjecture of Erdős that the structural version of the theorem holds for infinite graphs as well. This theorem had a big impact on the development of infinite connectivity theory.

In the area of structural graph theory, the study of the duality between connectivity and tree structure is a common theme. Such type of duality theorems assert a dichotomy between the existence of a 'highly connected part' in a graph and the non-existence of some kind of tree structure with certain properties, which, if it exists, clearly precludes the existence of such a 'highly connected part'.

The most prominent such tree structures with regard to finite graphs are tree-decompositions. A *tree-decomposition of a graph G* consists of a *decomposition tree T* as well as for each node t of the tree a *part P_t* , that is an induced subgraph of G , satisfying the following two properties. These parts cover G , and they are

organised ‘like the decomposition tree’ in the following way. We demand that for every edge tt' of the decomposition tree the set $V(P_t) \cap V(P_{t'})$ separates in G the union of the parts of the component of $T - tt'$ containing t from the union of the parts of the component of $T - tt'$ containing t' .

This last property says that an edge of the decomposition tree naturally defines a separation of the graph. A *separation of the graph* is an ordered pair of non-empty vertex sets of that graph, called the *sides of that separation*, such that the subgraphs induced by the sides cover the whole graph. The *order* of the separation is then the size of the intersection of its sides.

Any tree-decomposition of a graph into small parts witnesses that the graph cannot contain, for example, large cliques or grids, or large clique or grid minors. All these dense objects in a graph have the property that they orient the low-order separations of the graph by lying ‘mostly’ on one side of any given low-order separation. For such a dense structure in a graph these orientations of separations are consistent with each other: no two of them ‘disagree’ about where the dense object lies by being oriented away from each other.

This led Robertson and Seymour to the introduction of tangles as an interpretation of such ‘highly connected parts’ of a graph [41]. Formally, *tangles* are just orientations of the low-order separations of a graph satisfying certain consistency conditions. This more abstract look at these ‘highly connected parts’ inspired much of the recent research into the aforementioned type of duality theorems.

In infinite graphs, Halin’s concept of an end defines such a tangle. An *end* of a graph is an equivalence class of rays, i.e. one-way infinite paths, where two rays are said to be equivalent, if they cannot be separated by a finite vertex set. Hence these ends define an orientation of all the finite-order separations of a graph, which turns out to be such a tangle, and different ends define different such tangles.

The notion of tree-decomposition has only limited use in infinite graphs with respect to these kind of duality theorems, since tree-decompositions with certain properties may not exist for reasons other than containing the ‘highly connected part’ which its existence would preclude. In the same manner as with tangles, the more abstracted view regarding separations has turned out to be very useful in this context as well. While we saw above that each edge of a tree-decomposition naturally defines a separation of the graph, the set of all these separations yields a *nested separation system*, where any two separations that the edges induce are

nested in the sense that one side of the first separation is a subset of a side of the second separation, and the other side of the second separation is a subset of the other side of the first separation. In finite graphs, these nested separation systems turn out to be equivalent to tree-decompositions in the sense that these objects can be translated into each other. Many duality theorems in infinite graphs feature these nested separation system as the relevant tree structure.

1.2. Overview of the chapters in this dissertation

This dissertation deals with different aspects of connectivity and tree structure in infinite graphs, which make it part of the research area of structural infinite graph theory. It consists of two parts: simple graphs are considered in Part I, and directed graphs (or *digraphs*) are considered in Part II. Part I consists of Chapters 2, 3 and 4, while Part II consists of Chapters 5 and 6.

We now give a brief overview of the results of these chapters, although each of them will feature its own more comprehensive introduction.

1.2.1. Chapter 2: Representations of infinite tree sets

Tree sets are abstract structures that can be used to model various tree-shaped objects in combinatorics. In a more axiomatic way they generalise the notion of nested separation systems we mentioned earlier. Their definition is based on the notion of an *abstract separation system*, which consists of a partially ordered set, whose elements we call (abstract) separations, together with an order-reversing involution. Finite tree sets can be represented by finite graph-theoretical trees in the same way as nested separation system of graphs can. In this chapter we extend this representation theory to infinite tree sets.

In the first part of the chapter, we characterise those tree sets that can be represented by infinite trees; these are precisely those tree sets which are regular, i.e. no separation is comparable with its reverse orientation, and which do not contain a chain of order type $\omega + 1$.

Then we introduce and study a topological generalisation of infinite trees which can have ‘limit edges’, so called *tree-like spaces*. The definition of these tree-like spaces is based on graph-like spaces introduced by Thomassen and Vella [49], and

further studied by Bowler, Carmesin and Christian [6]. Tree-like spaces give rise to regular tree sets in a natural way too, and we show that every regular tree set can be represented by a tree-like space.

Thus, tree sets turn out to be the fitting analogue to the tree-decompositions of finite graphs for capturing the tree-likeness of infinite graphs that occurs in their connectivity duality theorems mentioned earlier.

1.2.2. Chapter 3: Infinite end-devouring sets of rays with prescribed start vertices

In this chapter, we turn our focus towards the notion of ends.

An interesting property of the rays in a normal spanning tree of a graph is that it meets every ray in the end ω it is contained in. We say a ray with this property *devours the end* ω . Georgakopoulos [23] introduced this concept for families of rays: Given an end ω of a graph, we call a family of rays ω -*devouring* if every ray in ω meets at least one of the rays in the family.

In his efforts to look for Hamilton circles in locally finite graphs, Georgakopoulos proved the following result [23]. Given a finite family of disjoint rays in an end ω of countable degree, there exists a family of disjoint rays with the exact same set of start vertices which devours the end ω .

Georgakopoulos then asked the question, whether this result can be generalised to infinite families of rays. Since any maximal disjoint family of rays in an end trivially devours that end, the difficulties in this question lie mainly with the prescribed set of start vertices.

In this chapter we prove this conjecture of Georgakopoulos with a construction that independently gives a proof of his original result. Afterwards we discuss the problems that may arise for possible generalisations of this result to ends of uncountable degree.

1.2.3. Chapter 4: Characterising k -connected sets in infinite graphs

One aspect of connectivity in graphs is the concept of vertex sets that are k -connected in that graph for a positive integer k . A set X of vertices of a graph G is

k -connected in G if any two of its subsets of the same size $\ell \leq k$ can be connected by ℓ disjoint paths in the whole graph.

These k -connected sets in finite graphs have been studied in connection with tree-width first by Robertson, Seymour and Thomas [43], and later by Diestel, Gorbunov, Jensen and Thomassen [13]. More recently, Geelen and Joeris [22,31] proved a duality theorem about these k -connected sets in finite graphs with tree-decompositions whose adhesion is less than k . (The *adhesion* of a tree-decomposition is the maximum order of the separations induced by the edges of the decomposition tree.)

The first main result of this chapter is a generalisation of this duality theorem to infinite graphs: we prove for a positive integer k and a cardinal κ that if a graph contains no k -connected set of κ vertices, then there is a nested separation system containing only separations of order less than k which has width less than κ , where we will define the width of a nested separation system as a natural analogue of the width of a tree-decomposition. Once more, such a tree structure is a natural obstruction to the existence of k -connected sets of cardinality κ .

Geelen and Joeris also provided a structural description of these k -connected sets in finite graphs in terms of certain ‘unavoidable’ minors. As the second main result we generalise this result to infinite graphs as well. For fixed positive integer k and cardinal κ , we find a finite list of ‘unavoidable’ minors (as well as a finite list of ‘unavoidable’ topological minors) such that a graph contains a k -connected set of size κ if and only if it contains one of these finitely many graphs as a (topological) minor. This extends earlier work of Halin [29], as well as of Oporowski, Oxley and Thomas [37] on such a question with slightly different connectivity notions.

1.2.4. Chapter 5: An analogue of Edmonds’ branching theorem for infinite digraphs

In this chapter of the dissertation we focus on a different aspect of the connection between high connectivity and trees, namely, tree packing theorems. Independently, Nash-Williams [36] and Tutte [51] proved a famous tree-packing theorem, which implies that every finite $2k$ -edge-connected undirected multigraph has k edge-disjoint spanning trees. A counterexample by Aharoni and Thomassen [3] shows that a straightforward generalisation of this result to infinite graphs fails.

Spanning trees in a graph have as a defining property that they are precisely those minimal edge sets of that graph which meet all its ends. While Tutte [51] observed that his packing result can be generalised to infinite graphs via packings with minimal edge sets that meet all finite cuts, this generalisation remained not very well motivated until the rise of topological infinite graph theory. This started in Hamburg around 2000, when Diestel and his group developed a topological framework where they consider the Freudenthal compactification of a locally finite graph G . The edge sets considered by Tutte turned out to be precisely the topological spanning trees in this setting.

In finite directed graphs there is a similar packing result by Edmonds [19]. A *cut* in a digraph is the edge set between a bipartition of the vertex set of the digraph, where we refer to the bipartition classes as the *sides* of the cut. *Edmonds branching theorem* says that for a packing of k edge disjoint spanning arborescences, i.e. spanning trees that are directed away from a fixed common root r , if and only if every cut of the digraph contains at least k edges directed from the side containing the root to the other side. The example of Aharoni and Thomassen [3] also shows that a straightforward generalisation of this result to infinite digraphs fails. While there exist results by Thomassen and Joó respectively that generalise Edmonds' result to certain classes of infinite digraphs, in this chapter we focus on a generalisation in the spirit of Tutte's approach.

We introduce the notion of *spanning pseudo-arborescences*. These are edge sets which are minimal in containing from every cut an edge that is directed from the side containing the root to the other side. We prove a corresponding packing result. Finally, we verify some tree-like properties for these objects, but give also an example that their underlying graphs do not in general correspond to topological trees in the Freudenthal compactification of the underlying multigraph of the digraph. For this we generalise several concepts of this topological framework to directed graphs.

1.2.5. Chapter 6: On the infinite Lucchesi-Younger conjecture

In this final chapter of the dissertation we consider an aspect of connectivity exclusive to digraphs, namely, strong connectivity. A digraph is *strongly connected*

if for any two vertices there exist directed paths in both directions between these vertices. There is a corresponding notion of *weak connectivity*, where we only require that the underlying undirected multigraph of a digraph is connected.

A trivial obstruction for a digraph to being strongly connected is the existence of a cut in that digraph whose edges are all directed from the same side to the other side. In fact, a digraph is strongly connected, if and only if no such *dicut* exists.

We call an edge set a *dijoin* if it meets every dicut of the digraph. Contracting the edges of a dijoin always yields a strongly connected digraph. Thus, in a sense, these dijoins provide a kind of measure on ‘how far away’ a digraph is from being strongly connected.

A well-known min-max theorem of Lucchesi and Younger [35] states that in every finite digraph the least size of a dijoin equals the maximum number of disjoint dicuts in that digraph. As with Menger’s theorem, there is an obvious structural reformulation of this theorem, similar to the approach of Erdős. It says that in every finite digraph there is a set of disjoint dicuts together with a dijoin consisting of precisely one edge from each of the dicuts in that set.

After giving an example that a straightforward generalisation of this theorem to infinite digraphs fails, we work on a conjecture of Heuer stating that every digraph contains a set of disjoint finite dicuts together with an edge set meeting every finite dicut, which we call a *finitary dijoin*, consisting of precisely one edge from each of the dicuts in that set. We call this conjecture the Infinite Lucchesi-Younger conjecture.

One of the main results of this chapter is that it suffices to prove the conjecture for countable digraphs whose underlying undirected multigraph is 2-connected. Moreover, we verify several special cases of the conjecture.

In the finite case, one can always choose the set of disjoint dicuts to be nested. Thus, they define a directed tree structure in the same way as nested separation systems do. While we provide some evidence that a nested version of Heuer’s conjecture may be strictly stronger than his original conjecture, all our results are applicable to the nested conjecture as well.

1.3. Preliminaries

1.3.1. Basic notation

For basic facts about finite and infinite graphs we refer the reader to [9]. If not stated differently, we also use the notation of [9]. Especially for facts and notations about directed graphs we refer to [4].

Throughout this thesis we will consider different types of graphs. For a graph G we denote by $V(G)$ the vertex set, and by $E(G)$ its edge set.

In Part I we will consider simple graphs, i.e. undirected graphs with no multiple edges or loops. We write an edge as a string uv of its endvertices u and v .

In Part II we will consider mostly directed graphs, which we also call *digraphs*, but also undirected multigraphs. In general, we allow our digraphs to have parallel edges, but no loops if we do not explicitly mention them. Similarly, all undirected multigraphs we consider do not have loops if nothing else is explicitly stated. Sometimes we write uv for edge directed from vertex u to vertex v , although this might not uniquely determine an edge in case of parallel edges. In parts where a finer distinction becomes important we shall clarify the situation, and we will point out specifically any difficulties that may arise by this abuse of notation. For an edge uv of a digraph we furthermore denote the vertex u as the *tail* of uv and v as the *head* of uv . We denote the underlying undirected multigraph of a digraph D by $\text{Un}(D)$.

In this thesis we consider both finite and infinite cardinals. As usual, for an infinite cardinal κ we define its *cofinality*, denoted by $\text{cf } \kappa$, as the smallest infinite cardinal λ such that there is a set $X \subseteq \{Y \subseteq \kappa \mid |Y| < \kappa\}$ such that $|X| = \lambda$ and $\bigcup X = \kappa$. We distinguish infinite cardinals κ to *regular cardinals*, i.e. cardinals where $\text{cf } \kappa = \kappa$, and *singular cardinals*, i.e. cardinals where $\text{cf } \kappa < \kappa$. Note that $\text{cf } \kappa$ is always a regular cardinal. For more information on infinite cardinals and ordinals, we refer the reader to [33].

Let G be a graph. For two disjoint vertex sets X, Y of a graph G we denote by $E_G(X, Y)$ the set of all edges of G having one of their endvertices in X and the other in Y . Moreover, if G is a digraph, we define

$$\vec{E}_G(X, Y) := \{xy \in E(X, Y) \mid x \in X \text{ and } y \in Y\}.$$

We make the following definitions for a set $X \subseteq V(G)$.

- $\delta_G(X) := E(X, V(G) \setminus X)$, the *set of incident edges of X* ;
- $N_G(X) := \{y \in V(G) \setminus X \mid \exists x \in X : xy \in \delta_G(X) \text{ or } yx \in \delta_G(X)\}$, the *neighbourhood of X* ;
- $\partial_G(X) := \{x \in X \mid \exists y \in V(G) \setminus X : xy \in \delta_G(X) \text{ or } yx \in \delta_G(X)\}$, the *boundary of X* ; and
- $d_G(X) := |\delta_G(X)|$, the *degree of X* ;

as well as for specifically digraphs

- $\delta_G^-(X) := \overrightarrow{E}(V(G) \setminus X, X)$, the *set of in-going edges of X* ;
- $\delta_G^+(X) := \overrightarrow{E}(X, V(G) \setminus X)$, the *set of out-going edges of X* ;
- $N_G^-(X) := \{y \in V(G) \setminus X \mid \exists x \in X : yx \in \delta_G^-(X)\}$, the *in-neighbourhood of X* ;
- $N_G^+(X) := \{y \in V(G) \setminus X \mid \exists x \in X : xy \in \delta_G^+(X)\}$, the *out-neighbourhood of X* ;
- $\partial_G^-(X) := \{x \in X \mid \exists y \in V(G) \setminus X : yx \in \delta_G^-(X)\}$, the *in-boundary of X* ;
- $\partial_G^+(X) := \{x \in X \mid \exists y \in V(G) \setminus X : xy \in \delta_G^+(X)\}$, the *out-boundary of X* ;
- $d_G^-(X) := |\delta_G^-(X)|$, the *in-degree of X* ;
- $d_G^+(X) := |\delta_G^+(X)|$, the *out-degree of X* ; and

We will usually omit the subscript if the graph we are talking about is clear from the context. If $X = \{v\}$ is a singleton we will replace X by v for these notions.

A vertex of in-degree 0 in a digraph is a *source*, and a vertex of in-degree 1 in a digraph is a *sink*.

We call a graph G *locally finite* if each vertex of G has finite degree.

Let G and H be two graphs. The *union* $G \cup H$ of G and H is the graph with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H)$. The *Cartesian product* $G \times H$ of G and H is the graph with vertex set $V(G) \times V(H)$ such that two

vertices $(g_1, h_1), (g_2, h_2) \in V(G \times H)$ are adjacent if and only if either $h_1 = h_2$ and $g_1g_2 \in E(G)$ or $g_1 = g_2$ and $h_1h_2 \in E(H)$ holds.

Unless otherwise specified, a *path* in this thesis is a finite path. The *length* of a path is the size of its edge set. A path is *trivial*, if it only contains only one vertex, which we will call its *endvertex*. Otherwise, the two vertices of degree 1 in the path are its *endvertices*. The other vertices are called the *inner vertices* of the path.

Let $A, B \subseteq V(G)$ be two (not necessarily disjoint) vertex sets in a graph G . An A - B *path* is a path whose inner vertices are disjoint from $A \cup B$ such that one of its end vertices lies in A and the other lies in B . In particular, a trivial path whose endvertex is in $A \cap B$ is also an A - B path. For convenience, by a slight abuse of notation, if $A = \{a\}$ (or $B = \{b\}$) is a singleton we will replace A by a (or B by b respectively) for this notion.

A one-way infinite path R is called a *ray* and a two-way infinite path D is called a *double ray*. The unique vertex of degree 1 of R is its *start vertex*. A subgraph of R (or D) that is a ray itself is called a *tail* of R (or D respectively).

Given a path or ray P containing two vertices v and w we denote the unique v - w path in P by vPw .

Given a ray R and $v \in V(R)$, we write vR for the tail of R with start vertex v . A finite path $P \subseteq R$ (or $P \subseteq D$) is a *segment* of R (or D respectively). If v and w are the endvertices of P , then we denote P also by vRw (or vDw respectively). If v is the end vertex of vRw whose distance is closer to the start vertex of R , then v is called the *bottom vertex of vRw* and w is called the *top vertex of vRw* . If additionally v is the start vertex of R , then we call vRw an *initial segment* of R and denote it by Rw .

We use the following notion to abbreviate *concatenations* of paths and rays. Let P be a v - w path for two vertices v and w , and let Q be either a ray or another path such that $V(P) \cap V(Q) = \{w\}$. Then we write PQ for the path or ray $P \cup Q$, respectively. We omit writing brackets when stating concatenations of more than two paths or rays.

An *end* of a graph G is an equivalence class of rays, where two rays are *equivalent* if they cannot be separated by deleting finitely many vertices of G . We denote the set of ends of G by $\Omega(G)$. A ray being an element of an end $\omega \in \Omega(G)$ is called an ω -*ray*. A double ray all whose tails are elements of ω is called an ω -*double ray*.

For an end $\omega \in \Omega(G)$ let $\deg(\omega)$ denote the *degree* of ω , that is the supremum of the set $\{|\mathcal{R}| \mid \mathcal{R} \text{ is a set of disjoint } \omega\text{-rays}\}$. Note for each end ω there is in fact a set \mathcal{R} of vertex disjoint ω -rays with $|\mathcal{R}| = \deg(\omega)$ [30, Satz 1].

In digraphs all these terms are identically defined as for the underlying multigraph, although we will often add the adjective *undirected* to them to avoid confusion.

We will also consider directed versions for most of these terms, based on *directed paths*, which are either trivial paths or paths with a unique source and a unique sink. We call the source its *start vertex* and the sink its *endvertex*. If P consists only of a single vertex, we call that vertex the *endvertex* of P .

We call a ray R a *backwards directed ray* if it has a unique source while $d^-(w) = d^+(w) = 1$ holds for every other vertex $w \in V(R) \setminus \{v\}$. We call the source its *start vertex*. A *forwards directed ray* is analogously defined with a unique sink and by interchanging d^- and d^+ in the second condition. We call the source its *end vertex*.

We define the *ends of a digraph* G precisely as the ends of its underlying multigraph. The set of all ends of G is also denoted by $\Omega(G)$. We say that a directed ray R of D is *contained* in some end $\omega \in \Omega(D)$ if the underlying ray of R is contained in the end ω of the underlying multigraph of D .

Let $A, B \subseteq V(G)$ be two (not necessarily disjoint) vertex sets in a graph G . An *A-B separator* is a set S of vertices such that $A \setminus S$ and $B \setminus S$ lie in different components of $G - S$. We also say S *separates* A and B . As before, by a slight abuse of notation, if $A = \{a\}$ (or $B = \{b\}$) is a singleton we will replace A by a (or B by b respectively) for this notion.

For a graph or digraph G a *bipartition* (X, Y) of $V(G)$, i.e. a pair of vertex sets with $X \cup Y = V(G)$ and $X \cap Y = \emptyset$, we call the edge set $E(X, Y)$ (as defined above) a *cut* of G and refer to X and Y as the *sides* of the cut. Moreover, by writing $E(M, N)$ and calling it a cut of D we implicitly assume M and N to be the sides of that cut, and by calling an edge set B a cut we implicitly assume that B is of the form $E(M, N)$ for suitable sets M and N .

1.3.2. The compactness principle in combinatorics

A very useful tool in infinite combinatorics is the *compactness principle in combinatorics*. We give a short overview for different versions of this principle. All of these versions use some amount of the axiom of choice (usually in the form of Tychonoff's Theorem, which motivates the name for the principle). A discussion of some of these versions can be found in [9, Appendix A].

A weaker version of this principle is due to Kőnig, but it still is powerful enough for many applications.

Lemma (Kőnig's Infinity Lemma 1927). [9, Lemma 8.2.1] *Let $(X_n \mid n \in \mathbb{N})$ be a family of disjoint non-empty finite sets and let G be a graph with vertex set $\bigcup_{n \in \mathbb{N}} X_n$. Assume that for every $n \in \mathbb{N}$ that every vertex $v \in X_{n+1}$ has a neighbour in V_n . Then G contains a ray $v_0 v_1 v_2 \cdots$ with $v_n \in X_n$ for all $n \in \mathbb{N}$.*

One of the earliest versions of a stronger version of the compactness principle is due to Rado. We will not use this version but state it to give some historical context.

A *choice function* for a family $(X_i \mid i \in I)$ is a map $f : I \rightarrow \bigcup\{X_i \mid i \in I\}$ such that $f(i) \in X_i$ for all $i \in I$.

Lemma (Rado's Selection Principle, 1949). [39, 40] *Let $(X_i \mid i \in I)$ be a family of finite non-empty sets. Suppose that for every finite $J \subseteq I$ there is a choice function f_J for the family $(A_j \mid j \in J)$. Then there is a choice function f for the family $(X_i \mid i \in I)$ such that for every finite $J \subseteq I$ there is a finite $K \subseteq I$ with $J \subseteq K$ such that $f_K(j) = f(j)$ for all $j \in J$.*

A straightforward generalisation of Kőnig's Infinity Lemma is the following Generalised Infinity Lemma, which has its roots in category theory.

A partially ordered set (P, \leq) is *directed* if any two elements have a common upper bound, i.e. for any $p, q \in P$ there is an $r \in P$ with $p \leq r$ and $q \leq r$. A *directed inverse system* consists of

- a directed poset P ;
- a family of sets $(X_p \mid p \in P)$;
- for all $p, q \in P$ with $p < q$ a map $f_{q,p} : X_q \rightarrow X_p$;

such that the maps are *compatible*, meaning $f_{q,p} \circ f_{r,q} = f_{r,p}$ for all $p, q, r \in P$ with $p < q < r$.

The *inverse limit* of such a directed inverse system is the set

$$\varprojlim (X_p \mid p \in P) = \left\{ (x_p \mid p \in P) \in \prod_{p \in P} X_p \mid f_{q,p}(x_q) = x_p \right\}.$$

Lemma (Generalised Infinity Lemma). *The inverse limit of any directed inverse system of non-empty finite sets is non-empty.*

Lastly, we will state two more versions of the compactness principle which both are commonly used in infinite combinatorics.

Let X be any set and S a finite set. Let \mathcal{F} be a set of finite subsets of X . For every $Y \in \mathcal{F}$ let $\mathcal{A}(Y)$ be a set of functions from Y to S , which we call the *admissible functions of Y* . We call a set $\mathcal{Y} \subseteq \mathcal{F}$ *compatible* if there exists a function $f : X \rightarrow S$ all whose restrictions to the sets in \mathcal{Y} are admissible, i.e. which satisfies $f \upharpoonright Y \in \mathcal{A}(Y)$ for all $Y \in \mathcal{Y}$.

Lemma (Compactness Principle, version 1). *\mathcal{F} is compatible if every finite $Y \subseteq \mathcal{F}$ is compatible.*

Let $(X_i \mid i \in I)$ be a family of finite sets. A *constraint* is a pair (J, \mathcal{K}) where J is a finite subset of I and $\mathcal{K} \subseteq \prod_{i \in J} X_j$. An element $x \in \prod_{i \in I} X_i$ *satisfies* a constraint (J, \mathcal{K}) if $(x_i \mid i \in J) \in \mathcal{K}$. A set \mathcal{C} of constraints is *satisfiable* if there is an $x \in \prod_{i \in I} X_i$ satisfying every constraint in \mathcal{C} .

Lemma (Compactness Principle, version 2). *A set of constraints is satisfiable if and only if every finite subset is satisfiable.*

In many applications we will not spell out the precise translation of the problem to one of these versions.

Part I.

Undirected graphs

2. Representations of infinite tree sets

2.1. Introduction

Separations of graphs have been studied in the context of structural graph theory for a long time. For instance every edge of the decomposition tree of a tree-decomposition of a graph defines a separation in a natural way*. The separations obtained in this way have an additional important property: they are nested[†] with each other. Looking at nested sets of separations of a graph has since been a useful way to study tree-decompositions, and especially in infinite graphs they offer an analogue when a tree-decomposition with a certain desired property may not exist (see [42] for example).

While any tree-decomposition of a graph into small parts witnesses that the graph has low tree-width, there are various dense objects that force high tree-width in a graph. Among these are large cliques and clique minors, large grids and grid minors as well as high-order brambles. All these dense objects in a graph have the property that they orient its low-order separations by lying mostly on one side of any given low-order separation. For such a dense structure in a graph these orientations of separations are consistent with each other: no two of them ‘disagree’ about where the dense object lies by pointing away from each other.

In [41] Robertson and Seymour proposed the notion of tangles, which are such families of consistently oriented separations up to a certain order. These tangles can be studied in their own right, instead of any dense objects that may induce them. By varying the strength of the consistency conditions one can model different kinds of dense objects, and the resulting consistent orientations give rise to different

*As the sides of the separation, consider the union of the parts corresponding to the components of the tree after deleting the edge.

[†]Two separations are nested if a side of the first separation is a subset of a side of the second separation, and the other side of the second separation is a subset of the other side of the first separation.

types of tangles.

To talk about these separations systems one does not even need an underlying graph structure or ground set: they can be formulated in a purely axiomatic way, see Diestel [10]. Such a separation system is simply a partially ordered set with an order-reversing involution. The notions of consistency of separations that come from dense substructures in graphs can be translated into this setting as well. The tangles of graphs then become abstract tangles, and the tree-like structures become nested systems of separations, so-called *tree sets* [12]. This abstract framework turns out to be no less powerful, even for graphs alone, than ordinary graph separations. In [16] Diestel and Oum established an abstract duality theorem for separation systems which easily implies (see [17]) all the classical duality results from graph- and matroid theory, such as the tree-width duality theorem by Seymour and Thomas [44]. The unified duality theorem asserts that for any sensible notion of consistency a separation system contains either an abstract tangle or a tree set witnessing that no such tangle exists. Furthermore this abstract notion of separation systems can be applied in fields outside of graph theory, for instance in image analysis [18].

Tree sets are also interesting objects in their own right: they are flexible enough to model a whole range of other ‘tree-like’ structures in discrete mathematics, such as ordinary graph trees, order trees and nested systems of bipartitions of sets [12].

In fact, tree sets and graph-theoretic trees are related even more closely than that: for any tree T the set \vec{E} of oriented edges of T admits a natural partial order, which in fact turns \vec{E} into a tree set, the *edge tree set* of T . As was shown in [12], these edge tree sets of graph-theoretical trees are rich enough to represent all finite tree sets: every finite tree set is isomorphic to the edge tree set of a suitable tree.

In this chapter we extend the analysis of representations of tree sets to infinite tree sets. The definition of an edge tree set of a graph-theoretical tree straightforwardly extends to infinite trees. From the structure of these it is clear that the edge tree set of a tree T cannot contain a chain of order type $\omega + 1$. We will show that this is the only obstruction for a tree set to being representable by the edge tree set of a (possibly infinite) tree:

Theorem 2.1.1. *Every tree set without a chain of order type $\omega + 1$ is isomorphic to the edge tree set of a suitable tree.*

Secondly, we would like to represent infinite tree sets that do contain a chain of order type $\omega + 1$ by edge tree sets of an adequate tree structure as well. To achieve this we turn to the notion of *graph-like spaces* introduced by Thomassen and Vella [49] and further studied by Bowler, Carmesin and Christian [6]: these are topological spaces with a clearly defined structure of vertices and edges, which can be seen as a limit object of finite graphs. In particular, for a chain of any order type, there exists a graph-like space containing a ‘path’ whose edges form a chain of that order type. Therefore the *tree-like spaces*, those graph-like spaces which have a tree-like structure, overcome the obstacle of chains of order type $\omega + 1$ which prevented the edge tree sets of infinite trees from representing all infinite tree sets: unlike graph-theoretic trees, tree-like spaces can have limit edges. And indeed we will prove in this chapter that the edge tree sets of tree-like spaces can be used to represent all tree sets.

Theorem 2.1.2. *Every tree set is isomorphic to the edge tree set of a suitable tree-like space.*

This chapter is organised as follows. In Section 2.2 we recall the basic definitions of abstract separation systems and tree sets and establish a couple of elementary lemmas we will use throughout the chapter. Following that, in Section 2.3, we formally define the edge tree set of a tree and prove Theorem 1. In Section 2.4, we introduce the concept of *tree-like spaces* which generalise infinite graph-theoretical trees. We define edge tree sets of tree-like spaces analogously to edge tree sets of graph-theoretical trees and then prove Theorem 2. In order to do this we need a result linking the two concepts of connectivity in graph-like spaces: topological connectivity and ‘pseudo-arc connectivity’, the analogue of graph-theoretical connectivity for graph-like spaces. In Section 2.4 we make use of the fact that for compact graph-like spaces these two notion of connectivity are equivalent, and give a proof of this fact in Section 2.5.

2.2. Separation systems

An abstract *separation system* $\vec{S} = (\vec{S}, \leq, *)$ is a partially ordered set with an order-reversing involution $*$. An element $\vec{s} \in \vec{S}$ is called an *oriented separation*, and its *inverse* $(\vec{s})^*$ is denoted as \vec{s} , and vice versa. The pair $s = \{\vec{s}, \vec{s}\}$ is an

unoriented separation[‡], with orientations \vec{s} and $\vec{\bar{s}}$, and the set of all such pairs is denoted as S . The assumption that $*$ is order-reversing means that for all $\vec{s}, \vec{r} \in \vec{S}$ we have $\vec{s} \leq \vec{r}$ if and only if $\vec{\bar{s}} \geq \vec{\bar{r}}$. If S' is a set of unoriented separations, we write \vec{S}' for the set $\cup S'$ of all orientations of separations in S' .

A separation \vec{s} is *small* and its inverse $\vec{\bar{s}}$ *co-small* if $\vec{s} \leq \vec{\bar{s}}$. If neither \vec{s} nor $\vec{\bar{s}}$ is small then s is *regular*, and we call both \vec{s} and $\vec{\bar{s}}$ regular as well.

A separation $\vec{s} \in \vec{S}$ is *trivial in \vec{S}* and its inverse $\vec{\bar{s}}$ is *co-trivial in \vec{S}* if there is some $\vec{r} \in \vec{S}$ with $\vec{s} \leq \vec{r}, \vec{\bar{s}}$ and $s \neq r$. In this case r is the *witness* of the triviality of \vec{s} . If neither \vec{s} nor $\vec{\bar{s}}$ is trivial in \vec{S} we call s *nontrivial*. If \vec{s} is a trivial separation with witness r then \vec{s} is small as $\vec{s} \leq \vec{r} \leq \vec{\bar{s}}$. Conversely every separation that lies below a small separation is trivial: if \vec{s} is small and $r \neq s$ has an orientation $\vec{r} \leq \vec{s}$, then \vec{r} is trivial as $\vec{r} < \vec{s} \leq \vec{\bar{s}}$.

Two unoriented separations s and r are *nested* if they have comparable orientations. Otherwise r and s *cross*. A set S' of separations is nested if all of its elements are pairwise nested.

A *tree set* is a nested separation system with no trivial elements. It is *regular* if all of its elements are regular, i.e. if no $\vec{s} \in \tau$ is small.

An *orientation* of a set \vec{S}' or S' of separations is a set $O \subseteq \vec{S}'$ with $|O \cap s| = 1$ for every $s \in S'$. An orientation is *consistent* if $\vec{\bar{s}} \leq \vec{r}$ implies $r = s$ for all $\vec{r}, \vec{s} \in O$. A *partial orientation* of \vec{S} is an orientation of a subset of \vec{S} . A partial orientation P *extends* another partial orientation Q if $Q \subseteq P$.

For a tree set τ an orientation O of τ is *splitting* if it is consistent and has the property that for every $\vec{r} \in O$ there is some maximal element \vec{s} of O with $\vec{r} \leq \vec{s}$.

Consistent orientations of a tree set τ can be thought of as the ‘vertices’ of a tree set, an idea that we will make more precise in the next sections. In the context of infinite tree sets, the non-splitting orientations can be thought of as ‘limit vertices’ or ‘ends’ of the tree set.

A subset $\sigma \subseteq \tau$ is a *star* if $\vec{r} \leq \vec{\bar{s}}$ for all $\vec{r}, \vec{s} \in \sigma$ with $\vec{r} \neq \vec{s}$. For example, the set of maximal elements of a consistent orientation of a tree set is always a star:

Lemma 2.2.1. *Let O be a consistent orientation of a tree set τ . Then the set σ of the maximal elements of O is a star.*

[‡]To improve readability ‘oriented’ and ‘unoriented’ will often be omitted if the type of separation follows from the context.

Proof. Let $\vec{r}, \vec{s} \in \sigma$ with $\vec{r} \neq \vec{s}$ be given. Then neither $\vec{r} \leq \vec{s}$ nor $\vec{r} \geq \vec{s}$ as both are maximal elements of O . The consistency of O implies that $\vec{r} \not\leq \vec{s}$, so $\vec{r} \leq \vec{s}$ is the only possible relation and hence σ is a star. \square

A star $\sigma \subseteq \tau$ *splits* τ , or is a *splitting star* of τ , if it is the set of maximal elements of a splitting orientation of τ . Note that every element of a finite tree set lies in a splitting star, but infinite tree sets can have elements that lie in no splitting star; see Example 2.2.3 and Lemma 2.2.4 below.

More generally, given a partial orientation P of τ , is it possible to extend it to a consistent orientation of τ ? Of course P needs to be consistent itself for this to be possible. The next Lemma shows that under this necessary assumption it is always possible to extend a partial orientation to all of τ . In particular, every element of a tree set induces a consistent orientation in which it is a maximal element. This orientation is in fact unique:

Lemma 2.2.2 (Extension Lemma). [10] *Let S be a set of separations, and let P be a consistent partial orientation of S .*

- (i) *P extends to a consistent orientation O of S if and only if no element of P is co-trivial in S .*
- (ii) *If \vec{p} is maximal in P , then O in (i) can be chosen with \vec{p} maximal in O if and only if \vec{p} is nontrivial in \vec{S} .*
- (iii) *If S is nested, then the orientation O in (ii) is unique.*

The last part of the Extension Lemma implies that every element \vec{s} of a tree set τ is maximal in exactly one consistent orientation O of τ . Hence \vec{s} lies in a splitting star if and only if this O is splitting.

In an infinite tree set there might be elements that do not lie in a splitting star:

Example 2.2.3. Let τ be the tree set with ground set

$$\{\vec{s}_n \mid n \in \mathbb{N}\} \cup \{\vec{s}_n \mid n \in \mathbb{N}\} \cup \{\vec{t}, \vec{t}\},$$

where $\vec{s}_i \leq \vec{s}_j$ and $\vec{s}_i \geq \vec{s}_j$ whenever $i \leq j$, as well as $\vec{s}_n \leq \vec{t}$ and $\vec{s}_n \geq \vec{t}$ for all $n \in \mathbb{N}$. The separation \vec{t} is maximal in the orientation

$$O = \{\vec{s}_n \mid n \in \mathbb{N}\} \cup \{\vec{t}\},$$

which is not splitting as no \vec{s}_n lies below a maximal element of O . Hence \vec{t} does not lie in a splitting star of τ .

In the above example the chain $C = \{\vec{s}_n \mid n \in \mathbb{N}\} \cup \{\vec{t}\}$ has order-type $\omega + 1$. But these $\omega + 1$ chains turn out to be the only obstruction for separations not being elements of splitting stars, as the following lemma shows. Let us call a tree set that does not contain a chain of order type $\omega + 1$ *tame*.

Lemma 2.2.4. *Every element of a tame tree set τ lies in some splitting star of τ .*

Proof. For every $\vec{t} \in \tau$ we can apply the Extension Lemma 2.2.2 to $P := \{\vec{t}\}$ to find that there is a unique consistent orientation O of τ in which \vec{t} is a maximal element. Thus \vec{t} lies in a splitting star if and only if this orientation O is splitting. Let us show that for every $\vec{t} \in \tau$ this orientation O splits τ unless O contains a chain of order type ω for which \vec{t} is an upper bound; this directly implies the claim since every such chain in O together with \vec{t} is a chain of order type $\omega + 1$ in τ .

So let $\vec{t} \in \tau$ be given and consider the unique consistent orientation O of τ in which \vec{t} is maximal. Suppose that O does not split τ , i.e. that there is some $\vec{s} \in O$ which does not lie below any maximal element of O . Consider the set $C \subseteq O$ of all elements \vec{r} of O with $\vec{r} \geq \vec{s}$. Since \vec{s} and hence no element of C can lie below \vec{t} we must have $\vec{r} \leq \vec{t}$ for all $\vec{r} \in C$ since τ is nested. Thus \vec{t} is an upper bound for C . Now if C has a maximal element then this separation is also a maximal element of O , contrary to our assumption about \vec{s} ; therefore C cannot have a maximal element and hence contains a chain of order type ω , as claimed. \square

A direct consequence of Lemma 2.2.4 is that every element of a finite tree set lies in a splitting star.

Given two separation systems R and S , a map $f: R \rightarrow S$ is a *homomorphism* of separation systems if it *commutes with the involution*, i.e. $(f(\vec{r}))^* = f(\vec{r}^*)$ for all $\vec{r} \in R$, and is *order-preserving*, i.e. $f(\vec{r}_1) \leq f(\vec{r}_2)$ whenever $\vec{r}_1 \leq \vec{r}_2$ for all $\vec{r}_1, \vec{r}_2 \in R$. Please note that the condition for f to be order-preserving is not ‘if and only if’: it is allowed that $f(\vec{r}_1) \leq f(\vec{r}_2)$ for incomparable $\vec{r}_1, \vec{r}_2 \in R$. Furthermore f need not be injective.

As all trivial separations are small every regular nested separation system is a tree set. These two properties, regular and nested, are preserved by homomorphisms of

separations systems, albeit in different directions: the image of nested separations is nested, and the preimage of regular separations is regular.

Lemma 2.2.5. *Let $f: R \rightarrow S$ be a homomorphism of separation systems. If S is regular then so is R ; and if R is nested then so is its image in S .*

Proof. First suppose that some $\vec{r} \in R$ is small, that is, that $\vec{r} \leq \tilde{r}$. Then

$$f(\vec{r}) \leq f(\tilde{r}) = (f(\vec{r}))^*,$$

so S contains a small element. Therefore if S is regular then R must be regular as well.

Now suppose that R is nested consider two unoriented separations $s, s' \in S$ and for which there are $r, r' \in R$ with $s = f(r)$ and $s' = f(r')$. Since R is nested r and r' have comparable orientations, say $\vec{r} \leq \vec{r}'$. Then $\vec{s} := f(\vec{r}) \leq f(\vec{r}') =: \vec{s}'$, showing that s and s' are nested. Hence if R is nested its image in S is nested too. \square

A bijection $f: R \rightarrow S$ is an *isomorphism* of separation systems if both f and its inverse map are homomorphisms of separation systems. Two separation systems R and S are *isomorphic*, denoted as $R \cong S$, if there is an isomorphism $f: R \rightarrow S$ of separation systems. If one of R and S (and thus both) is a tree set we call f an *isomorphism of tree sets*.

Lemma 2.2.5 makes it possible to show that a homomorphism $f: R \rightarrow S$ of separation systems is an isomorphism of tree sets without knowing beforehand that either R or S is a tree set:

Lemma 2.2.6. *Let $f: R \rightarrow S$ be a bijective homomorphism of separation systems. If R is nested and S regular then f is an isomorphism of tree sets.*

Proof. From Lemma 2.2.5 it follows that both R and S are regular and nested, which means they are regular tree sets. Therefore all we need to show is that the inverse of f is order-preserving, i.e. that $\vec{r}_1 \leq \vec{r}_2$ whenever $f(\vec{r}_1) \leq f(\vec{r}_2)$. Let $\vec{r}_1, \vec{r}_2 \in R$ with $f(\vec{r}_1) \leq f(\vec{r}_2)$ be given. As R is nested, r_1 and r_2 have comparable orientations. If $\vec{r}_1 \geq \vec{r}_2$, then $f(\vec{r}_1) = f(\vec{r}_2)$, implying $\vec{r}_1 = \vec{r}_2$ and hence the claim. If $\vec{r}_1 \leq \tilde{r}_2$, then $f(\vec{r}_1) \leq f(\vec{r}_2), f(\tilde{r}_2)$, contradicting the fact that S is a regular tree set. Finally, if $\vec{r}_1 \geq \tilde{r}_2$, then $f(\tilde{r}_2) \leq f(\vec{r}_2)$, contradicting the fact that S is regular. Hence $\vec{r}_1 \leq \vec{r}_2$, as desired. \square

2.3. Regular tame tree sets and graph-theoretical trees

Every graph-theoretical tree T naturally gives rise to a tree set, its *edge tree set* $\tau(T)$ of T (see below for a formal definition). However, while every tree gives rise to a tree set, not every tree set ‘comes from’ a tree. In this section we characterise those infinite tree sets that arise from graph-theoretical trees as the tree sets which are both regular and tame, i.e. contain no chain of order-type $\omega + 1$. More precisely, given a regular tame tree set τ we will define a corresponding tree $T(\tau)$. These definitions in turn should be able to capture the essence of what it means to be ‘tree-like’. More precisely we want the following properties:

- the tree constructed from the edge tree set of T is isomorphic to T ;
- the edge tree set of the tree constructed from τ is isomorphic to τ .

2.3.1. The edge tree set of a tree

Let $T = (V, E)$ be a graph-theoretical tree, finite or infinite. Let $\vec{E}(T)$ be the set of oriented edges of T , that is

$$\vec{E}(T) = \{(x, y) \mid \{x, y\} \in E(T)\}.$$

We define an involution $*$ by setting $(x, y)^* := (y, x)$ for all edges $xy \in E(T)$, and a partial order \leq on $\vec{E}(T)$ by setting $(x, y) < (v, w)$ for edges $xy, vw \in E(T)$ if and only if $\{x, y\} \neq \{v, w\}$ and the unique $\{x, y\}$ – $\{v, w\}$ -path in T joins y to v . Then the *edge tree set* $\tau(T)$ is the separation system $(\vec{E}(T), \leq, *)$. It is straightforward to check that $\tau(T)$ is indeed a regular tree set.

Note that every maximal chain in $\tau(T)$ corresponds to the edge set of a path, ray or double ray in T . Hence $\tau(T)$ does not contain any chain of length $\omega + 1$ and hence is tame.

If T is the decomposition tree of a tree-decomposition of a graph G , then the tree set $\tau(T)$ is isomorphic to the tree set formed by the separations of G that correspond[§] to the edges of T (with some pathological exceptions). This

[§]An edge e of the decomposition tree T of a tree-decomposition naturally defines a graph separation by considering the union of the parts in the respective components of $T - e$ as the sides of that separation.

relationship between tree-decompositions and tree sets was further explored in [12].

2.3.2. The tree of a regular tame tree set

Let τ be a regular tame tree set. Our aim is to construct a corresponding graph-theoretical tree $T(\tau)$. Recall that a consistent orientation O of τ is called splitting if every element of O lies below some maximal element of O . By the uniqueness part of the Extension Lemma 2.2.2, every splitting star extends to exactly one splitting orientation. Write \overline{O} for the set of all splitting orientations of τ . We will use \overline{O} as the vertex set of $T(\tau)$. Moreover note that it will turn out that the non-splitting orientations will precisely correspond to the ends of $T(\tau)$.

Let us show first that, for any two splitting stars, each of them contains exactly one element that is inconsistent with the other star. We will later use this little fact when we define the edges of our tree.

Lemma 2.3.1. *Let σ_1, σ_2 be two distinct splitting stars of τ and $O_2 \in \overline{O}$ the orientation inducing σ_2 . Then there is exactly one $\vec{s} \in \sigma_1$ with $\vec{s} \in O_2$.*

Proof. There is at least one such \vec{s} as O_2 does not induce σ_1 . For any two $\vec{r}, \vec{s} \in \sigma$ the set $\{\vec{r}, \vec{s}\}$ is inconsistent, so there is at most one $\vec{s} \in \sigma_1$ with $\vec{s} \in O_2$. \square

Note that this lemma holds for every tree set as the proof did not use any assumptions on τ .

Our assumption that τ is tame implies the following sufficient condition for a consistent orientation to be splitting:

Lemma 2.3.2. *Let O be a consistent orientation of τ with at least one maximal element. Then O splits τ .*

Proof. Let \vec{t} be a maximal element of O . By Lemma 2.2.4 \vec{t} lies in a splitting star of τ , i.e. is a maximal element of a consistent orientation that splits τ . By the Extension Lemma 2.2.2, O is the only consistent orientation of τ of which \vec{t} is a maximal element; hence O must be splitting. \square

Together with the Extension Lemma 2.2.2 this immediately implies the following:

Corollary 2.3.3. *Every $\vec{s} \in \tau$ lies in exactly one splitting star of τ . Equivalently every $\vec{s} \in \tau$ is maximal in exactly one consistent orientation O and $O \in \overline{O}$.*

Proof. For $\vec{s} \in \tau$ apply the Extension Lemma 2.2.2 to $\{\vec{s}\}$ to obtain a unique consistent orientation O of τ in which \vec{s} is a maximal element. It then follows from Lemma 2.3.2 that O is splitting. \square

For $\vec{s} \in \tau$ write $O(\vec{s})$ for the unique consistent orientation of τ in which \vec{s} is maximal. Then Lemma 2.3.1 together with Corollary 2.3.3 says that for distinct $O, O' \in \overline{\mathcal{O}}$ there is at most one $\vec{s} \in O'$ with $O(\vec{s}) = O$.

Now we define the graph $T(\tau)$. Let $V(T(\tau)) = \overline{\mathcal{O}}$ and

$$E(T(\tau)) = \left\{ \{O(\vec{s}), O(\vec{s}')\} \mid \vec{s} \in \tau \right\}.$$

We call $T(\tau)$ the *tree corresponding to τ* , where τ is a regular tame tree set.

First note that $T(\tau)$ does not contain any loops and hence is indeed a simple graph since $O(\vec{s})$ and $O(\vec{s}')$ are different for any $\vec{s} \in \tau$.

We need to check that $T(\tau)$ is a tree.

Lemma 2.3.4. *$T(\tau)$ does not contain any cycles.*

Proof. For $O \in \overline{\mathcal{O}}$ the set of incoming edges is precisely the splitting star induced by O . If $\vec{s}_1, \dots, \vec{s}_k$ are the edges of an oriented cycle in \vec{T} , then each of these and the inverse of its cyclic successor lie in a common splitting star. Hence $\vec{s}_1 \leq \vec{s}_2 \leq \dots \leq \vec{s}_k \leq \vec{s}_1$ by the star property, a contradiction. \square

To prove that $T(\tau)$ is connected, our strategy is as follows. To find a path from $O \in \overline{\mathcal{O}}$ to $O' \in \overline{\mathcal{O}}$ we use Lemma 2.3.1 to find $\vec{s} \in O$ which is maximal in O with $\vec{s}' \in O'$. Then we consider $O^* := (O \cup \{\vec{s}'\}) \setminus \{\vec{s}\}$. This orientation is again in $\overline{\mathcal{O}}$ and a neighbour of O in $T(\tau)$. If $O^* = O'$ we are done; otherwise we can iterate the process with O^* and O' . Either this process terminates after finitely many steps, in which case we found a path from O to O' , or it continues indefinitely. In the latter case the infinitely many separations we inverted form a chain with an upper bound in O' , which would yield a chain of order type $\omega + 1$.

The next short Lemma forms the basis of this iterative flipping process.

Lemma 2.3.5. *Let $\vec{s}_1, \dots, \vec{s}_n, \vec{s}' \in \tau$ be distinct separations with $O(\vec{s}_{k+1}) = O(\vec{s}_k)$ for all $k \in \mathbb{N}$ with $1 \leq k < n$ and $\vec{s}_n < \vec{s}'$. Then there is a separation $\vec{s}_{n+1} \in \tau$ with $O(\vec{s}_{n+1}) = O(\vec{s}_n)$ and $\vec{s}_{n+1} \leq \vec{s}'$.*

Proof. Let \vec{s}_{n+1} be the unique separation in $O(\vec{s}')$ with $O(\vec{s}_{n+1}) = O(\vec{s}_n)$. Then $\vec{s}_n \leq \vec{s}_{n+1}$ by the star property. Hence if $\vec{s}_{n+1} \leq \vec{s}'$, then \vec{s}_n would be trivial, therefore $\vec{s}_{n+1} \leq \vec{s}'$ as desired. \square

For $\vec{s}_1, \dots, \vec{s}_n, \vec{s}'$ and \vec{s}_{n+1} as in Lemma 2.3.5 there is an edge between $O(\vec{s}_k)$ and $O(\vec{s}_{k+1})$ for every $1 \leq k \leq n$. Additionally if $\vec{s}_{n+1} \neq \vec{s}'$ then $\vec{s}_1, \dots, \vec{s}_{n+1}, \vec{s}'$ again fulfill the assumptions of the lemma, so it can be used iteratively.

Furthermore note that $\vec{s}_1 \leq \vec{s}_2 \leq \dots \leq \vec{s}_n \leq \vec{s}_{n+1}$, so if this iteration does not terminate the \vec{s}_k form an infinite chain. From this we now prove that $T(\tau)$ is connected.

Lemma 2.3.6. *$T(\tau)$ is connected.*

Proof. Let $O, O' \in \overline{\mathcal{O}}$ be distinct orientations. Let \vec{s}_1 be the unique separation in O' with $O = O(\vec{s}_1)$, and \vec{s}' the unique separation in O with $O' = O(\vec{s}')$. Then $\vec{s}_1 \leq \vec{s}'$, and if $\vec{s}_1 = \vec{s}'$ then O and O' are joined by an edge in $T(\tau)$. Otherwise the assumptions of Lemma 2.3.5 are met for $n = 1$. Applying Lemma 2.3.5 iteratively either yields $\vec{s}_{n+1} = \vec{s}'$ for some $n \in \mathbb{N}$, in which case we found a path in $T(\tau)$ joining O and O' , or we obtain a strictly increasing sequence $(\vec{s}_n)_{n \in \mathbb{N}}$ with $\vec{s}_n \leq \vec{s}'$ for all $n \in \mathbb{N}$, that is, a chain of order type $\omega + 1$. \square

2.3.3. Regular tame tree sets and trees – A characterisation

Finally we will prove that the given constructions of the previous subsections agree with each other.

Lemma 2.3.7. *Any regular tame tree set τ' is isomorphic to $\tau(T(\tau'))$.*

Proof. Let $\varphi: \tau' \rightarrow \tau(T(\tau'))$ be the map defined by $\varphi(\vec{s}) = (O(\vec{s}), O(\vec{s}'))$. This is a bijection by Corollary 2.3.3. Note that for $\vec{s} \in \tau'$ the orientations $O(\vec{s})$ and $O(\vec{s}')$ differ only in s by consistency and are thus adjacent in T .

As τ' and $\tau(T(\tau'))$ are regular tree sets all we need to show is that φ is a homomorphism of separation systems. Then φ will be an isomorphism of tree sets by Lemma 2.2.6.

It is clear from the definition that φ commutes with the involution. Therefore it suffices to show that φ is order-preserving.

Let $\vec{s}, \vec{s}' \in \tau'$ be two separations with $\vec{s} < \vec{s}'$. We need to show that the unique $\{O(\vec{s}), O(\vec{s}')\}$ - $\{O(\vec{s}'), O(\vec{s})\}$ -path in $T(\tau)$ joins $O(\vec{s})$ and $O(\vec{s}')$. Redoing the proof of Lemma 2.3.6 with $O = O(\vec{s})$ and $O' = O(\vec{s}')$ constructs a $O(\vec{s})$ - $O(\vec{s}')$ -path every one of whose nodes contains \vec{s} and \vec{s}' by consistency. Hence $\varphi(\vec{s}) < \varphi(\vec{s}')$ as desired. \square

Lemma 2.3.8. *Any graph-theoretic tree T' is isomorphic to $T(\tau(T'))$.*

Proof. If $|V(T')| = 1$, then $\tau(T')$ is empty and hence $|V(T(\tau(T')))| = 1$.

Otherwise, for each node $v \in V(T')$ there is at some oriented edge $(w, v) \in \vec{E}(T')$ pointing towards that node. Let $\varphi: T' \rightarrow T(\tau(T'))$ be defined by $\varphi(v) := O((w, v))$. This map is well-defined since the edges directed towards a node $v \in V(T')$ form a splitting star with the same maximal elements yielding the unique consistent orientation containing all these oriented edges (cf. Corollary 2.3.3).

Similarly, given some $O = O((w, v)) \in V(T(\tau(T')))$, we obtain $\varphi(v) = O$ and hence that φ is surjective. By construction there is an edge between $O((v, w))$ and $O((w, v))$ for any edge $vw \in E(T)$ and similarly no edge between $O((v, w))$ and O if (w, v) is not maximal in O . \square

Hence we have proven our main theorem of this section:

Theorem 2.3.9. *1. A tree set is isomorphic to the edge tree set of a tree if and only if it is regular and tame.*

2. Any regular and tame tree set τ' is isomorphic to $\tau(T(\tau'))$.

3. Any graph-theoretic tree T' is isomorphic to $T(\tau(T'))$. \square

Additionally, for distinct but comparable tree sets, we can say precisely in which way the corresponding trees from Theorem 2.3.9 above are comparable: one will be a minor of the other.

Theorem 2.3.10. *Let T_1, T_2 be trees and τ_1, τ_2 be regular tame tree-sets.*

1. If $\tau_1 \subseteq \tau_2$, then $T(\tau_1)$ is a minor of $T(\tau_2)$.

2. If T_1 is a minor of T_2 , then $\tau(T_1)$ is isomorphic to a subset of $\tau(T_2)$.

Theorem 2.3.10 is a special case of Theorems 2.4.14 and 2.4.15 from the next section and hence we will omit its proof here.

2.4. Regular tree sets and tree-like spaces

2.4.1. Graph-like spaces

As we have seen in Section 2.3, not every tree set, even regular, can be represented as the edge tree set of a tree. In this section we find a (topological) relaxation of the notion of a (graph-theoretical) tree, to be called *tree-like spaces*. Like trees, these tree-like spaces give rise to a regular edge tree set in a natural way, but which are just general enough that, conversely, every regular tree set can be represented as the edge tree set of a tree-like space.

The concept of graph-like spaces was first introduced in [49] by Thomassen and Vella, and further studied in [6] by Bowler, Carmesin and Christian. In [6] the authors discuss the connections between graph-like spaces and graphic matroids, which are of no interest to us here. Instead we determine when a graph-like space is tree-like, and then show that every regular tree set can be represented as the edge tree set of a tree-like space.

Graph-like spaces are limit objects of graphs that are not themselves graphs. In short they consist of the usual vertices and edges, together with a topology that allows the vertices and edges to be limits of each other. The formal definition is as follows.

Definition 2.4.1. [6] A *graph-like space* G is a topological space (also denoted by G) together with a *vertex set* $V(G)$, an *edge set* $E(G)$ and for each $e \in E(G)$ a continuous map $\iota_e^G : [0, 1] \rightarrow G$ (the superscript may be omitted if G is clear from the context) such that:

- The underlying set of G is $V(G) \dot{\cup} [(0, 1) \times E(G)]$.
- For any $x \in (0, 1)$ and $e \in E(G)$ we have $\iota_e(x) = (x, e)$.
- $\iota_e(0)$ and $\iota_e(1)$ are vertices (called the *end-vertices* of e).
- $\iota_e \upharpoonright_{(0,1)}$ is an open map.
- For any two distinct $v, v' \in V(G)$, there are disjoint open subsets U, U' of G partitioning $V(G)$ and with $v \in U$ and $v' \in U'$.

The *inner points* of the edge e are the elements of $(0, 1) \times \{e\}$.

Note that G is always Hausdorff. For an edge $e \in E(G)$ the definition of graph-like space allows $\iota_e(0) = \iota_e(1)$. We call such an edge a *loop*. In our discussions of graph-like spaces loops are irrelevant, so the reader may imagine all graph-like spaces to be loop-free.

If U and U' are disjoint open subsets of G partitioning $V(G)$ we call the set of edges with end-vertices in both U and U' a *topological cut* of G and say that the pair (U, U') *induces* that cut. The last property of graph-like spaces then says that any two vertices can be separated by a topological cut.

A graph-like space G' is a *sub-graph-like space* of a graph-like space G if $V(G') \subseteq V(G)$, $E(G') \subseteq E(G)$ and G' is a subspace of G (as topological spaces). By slight abuse of notation we will write $G' \subseteq G$ to say that G' is a sub-graph-like space of G .

Let G be a graph-like space and $F \subseteq E(G)$ a set of edges of G . We write $G - F$ for the sub-graph-like space $G \setminus \{(x, e) \mid x \in (0, 1), e \in F\}$ with the same vertex set as G , with edge set $E(G) \setminus F$ and $\iota_e^{G-F} = \iota_e^G$ for all $e \in E(G) \setminus F$. We abbreviate $G - \{e\}$ as $G - e$. Given a set $W \subseteq V(G)$ of non-end-vertices we write $G - W$ for the sub-graph-like space $G \setminus W$ with $V(G - W) := V(G) \setminus W$, $E(G - W) := E(G)$ and $\iota_e^{G-F} = \iota_e^G$ for all $e \in E(G)$.

For reasons of cardinality arc-connectedness is not a very useful notion in graph-like spaces. Instead we work with an adapted concept of arcs. A graph-like space P is a *pseudo-arc* if P is a compact connected graph-like space with a *start-vertex* a and an *end-vertex* b satisfying the following:

- for each $e \in E(P)$ the vertices a and b are separated in $P - e$;
- for any two $x, y \in V(P)$ there is an edge $e \in E$ such that x and y are separated in $P - e$.

If P contains an edge then $a \neq b$; otherwise we call P *trivial*. A graph-like space G is *pseudo-arc-connected* if for all vertices $a, b \in V(G)$ there is a pseudo-arc $P \subseteq G$ with *start-vertex* a and *end-vertex* b .

The adapted notion of circles is analogous. A graph-like space is a *pseudo-circle* if it is a compact connected graph-like space with at least one edge satisfying the following:

- removing any edge from C does not disconnect C but removing any pair does;

- any two vertices of C can be separated in C by removing a pair of edges.

Pseudo-arcs and pseudo-circles are related as follows:

Lemma 2.4.2. [6] *Let G be a graph-like space, C a pseudo-circle in G and $e \in E(C)$. Then $C - e$ is a pseudo-arc in G joining the end-vertices of e .*

Conversely, let P and Q be nontrivial non-loop pseudo-arcs in G that meet precisely in their end-vertices. Then $P \cup Q$ is a pseudo-circle in G .

Given two graph-like spaces G_1, G_2 , a map $\varphi : G_1 \rightarrow G_2$ is an *isomorphism of graph-like spaces* if it is a homeomorphism (for the topological spaces) and it induces a bijection between $V(G_1)$ and $V(G_2)$.

Let G be a graph-like space and $F \subseteq E(G)$ a set of edges of G . We define a relation \sim'_F on G via

$$\iota_e(x) \sim'_F \iota_e(y) \text{ for all } e \in F \text{ and } x, y \in [0, 1].$$

Let \sim_F denote the minimal equivalence relation that extends the transitive and reflexive closure of \sim'_F such that the resulting quotient space $G/F := G/\sim_F$ is Hausdorff.

Remark 2.4.3. The *contraction* G/F of F in G is a graph-like space with vertex set $V(G/F) := \{[v] \in G/\sim_F \mid v \in V(G)\}$, edge set[¶] $E(G/F) := E(G) \setminus F$ and for each edge $e \in E(G) \setminus F$ the map $\iota_e^{G/F} := \iota_e^G$.

One can also easily show that each equivalence class with respect to \sim_F is connected in G . Moreover, we write $G.F$ for $G/(E(G) \setminus F)$ for the *contraction* to F in G .

We say that a graph-like space G' is a *minor* of graph-like space G if there are disjoint edge sets $F_1, F_2 \subseteq E(G)$ and a set $W \subseteq V(G/F_1) - F_2$ of non-end-vertices such that G' is isomorphic to $((G/F_1) - F_2) - W$.

We will also need the following fact about graph-like spaces:

Theorem 2.4.4. *A compact graph-like space is connected if and only if it is pseudo-arc connected.*

As the proof of Theorem 2.4.4 is relatively long and does not involve any tree-like spaces or other tree structures, we shall use Theorem 2.4.4 in this section without proof. Section 2.5 will then be devoted entirely to proving Theorem 2.4.4.

[¶]This is a slight abuse of notation since technically the inner points of an edge e in the quotient space are of the form $\{(x, e)\}$ and not (x, e) .

2.4.2. Tree-like spaces

There are many different equivalent ways of defining the graph-theoretical trees, which is an easy exercise to prove.

Proposition 2.4.5. *For a graph $T = (V, E)$ the following are equivalent.*

- (i) *For any two vertices $a, b \in V(T)$ there is a unique path in T from a to b ;*
- (ii) *T is connected but $T - e$ is not for any edge $e \in E(T)$;*
- (iii) *T is connected and contains no cycle.*
- (iv) *T contains no cycle but every graph T' with $V(T') = V(T)$ and $T' - F = T$ for some non-empty $F \subseteq E(T') \setminus E(T)$ does.*

A graph T is a tree if it has one (and thus all) of the above properties. In some situations one of these properties is easier to work with than the others, and their equivalence is used implicitly in many places in graph theory.

The above properties can be translated into the setting of graph-like spaces to say when a graph-like space is tree-like as follows:

Definition 2.4.6. A compact loop-free graph-like space G is a *tree-like space* if one of the following conditions holds:

- (i) For any two vertices $a, b \in V(G)$ there is a unique pseudo-arc in T from a to b ;
- (ii) G is connected but $G - e$ is not for any edge $e \in E(G)$;
- (iii) G is connected and contains no pseudo-circle;
- (iv) G contains no pseudo-circle but every graph-like space G' with $V(G') = V(G)$ and $G' - F = G$ for some non-empty $F \subseteq E(G') \setminus E(G)$ does.

Analogous to Proposition 2.4.5, we prove the following proposition.

Proposition 2.4.7. *For compact loop-free graph-like spaces the conditions in Definition 2.4.6 are equivalent.*

The argument is very similar to the proof of Proposition 2.4.5, but one additional technical lemma is needed: if two vertices a and b of a graph G are joined by two different paths it is obvious that some edge $e \in E(G)$ lies on exactly one of the two paths. However for graph-like spaces and pseudo-arcs this intuitive fact requires a surprising amount of set-up to prove (see [6]).

We forego this technical set-up and simply use the following lemma:

Lemma 2.4.8. [6, Remark 4.4] *Any nontrivial pseudo-arc in a graph-like space is the closure of the inner points of its edges.*

Lemma 2.4.8 immediately implies that if two vertices a and b of a graph-like space G are joined by two distinct pseudo-arcs P and Q then there is an edge $e \in E(G)$ which lies on exactly one of the two pseudo-arcs. In fact slightly more is true: both P and Q contain an edge that does not lie on the other pseudo-arc. For if the edge set of Q was a proper subset of the edge set of P then Q would be disconnected as the removal of any edge from P separates a and b in P .

Proof of Proposition 2.4.7 (based on Theorem 2.4.4).

(i) \Rightarrow (iv): Let G be a compact loop-free graph-like space with property (i). Suppose C is a pseudo-circle in G ; then for any $e \in E(C)$ both e and $C - e$ define pseudo-arcs in G joining the end-vertices of e , contradicting (i). Now let G' be a graph-like space with $V(G') = V(G)$ and $G' - F = G$ for some non-empty $F \subseteq E(G') \setminus E(G)$. Let $e \in F$ be an edge with end-vertices a and b . Then e defines a pseudo-arc P between a and b in G' . Let Q be the unique pseudo-arc in G joining a and b . Then P and Q intersect only in a and b , and hence their union is a pseudo-circle in G' by Lemma 2.4.2.

(iv) \Rightarrow (iii): Let G be a compact loop-free graph-like space with property (iv). Suppose G is not connected. Then G is not pseudo-arc connected by Theorem 2.4.4. Let a and b be a pair of vertices that are not connected by any pseudo-arc in G . In particular there is no edge between a and b . Let G' be a graph-like space with $V(G') = V(G)$ such that $G = G' - \{e\}$, where e is an edge in G' joining a and b . Then G' contains a pseudo-circle C , which has to contain e as otherwise C would be a pseudo-circle in G . But then by Lemma 2.4.2 $C - e \subseteq G$ is a pseudo-arc between the end-vertices of e , showing that a and b are joined by a pseudo-arc in G .

(iii) \Rightarrow (ii): Let G be a compact loop-free graph-like space with property (iii). Suppose $G - e$ is still connected for some $e \in E(G)$ with end-vertices a and b . Then $G - e$ contains a pseudo-arc P between a and b by Theorem 2.4.4, which together with e forms a pseudo-circle by Lemma 2.4.2.

(ii) \Rightarrow (i): Let G be a compact loop-free graph-like space with property (ii). Theorem 2.4.4 implies that G is pseudo-arc connected. For the uniqueness suppose G contains two different pseudo-arcs P and Q between two vertices a and b . Lemma 2.4.8 implies that there is an edge $e \in E(G)$ which lies on exactly one of the two pseudo-arcs. But then $G - e$ is still pseudo-arc connected^{||} and therefore connected, a contradiction. \square

Similarly to graph-theoretical trees every tree-like space gives rise to a regular tree set, see Subsection 2.4.3. We will show that the tree-like spaces are rich enough that one can obtain every regular tree set from them. This is in contrast to Section 2.3 where we showed that the regular tree sets coming from trees are precisely those with no chain of order type $\omega + 1$. This restriction was owed to the fact that graph-theoretical trees cannot have edges that are the limit of other edges. But tree-like spaces *can* have limit edges, so this is no longer a restriction.

In Subsection 2.4.4 we construct a corresponding regular tree set for a given tree-like space, and in Subsection 2.4.5 we will prove the characterisation analogously to the one in Section 2.3 by showing:

- the tree-like space constructed from the edge tree set of a tree like space T is isomorphic to T ;
- the edge tree set of the tree-like space constructed from a regular tree set τ is isomorphic to τ .

2.4.3. The edge tree set of a tree-like space

For a tree-like space T we can define the *edge tree set* $\tau(T)$ in a way that is very similar to the definition of $\tau(T)$ in Section 2.3. Let

$$\vec{E}(T) := \left\{ (\iota_e(0), \iota_e(1)) \mid e \in E(T) \right\} \cup \left\{ (\iota_e(1), \iota_e(0)) \mid e \in E(T) \right\}$$

^{||}See Lemma 4.16 in [6].

be the set of *oriented edges* of T . As tree-like spaces cannot contain loops every element of $\vec{E}(T)$ is a pair of two distinct vertices of T . For vertices $u, v \in V(T)$ let $P(u, v)$ be the unique pseudo-arc in T with end-vertices u and v . Then $\tau(T) := (\vec{E}(T), \leq, *)$ becomes a separation system by setting $(x, y)^* := (y, x)$ and $(x, y) < (v, w)$ for $(x, y), (v, w) \in \vec{E}(T)$ with $\{x, y\} \neq \{v, w\}$ whenever

$$P(y, v) \subseteq P(x, v) \subseteq P(x, w).$$

It is straightforward to check that $\tau(T)$ is a regular tree set.

2.4.4. The tree-like space of a tree set

Let $\tau = (\vec{E}, \leq, *)$ be a regular tree set; we define the *tree-like space corresponding to τ* , denoted $T(\tau)$. Let $V := \mathcal{O}(\tau)$ be the set of consistent orientations and E the set of unoriented separations of τ . As in Section 2.3 let $O(\vec{s})$ be the unique $O \in \mathcal{O}(\tau)$ in which \vec{s} is maximal. We define the tree-like space $T(\tau)$ with vertex set V and edge set E , that is with ground set $V \cup ((0, 1) \times E)$. For this we need to define the maps $\iota_e: [0, 1] \rightarrow T(\tau)$.

Fix any orientation O' of τ . For each $\vec{e} \in O'$ let $\iota_e: [0, 1] \rightarrow T$ be the map

$$\iota_e(x) = \begin{cases} O(\vec{e}), & x = 0 \\ (x, e), & 0 < x < 1 \\ O(\vec{e}), & x = 1 \end{cases} .$$

So far the definition of V and the adjacencies in $T(\tau)$ have been analogous to the construction from Section 2.3. But to make $T(\tau)$ into a graph-like space we also need to define a topology.

For $\vec{e} \in O'$ let $E^+(\vec{e})$ be the set of all $\vec{s} \in O'$ with $\vec{e} < \vec{s}$ or $\vec{e} < \vec{s}$, and $E^-(\vec{e})$ the set of all $\vec{s} \in O'$ with $\vec{s} < \vec{e}$. For $\vec{e} \in O'$ and $r \in (0, 1)$ set

$$S(\vec{e}, r) := \{O \in \mathcal{O}(\tau) \mid \vec{e} \in O\} \cup ((0, 1) \times E^+(\vec{e})) \cup ((r, 1) \times e)$$

and

$$S(\vec{e}, r) := \{O \in \mathcal{O}(\tau) \mid \vec{e} \in O\} \cup ((0, 1) \times E^-(\vec{e})) \cup ((0, r) \times e).$$

We define the sub-base of the topology on $T(\tau)$ as $\mathcal{S} := \{S(\vec{e}, r) \mid \vec{e} \in \tau, r \in (0, 1)\}$. Note that only the notation depends on the choice of O' but the topology on $T(\tau)$ does not. It is clear that $T(\tau)$ is a graph-like space: for any two vertices $a, b \in V$

pick any \vec{e} in the symmetric difference of a and b , viewed as orientations of τ . Then $S(\vec{e}, \frac{1}{2})$ and $S(\overleftarrow{e}, \frac{1}{2})$ are disjoint open sets partitioning V and $\{a, b\}$.

Lemma 2.4.9. *$T(\tau)$ is compact.*

Proof. By the Alexander sub-base theorem from general topology it suffices to show that any open covering of sets in \mathcal{S} has a finite sub-cover. Suppose that \mathcal{C} is a sub-basic open cover of $T(\tau)$ with no finite sub-cover. Let $E(\mathcal{C})$ be the set of all $\vec{e} \in \tau$ such that $S(\vec{e}, x) \in \mathcal{C}$ for some $x \in (0, 1)$. If $\vec{r} \leq \vec{s}$ for any $\vec{r}, \vec{s} \in E(\mathcal{C})$ then their corresponding sets in \mathcal{C} already cover all of $T(\tau)$, except possibly for $(0, 1) \times r$ if $\vec{r} = \vec{s}$, which can be finitely covered. Thus we may assume that $\vec{r} \not\leq \vec{s}$ for all $\vec{r}, \vec{s} \in E(\mathcal{C})$. Then the set

$$E^*(\mathcal{C}) := \{\overleftarrow{e} \mid \vec{e} \in E(\mathcal{C})\}$$

is a consistent partial orientation of τ , so by the Extension Lemma 2.2.2 there is an $O \in \mathcal{O}(\tau)$ with $E^*(\mathcal{C}) \subseteq O$. But $O \not\subseteq S(\vec{e}, r)$ for every $\vec{e} \in E(\mathcal{C})$ and $r \in (0, 1)$, so \mathcal{C} was not a cover of T . Therefore T is a compact graph-like space. \square

Lemma 2.4.10. *$T(\tau)$ is connected, but $T(\tau) - e$ is not for every $e \in E$.*

Proof. The latter follows immediately from the definition of \mathcal{S} : for any edge $e \in E$ the sets $S(\vec{e}, \frac{1}{2})$ and $S(\overleftarrow{e}, \frac{1}{2})$ define a partition of $T(\tau) - e$ into non-empty disjoint open sets.

To show that T is connected first note that any non-empty open set in T contains an inner point of an edge. Suppose that A, B are non-empty disjoint open sets partitioning T . For any edge $e \in E$ the image of ι_e in T is connected, hence every edge whose inner points meet A is completely contained in A , and similarly for B . Write τ_A for the set of $\vec{e} \in \tau$ with $\dot{e} \subseteq A$, and τ_B for the set of $\vec{e} \in \tau$ with $\dot{e} \subseteq B$. Then τ_A and τ_B partition τ and are closed under involution. Fix any $\vec{a} \in \tau_A$ and $\vec{b} \in \tau_B$ with $\vec{a} \leq \vec{b}$ and write $C := \{\vec{r} \in \tau \mid \vec{a} \leq \vec{r} \leq \vec{b}\}$ for the chain of elements between \vec{a} and \vec{b} . Let C_A be a maximal initial segment of C with $C_A \subseteq \tau_A$ and C_B a maximal initial segment of C^* with $C_B \subseteq \tau_B$, where C^* is the image of C under the involution. The set $C_A \cup C_B$ is a consistent partial orientation of τ , so by the Extension Lemma 2.2.2 there is an $O \in V$ with $C_A \cup C_B \subseteq O$. Suppose that $O \in A$, say. Let $X \subseteq \tau$ be minimal in size with the property that

$$O \in \mathcal{X} := \bigcap_{\vec{x} \in X} S(\vec{x}, r(\vec{x})) \subseteq A$$

for suitable $r(\vec{x}) \in (0, 1)$. From our assumptions it follows that such an X exists and is a finite subset of O , and the minimality implies that X is a star. Observe that $\vec{b} \subseteq S(\vec{x}, r(\vec{x}))$ for all $\vec{x} \in X$ with $\vec{x} < \vec{b}$. As \mathcal{X} does not meet B there must be a (unique) $\vec{x} \in X$ with $\vec{x} \geq \vec{b}$ and thus $\vec{x} \in C$. If $\vec{x} \in \tau_B$ then \mathcal{X} again meets B , hence $\vec{x} \in \tau_A$. As $\vec{x} \in O$ and thus $\vec{x} \notin C_A$, there is a $\vec{t} \in \tau_B \cap O$ with $\vec{x} \leq \vec{t}$. But then $\vec{t} \in \mathcal{X}$, a contradiction. Therefore $T(\tau)$ is connected. \square

Hence we have shown that $T(\tau)$ is indeed a tree-like space.

2.4.5. Regular tree sets and tree-like spaces – A characterisation

Lemma 2.4.11. *Any regular tree set τ' is isomorphic to $\tau(T(\tau'))$.*

Proof. For two vertices $u, v \in \mathcal{O}(\tau')$ the set $C = v \setminus u$ is a chain in τ' . Set

$$P(u, v) := \overline{\bigcup \{\vec{e} \mid \vec{e} \in C\}} \subseteq T(\tau').$$

Then $P(u, v) = P(v, u)$ and $P(u, v)$ is the unique pseudo-arc in T with u and v as end-vertices^{**}. Define the map $\varphi: \tau' \rightarrow \vec{E}(T(\tau'))$ as

$$\varphi(\vec{e}) := \begin{cases} (\iota_e(0), \iota_e(1)), & \vec{e} \in \iota_e(1) \\ (\iota_e(1), \iota_e(0)), & \vec{e} \in \iota_e(0) \end{cases}.$$

This is a bijection between τ' and $\vec{E}(T(\tau'))$ that commutes with the involution. The claim follows from Lemma 2.2.6 if we can show that φ is order-preserving. For this let $\vec{r}, \vec{s} \in \tau'$ with $\vec{r} < \vec{s}$. Let (x, y) be the end-vertices of $r \in E(T(\tau'))$ with $\vec{r} \in y$ and (v, w) the end-vertices of $s \in E(T(\tau'))$ with $\vec{s} \in w$. Then

$$v \setminus y = (v \setminus x) \setminus \{\vec{r}\}$$

and

$$v \setminus x = (w \setminus x) \setminus \{\vec{s}\},$$

so $P(y, v) \subseteq P(x, v) \subseteq P(x, w)$ and hence $\varphi(\vec{r}) = (x, y) \leq (v, w) = \varphi(\vec{s})$. \square

Lemma 2.4.12. *Any tree-like space T' is isomorphic to $T(\tau(T'))$.*

^{**}This follows immediately if one uses the machinery established in [6], which we do not introduce here. Alternatively one can show the connectedness of $P(u, v)$ by repeating the proof that $T(\tau')$ is connected, and verifying the other properties of a pseudo-arc directly.

Proof. For ease of notation, we may assume without loss of generality that the arbitrary orientation of $\tau(T')$ we fixed for the construction of $T(\tau(T'))$ is $\{(\iota_e^{T'}(0), \iota_e^{T'}(1)) \mid e \in E(T')\}$.

For every edge $e \in E(T')$ there is a unique $j(v, e) \in \{0, 1\}$ such that v is in the same component of $T' - e$ as $\iota_e^{T'}(j(v, e))$ by Proposition 2.4.7. We define a map $\varphi : V(T') \rightarrow V(T(\tau(T')))$ by setting $\varphi(v)$ to be the orientation

$$\{(\iota_e^{T'}(1 - j(v, e)), \iota_e^{T'}(j(v, e))) \mid e \in E(T)\}$$

of $\tau(T')$, which is easily verified to be consistent.

We extend φ to a map $T' \rightarrow T(\tau(T'))$ by setting $\varphi(r, e) := (r, \{\iota_e^{T'}(0), \iota_e^{T'}(1)\})$ for $r \in (0, 1)$ and $e \in E(T')$. It is easy to check that φ is a bijection and induces a bijection between $V(T')$ and $V(T(\tau(T')))$. Since T' is compact and $T(\tau(T'))$ is Hausdorff, we only need to check that φ is continuous. For each $e \in E(T')$ and each $r \in (0, 1)$ note that $T' \setminus \{r\}$ contains two connected components $C(e, r, 0)$ and $C(e, r, 1)$, where $C(e, r, j)$ denotes the component containing $\iota_e^{T'}(j)$. By construction, $\varphi(C(e, r, j)) = S((\iota_e^{T'}(1 - j), \iota_e^{T'}(j)), r)$ and hence the preimage of any subbasis element is open. \square

Altogether we have proven the main theorem of this section.

Theorem 2.4.13. *1. A tree set is isomorphic to the edge tree set of a tree-like space if and only if it is regular.*

2. Any regular tree set τ' is isomorphic to $\tau(T(\tau'))$.

3. Any tree-like space T' is isomorphic to $T(\tau(T'))$. \square

Additionally, for distinct but comparable tree sets, we can say precisely in which way the corresponding trees from Theorem 2.3.9 above are comparable: one will be a minor of the other.

Let us finish this section with two further results on how these constructions relate to substructures.

Theorem 2.4.14. *Let τ_1, τ_2 be regular tree-sets with $\tau_1 \subseteq \tau_2$. Then $T(\tau_1)$ is a minor of $T(\tau_2)$.*

Proof. We show that $T_1 := T(\tau_1)$ is isomorphic to $T_2 := T(\tau_2).E(T(\tau_1))$.

First we note that $\mathcal{O}(\tau_1) = \{O \cap \tau_1 \mid O \in \mathcal{O}(\tau_2)\}$. Moreover it immediately follows from the definitions that $O, O' \in \mathcal{O}(\tau_2)$ are representatives of the same vertex of T_2 if and only if $O \cap \tau_1 = O' \cap \tau_1$.

For ease of notation we may assume without loss of generality that the orientation of τ_1 that we chose in the construction of $T(\tau_1)$ is induced by the orientation we chose for τ_2 in the construction of $T(\tau_2)$. Let φ denote the concatenation of the identity from T_1 to $T(\tau_2)$ and the quotient map from $T(\tau_2)$ to T_2 . By the previous observations, this map is a bijection and induces a bijection between $V(T_1)$ and $V(T_2)$. By definition φ is continuous and hence shows that T_1 is isomorphic to T_2 . \square

Theorem 2.4.15. *Let T_1, T_2 be tree-like spaces where T_1 is a minor of T_2 . Then $\tau(T_1)$ is isomorphic to a subset of $\tau(T_2)$.*

Proof. For ease of notation we assume without loss of generality that $T_1 = T_2.E(T_1)$ and that $\iota_e^{T_2}(j) \in \iota_e^{T_1}(j)$ for all $e \in E(T_1)$ and $j \in \{0, 1\}$. We show that $\tau_1 := \tau(T_1)$ is isomorphic to $\tau_2 := \tau(T_2) \setminus \{(v, w) \mid v \in [w]\}$.

Let $\varphi : \tau_2 \rightarrow \tau_1$ be defined as $\varphi(v, w) = ([v], [w])$. It is easy to see that this map is well-defined, surjective and commutes with the involution. For the injectivity consider $(v_1, w_1), (v_2, w_2) \in \tau_2$ with $v_1 \in [v_2]$ and $w_1 \in [w_2]$ and let $e_i \in E(T_2)$ be such that $\{v_i, w_i\} = \{\iota_{e_i}^{T_2}(0), \iota_{e_i}^{T_2}(1)\}$ for $i \in \{1, 2\}$. Since $[v_2]$ and $[w_2]$ are both connected (as subspaces of T_2) but in different components of $T_2 - e_i$, we obtain that $e_1 = e_2$ and hence $(v_1, w_1) = (v_2, w_2)$.

Consider a pseudo-arc $P(v, w)$ in T_2 between any vertices v and w . It is not hard to verify that the unique pseudo-arc in T_1 between $[v]$ and $[w]$ has as its point set $\{[x] \in T_1 \mid x \in P(v, w)\}$. This observation implies that φ is order-preserving and hence an isomorphism by Lemma 2.2.6. \square

2.5. Proof of Theorem 2.4.4

Now we turn to the proof of Theorem 2.4.4. The backwards implication is clear as pseudo-arcs are connected.

For the remainder of this section let G be a compact connected graph-like space and a and b two vertices of G .

The strategy of the proof of the forward implication is as follows. Given vertices a and b which we want to connect with a pseudo-arc, first we find a minimal set L of edges which meets every a - b -cut (that is, every cut of G that separates a and b). We then want to show that the closure of these edges in G is the desired pseudo-arc. By minimality for every edge $e \in L$ there is a *signature cut*, that is, an a - b -cut for which e is the only cross-edge of L . This allows us to define a linear order on L : to compare two edges $e, f \in L$ we check on which side of e 's signature cut f lies. By extending this order to the points in the closure of L in G we can perform finite-intersection-arguments for suitable initial segments in order to prove connectedness.

We start off with a technical lemma that allows us to work with 'tidy' versions of our a - b -cuts. It also establishes that all topological cuts are finite if G is a compact graph-like space, which is important for the application of Zorn's Lemma.

Lemma 2.5.1. *Let C be a topological cut in G . Then there are disjoint open sets X and Y partitioning the vertices of G such that the edges in C are precisely those edges that are not completely in X or completely in Y . Furthermore, C is finite.*

Proof. Let X', Y' be two disjoint open sets inducing the topological cut C . Without loss of generality we may assume that every edge that meets exactly one of X', Y' is completely contained in that set. An edge that meets both X' and Y' cannot be partitioned by those two sets as it is connected. Consider the open covering F of G consisting of X', Y' and for each edge $e \in E(G)$ that meets both X' and Y' the set of inner points of e . No subsystem of F covers G , so by compactness F is a finite covering. Thus there are only finitely many edges meeting both X' and Y' , which also implies that C is finite. For every such edge e with both end-vertices in X' we can add the inner points of e to X' and delete the entire edge from Y' , and we can do the same thing for all such edges with both end-vertices in Y' . The resulting sets X and Y are still open and are as desired. \square

This lemma justifies the following formal definition of an a - b -cut.

A pair (A, B) of disjoint open sets in G is an a - b -cut if:

- (i) $a \in A$ and $b \in B$;
- (ii) $V(G) \subseteq A \cup B$;

- (iii) for every edge $e \in E(G)$ with both end-vertices in A we have $\dot{e} \in A$;
- (iv) for every edge $e \in E(G)$ with both end-vertices in B we have $\dot{e} \in B$.

That is, (A, B) is a cut separating a and b which is ‘clean’ in the sense of Lemma 2.5.1. In this case the set C of edges with end-vertices in both A and B is also called an a - b -cut, and we say that C is *induced* by (A, B) . The set of all a - b -cuts is denoted by $\mathcal{C}_{a,b}$. This set is non-empty: by the axioms of graph-like spaces there are open disjoint sets X, Y partitioning $V(G)$ and separating a and b , so the existence of an a - b -cut follows from Lemma 2.5.1.

Now we set up the application of Zorn’s Lemma to obtain a minimal set of edges that meets every a - b -cut. Let

$$X := \{e \in E(G) \mid e \in C \text{ for some } C \in \mathcal{C}_{a,b}\}.$$

This is non-empty as there is a $C \in \mathcal{C}_{a,b}$ which is non-empty by the connectedness of G . Now let

$$\mathcal{L} := \{L \subseteq X \mid L \cap C \neq \emptyset \text{ for all } C \in \mathcal{C}_{a,b}\}.$$

Since $X \in \mathcal{L}$, this set is non-empty as well. We order the elements of \mathcal{L} by inclusion. For any descending chain $(M_i \in \mathcal{L} \mid i \in I)$ the set $M := \bigcap_{i \in I} M_i$ is a lower bound in \mathcal{L} : for each $C \in \mathcal{C}_{a,b}$ every M_i contains at least one edge of C , but as C is finite, so does M . Therefore Zorn’s Lemma implies the existence of a minimal element $L \in \mathcal{L}$. We show that L is the set of edges of a pseudo-arc joining a and b .

For an edge $e \in L$ a $C \in \mathcal{C}_{a,b}$ is a *signature cut* of e if $L \cap C = \{e\}$. In that case we also call open disjoint sets (A, B) inducing C a signature cut of e . Such a cut exists for every $e \in L$ by the minimality of L .

Note that if (A, B) is a signature cut of an edge $e \in L$, then for any other $f \in L$ either $\dot{f} \subseteq A$ or $\dot{f} \subseteq B$.

For an edge $e \in L$ with end-vertices $x \neq y$ and a signature cut (A, B) of e we say that e *runs from* x to y if $x \in A$ and $y \in B$.

For two edges $e, f \in L$ we set $e < f$ if there is a signature cut (A, B) of e with $\dot{f} \subseteq B$. Furthermore, we set $e \leq f$ for all edges $e \in L$.

Before proceeding we need to check that neither the orientation of an edge $e \in L$ nor the definition of $e < f$ depends on the signature cut at hand, and that \leq is a linear order on L . The general strategy in the following proofs is this: assume

a counterexample to the claim exists. Consider the signature cuts of all edges involved, then for a contradiction find a suitable corner or union of corners of these cuts that is still an a - b -cut but contains no edge of L .

Lemma 2.5.2. *If $e \in L$ runs from x to y , then $x \in A$ and $y \in B$ for all signature cuts (A, B) of e . Furthermore if $e < f$ for $e, f \in L$, then $\overset{\circ}{f} \subseteq B$ for all signature cuts (A, B) of e .*

Proof. Suppose there is an $e \in L$ with end-vertices x, y and signature cuts (A_1, B_1) and (A_2, B_2) , for which $x \in A_1 \cap B_2$ and $y \in A_2 \cap B_1$. But then $(A_1 \cap A_2, B_1 \cup B_2)$ would induce an a - b -cut containing no edge of L : all edges of L apart from e have both their end-vertices either in $B_1 \cup B_2$ or in $A_1 \cap A_2$, and e has no end-vertex in $A_1 \cap A_2$. This contradicts the definition of L . Hence $x \in A$ and $y \in B$ for all signature cuts (A, B) of e .

Now suppose there are edges $e, f \in L$ and signature cuts $(A_1, B_1), (A_2, B_2)$ of e such that $f \in B_1 \cap A_2$. Let (A_3, B_3) be a signature cut of f . If $\overset{\circ}{e} \subseteq A_3$, then the bipartition $(A_1 \cup A_2 \cup A_3, B_1 \cap B_2 \cap B_3)$ induces an a - b -cut containing no edge of L . But if $\overset{\circ}{e} \subseteq B_3$ then $(A_1 \cap A_2 \cap A_3, B_1 \cup B_2 \cup B_3)$ induces an a - b -cut containing no edge of L , a contradiction. Hence if $e < f$, then $\overset{\circ}{f} \subseteq B$ for all signature cuts (A, B) of e . \square

Lemma 2.5.3. *The relation \leq on L is a linear order.*

Proof. It is reflexive: this is true by definition.

Every two edges of L are comparable: suppose there are two distinct edges $e, f \in L$ with respective signature cuts (A_1, B_1) and (A_2, B_2) , for which $\overset{\circ}{e} \subseteq A_2$ and $\overset{\circ}{f} \subseteq A_1$. Then $(A_1 \cap A_2, B_1 \cup B_2)$ induces an a - b -cut containing no edge of L , a contradiction.

It is antisymmetric: suppose there are two distinct edges $e, f \in L$ with respective signature cuts (A_1, B_1) and (A_2, B_2) , for which $\overset{\circ}{e} \subseteq B_2$ and $\overset{\circ}{f} \subseteq B_1$. Then $(A_1 \cup A_2, B_1 \cap B_2)$ induces an a - b -cut containing no edge of L , a contradiction.

It is transitive: suppose there are three distinct edges $e, f, g \in L$, $e < f$ and $f < g$, with signature cuts (A_1, B_1) of e and (A_2, B_2) of f for which $\overset{\circ}{f} \subseteq B_1$ and $\overset{\circ}{g} \subseteq B_2$ but $\overset{\circ}{g} \subseteq A_1$. Then $(A_1 \cup A_2, B_1 \cap B_2)$ is a signature cut of f (as $\overset{\circ}{e} \subseteq A_2$) with $\overset{\circ}{g} \subseteq A_1 \cup A_2$, which contradicts $f < g$. \square

Finally we define the pseudo-arc that shall join a and b . Write \bar{L} for

$$\bar{L} := \overline{\bigcup\{\dot{e} \mid e \in L\}}.$$

As G is compact \bar{L} is a compact subspace of G . Furthermore the removal of any edge $e \in L$ from \bar{L} (that is, removal of \dot{e}) separates a and b in \bar{L} as any signature cut of e witnesses.

To prove that \bar{L} is connected we perform finite-intersection arguments on suitable initial segments of \bar{L} . In order for this to be possible we first need to extend the order \leq on L to an order \prec on \bar{L} .

Let (A, B) be a signature cut of some $e \in L$ and $x \in \bar{L} \setminus \dot{e}$. Then we write $x \prec e$ if $x \in A$, and $x \succ e$ if $x \in B$. For $x, y \in \bar{L}$ we write $x \preceq y$ if any of the following holds:

- (i) there are edges $e, f \in L$ with $x \in \dot{e}$, $y \in \dot{f}$ and $e < f$;
- (ii) there is an edge $e \in L$ with $x \prec e \prec y$;
- (iii) there is an edge $e \in L$ with end-vertices v, w , running from v to w , such that $x, y \in \dot{e}$ and $\iota^{-1}(x) < \iota^{-1}(y)$ in the parametrization ι of e with $\iota(0) = v$ and $\iota(1) = w$.

In addition we set $x \preceq x$ for all $x \in \bar{L}$.

As for \leq we prove in the following lemma that \prec is well-defined in the sense that $x \prec e$ implies $x \in A$ for all signature cuts (A, B) of e .

Lemma 2.5.4. *If $x \prec e$ for $x \in \bar{L} \setminus \dot{e}$ and $e \in L$ then $x \in A$ for all signature cuts (A, B) of e .*

Proof. Suppose there are two signature cuts $(A_1, B_1), (A_2, B_2)$ of e with $x \in A_1, B_2$. If x is an end-vertex of e this is an immediate contradiction to Lemma 2.5.2. If x is not an end-vertex of e consider $D := (A_1 \cap B_2) \setminus \bar{\dot{e}}$. This is an open set containing x , so since $x \in \bar{L}$, there is an edge $f \neq e$ with $f \in L$ and $D \cap \dot{f} \neq \emptyset$. But then $\dot{f} \subseteq D$, contradicting Lemma 2.5.2 as well. \square

As one readily checks \preceq is a partial order on \bar{L} . If $x, y \in \bar{L}$ are incomparable then x and y are both vertices that are not the end-vertex of any edge in L . To show that \bar{L} is a pseudo-arc from a to b we need to show that any two vertices $x, y \in \bar{L}$

are separated in $\bar{L} - \hat{e}$ for some $e \in L$. That is, we need to show that \preceq is a linear order on \bar{L} . We shall achieve this with a finite intersection property argument for initial segments of \bar{L} .

Let $C \in \mathcal{C}_{a,b}$ be an a - b -cut and $L(C) := L \cap C = \{e_1, \dots, e_n\}$ with $e_1 < \dots < e_n$. For $k \in [n+1]$ the k -th *segment* of \bar{L} with regard to C is the set

$$S_C(k) := \{x \in \bar{L} \mid e_{k-1} \prec x \prec e_k\}$$

for $1 < k < n+1$, and $S_C(1) := \{x \in \bar{L} \mid x \prec e_1\}$ as well as

$$S_C(n+1) := \{x \in \bar{L} \mid x \succ e_n\}.$$

As in the analogous scenario with paths and cuts in graphs one would expect the segments of \bar{L} with regard to an a - b -cut (A, B) to alternate between being contained in A or in B . The next lemma shows that this is the case, and helps locate an edge which separates two given vertices in \bar{L} .

Lemma 2.5.5. *Let $C \in \mathcal{C}_{a,b}$ be induced by (A, B) with $L(C) = \{e_1, \dots, e_n\}$ and $e_1 < \dots < e_n$. For $k \in [n+1]$ the following statements hold.*

- (i) *If k is odd then $S_C(k) \subseteq A$;*
- (ii) *If k is even then $S_C(k) \subseteq B$.*

In particular, if an edge $e_k \in L(C)$ has end-vertices x, y with $x \preceq y$, then e_k runs from x to y if k is odd and from y to x if k is even.

Proof. For clarity we only consider the case where k is odd; the other case follows analogously.

First assume that $k = 1$. Suppose for a contradiction that there is an $x \in S_C(1)$ with $x \in B$. Let (A_1, B_1) be a signature cut of e_1 . Then $x \in B \cap A_1$ as $x \prec e_1$. Due to $x \in \bar{L}$ there has to be an edge $f \in L$ with $\overset{\circ}{f} \cap (B \cap A_1) \neq \emptyset$. This implies $\overset{\circ}{f} \subseteq B \cap A_1$ and in particular $e_1 \neq f$. Let (A_f, B_f) be a signature cut of f . Then $(A \cap A_1 \cap A_f, B \cup B_1 \cup B_f)$ is an a - b -cut not containing any edge of L : suppose $g \in L$ is an edge with end-vertices v, w such that $v \in A \cap A_1 \cap A_f$ and $w \in B \cup B_1 \cup B_f$. Then $w \in A_1 \cap A_f$ implying $w \in B$ and thus $g \in L(C)$, but also $g < e_1$, a contradiction.

If $k > 1$, then suppose for a contradiction that there is an $x \in S_C(k)$ with $x \in B$. Let (A_{k-1}, B_{k-1}) and (A_k, B_k) be signature cuts of e_{k-1} and e_k respectively. Then $x \in B \cap B_{k-1} \cap A_k$ as $e_{k-1} \prec x \prec e_k$. Due to $x \in \bar{L}$ there has to be an edge $f \in L$ with $\overset{\circ}{f} \cap (B \cap B_{k-1} \cap A_k) \neq \emptyset$. This implies $\overset{\circ}{f} \subseteq B \cap B_{k-1} \cap A_k$ and in particular $f \neq e_{k-1}, e_k$. Let (A_f, B_f) be a signature cut of f . Then

$$\left((B_{k-1} \cap B_f) \cap (A \cup (B \cap B_k)), A_{k-1} \cup A_f \cup (B \cap A_k) \right)$$

is an a - b -cut not containing any edge of L : suppose $g \in L$ is an edge with end-vertices v and w such that $v \in (B_{k-1} \cap B_f) \cap (A \cup (B \cap B_k))$ and $w \in A_{k-1} \cup A_f \cup (B \cap A_k)$. Then $w \in B_{k-1} \cap B_f$ and therefore $w \in B \cap A_k$, implying $v \in A_k$ and thus $v \in A$. Hence $g \in L(C)$ but $e_{k-1} < g < e_k$, a contradiction. \square

Lemma 2.5.5 indeed implies that any two vertices of \bar{L} can be separated by some $e \in L$.

Lemma 2.5.6. *Let $v \neq w$ be two vertices in \bar{L} . Then there is an edge $e \in L$ which separates v and w in \bar{L} .*

Proof. If C is an a - b -cut with v and w on different sides, then by Lemma 2.5.5 v and w are in different segments, $S_C(k_v)$ and $S_C(k_w)$, say. For $k := \min\{k_v, k_w\}$ the edge $e_k \in L(C)$ separates v and w in \bar{L} : as $x \prec e \prec y$ for any signature cut (A, B) of e we have $x \in A$ and $y \in B$, which gives a partition of $\bar{L} \setminus \overset{\circ}{e}$ into two relatively open sets.

It is thus left to show that an a - b -cut with v and w on different sides exists. Let (A, B) be any a - b -cut and (V, W) be a v - w -cut. If v and w are on different sides of (A, B) or if (V, W) is an a - b -cut we are done. If not, then $v, w \in A$ and $a, b \in V$, say. But then $(A \cap V, B \cup W)$ is the desired cut. \square

From this it follows that \preceq is in fact a linear order on \bar{L} . Next we prove that $a \in \bar{L}$ (which, surprisingly, is not obvious) by finding a minimum of \bar{L} and showing that this minimum has to be a .

Note that for any vertex $c \neq a$ there is an a - b -cut with c on the b -side: let (A, B) be an a - b -cut and (A', C) be an a - c -cut. Then $(A \cap A', B \cup C)$ is the desired cut.

Lemma 2.5.7. *The minimum of \bar{L} with regard to \preceq is a and the maximum is b . In particular $a, b \in \bar{L}$.*

Proof. We only show this for a .

If L has a minimum $m \in L$, let a' be the smaller one of its end-vertices (that is, m runs from a' to its other end-vertex). Then a' is the minimum of \bar{L} by Lemma 2.5.6. Suppose $a \neq a'$. Let C be an a - b -cut induced by (A, B) with $a' \in B$. Then $a' \notin S_C(1)$, so $e_1 \prec a'$ implying $e_1 < m$ a contradiction to the minimality of m .

If L does not have a minimum then for $e \in L$ set

$$X_e := \overline{\{f \mid f \in L, f < e\}}.$$

Then $X_e \subseteq \bar{L}$ for all $e \in L$. Since G is compact, \bar{L} has the finite intersection property. Therefore

$$X := \bigcap_{e \in L} X_e \neq \emptyset.$$

For any edge $e \in L$ no inner point $x \in \mathring{e}$ of e is in X , as $x \notin X_e$. Thus X contains a vertex a' . If there were another vertex $a'' \in X$, then a' and a'' could be separated by an edge $e \in L$ by Lemma 2.5.6 and one of them would not be in X_e . So $X = \{a'\}$. Suppose $a \neq a'$. Let C be an a - b -cut induced by (A, B) with $a' \in B$ and let $L(C) = \{e_1, \dots, e_n\}$ with $e_1 < \dots < e_n$. Then $a' \notin S_C(1)$ as $a' \in B$, so $e_1 \prec a'$. But this means $a' \notin X_{e_1}$, a contradiction. \square

The final property needed of \bar{L} to be a pseudo-arc joining a and b is that it is connected. The proof of this is similar to the proof of Lemma 2.5.7.

Lemma 2.5.8. *The subspace \bar{L} of G is connected.*

Proof. Suppose $X, Y \subseteq \bar{L}$ are two non-empty disjoint sets partitioning \bar{L} which are open in the subspace topology of \bar{L} with $a \in X$. As edges are connected, $\mathring{e} \subseteq X$ or $\mathring{e} \subseteq Y$ for all $e \in L$. Let $S := \{e \in L \mid \mathring{e} \subseteq Y\}$ and $\bar{S} := \overline{\{\mathring{e} \mid e \in S\}}$. Then S is non-empty as Y contains a point of \bar{L} and thus an inner point of an edge of L .

We aim to find a minimum of $Y = \bar{S}$ with regard to \preceq .

If S has a minimum $m \in S$ with regard to \leq then let y be the smaller one of its end-vertices. Then $y \in Y$ and $y \preceq z$ for all $z \in \bar{S}$.

If S does not have a minimum then for $e \in S$ set

$$R_e := \overline{\{f \mid f \in S, f < e\}}.$$

Every R_e is a non-empty closed subset of \bar{L} . By the finite intersection property $R := \bigcap_{e \in S} R_e$ is non-empty. For any edge $e \in S$ no inner point $x \in \mathring{e}$ of e is in R ,

as $y \notin R_e$. Thus R contains a vertex y . If there were another vertex $y' \in R$, then y and y' could be separated by an edge $e \in L$ by Lemma 2.5.6, with $y \prec e \prec y'$, say. This edge e cannot be in S as in that case y would not be in R_e . Thus $\dot{e} \subseteq X$. Let (A, B) be a signature cut of e . As $e < f$ for all $f \in S$ due to $e \prec y' \prec f$ we have $y \in A$ and

$$\bigcup \{\dot{f} \mid f \in S\} \subseteq B.$$

But then $A \cap \bar{L}$ witnesses that $y \notin \bar{S}$, a contradiction.

Therefore $R = \{y\}$ and y is the minimum of \bar{S} .

Now set

$$X' := \{x \in X \mid x \prec e \text{ for all } e \in S\}$$

and let $U := \{e \in L \mid \dot{e} \subseteq X'\}$. By a similar argument as above X' has a maximum x . Let y be the minimum of $Y = \bar{S}$ and $e \in L$ an edge separating x and y . If $y \prec e \prec x$ then either $e \in S$ and $x \notin X'$ or $e \in U$ and $y \notin Y$. So $x \prec e \prec y$, which implies $e \in U$. But this contradicts the fact that x is the maximum of X' . \square

We have succeeded in proving that \bar{L} is a pseudo-arc containing a and b . This concludes our proof of Theorem 2.4.4. \square

3. Infinite end-devouring sets of rays with prescribed start vertices

3.1. Introduction

Looking for spanning structures in infinite graphs such as spanning trees or Hamilton cycles often involves difficulties that are not present when one considers finite graphs. It turned out that the concept of ends of an infinite graph is crucial for questions dealing with such structures. Especially for locally finite graphs ends allow us to define these objects in a more general topological setting [11].

Nevertheless, the definition of an end of a graph is purely combinatorial: For any graph G we call two rays *equivalent* in G if they cannot be separated by finitely many vertices. It is easy to check that this defines an equivalence relation on the set of all rays in the graph G . The equivalence classes of this relation are called the *ends* of G and a ray contained in an end ω of G is referred to as an ω -ray.

When we focus on the structure of ends of an infinite graph G , we observe that *normal spanning trees* of G , i.e. rooted spanning trees of G such that the endvertices of every edge of G are comparable in the induced tree-order, have a powerful property: For any normal spanning tree T of G and every end ω of G there is a unique ω -ray in T which starts at the root of T and has the property that it meets every ω -ray of G , see [9, Section 8.2]. For any graph G , we say that an ω -ray with this property *devours* the end ω of G . Similarly, we say that a set of ω -rays *devours* ω if every ω -ray in G meets at least one ray out of the set. Note that if a set of ω -rays devours ω , then every ω -ray R meets the union of that set infinitely often, since each tail of R meets at least one ray out of the set.

End-devouring sets of rays are helpful for the construction of spanning structures such as infinite Hamilton cycles. For example, in a one-ended locally finite graph after removing an end-devouring set of rays, each component is finite. Thomassen [47] used this fact to show that the square of each locally finite 2-connected one-ended

graph contains a spanning ray. Georgakopoulos [23] generalised this to locally finite 2-connected graphs with arbitrary many ends by building some other kind of spanning structure with the help of an end-devouring set of rays, which he then used to construct an infinite Hamilton cycle in the square of such a graph. He proved the following proposition about the existence of finite sets of rays devouring any *countable end*, i.e., an end which does not contain uncountably many disjoint rays. Note that the property of an end being countable is equivalent to the existence of a finite or countably infinite set of rays devouring the end.

Proposition. [23] *Let G be a graph and ω be a countable end of G . If G has a set \mathcal{R} of $k \in \mathbb{N}$ disjoint ω -rays, then it also has a set \mathcal{R}' of k disjoint ω -rays that devours ω . Moreover, \mathcal{R}' can be chosen so that its rays have the same start vertices as the rays in \mathcal{R} .*

For the proof of this proposition Georgakopoulos recursively applies a construction similar to the one yielding normal spanning trees to find rays for the end-devouring set. However, this proof strategy does not suffice to give a version of this proposition for infinitely many rays. He conjectured that such a version remains true [23, Problem 1]. We confirm this conjecture with the following theorem, which also covers the proposition above.

Theorem 3.1.1. *Let G be a graph, ω a countable end of G and \mathcal{R} any set of disjoint ω -rays. Then there exists a set \mathcal{R}' of disjoint ω -rays that devours ω and the start vertices of the rays in \mathcal{R} and \mathcal{R}' are the same.*

Note that, in contrast to the proposition, the difficulty of Theorem 3.1.1 for an infinite set \mathcal{R} comes from fixing the set of start vertices, since any inclusion-maximal set of disjoint ω -rays devours ω .

The proof of Theorem 3.1.1 will feature in Section 3.2. In Section 3.3 we will see why this theorem does not immediately extend to ends that contain an uncountable set of disjoint rays. There we discuss an additional necessary condition on the set of start vertices.

3.2. Theorem

All graphs in this chapter are simple and undirected.

For a finite set M of vertices of a graph G and an end ω of G , let $C(M, \omega)$ denote the unique component of $G - M$ that contains a tail of every ω -ray.

For the proof of Theorem 3.1.1 we shall use the following characterisation of ω -rays.

Lemma 3.2.2. *Let G be a graph, ω an end of G and \mathcal{R}_{\max} an inclusion-maximal set of pairwise disjoint ω -rays. A ray $R \subseteq G$ is an ω -ray, if and only if it meets rays of \mathcal{R}_{\max} infinitely often.*

Proof. Let W denote the set $\bigcup\{V(R) \mid R \in \mathcal{R}_{\max}\}$.

If R is an ω -ray, then each tail of R meets a ray of \mathcal{R}_{\max} since \mathcal{R}_{\max} is inclusion-maximal. Hence R meets W infinitely often.

Suppose for a contradiction that R is an ω' -ray for an end $\omega' \neq \omega$ of G that contains infinitely many vertices of W . Let M be a finite set of vertices such that the two components $C := C(M, \omega)$ and $C' := C(M, \omega')$ of $G - M$ are different. By the pigeonhole principle there is either one ω -ray of \mathcal{R}_{\max} containing infinitely many vertices of $V(C') \cap V(R) \cap W$, or infinitely many disjoint rays of \mathcal{R}_{\max} containing those vertices. In both cases we get an ω -ray with a tail in C' , since we cannot leave C' infinitely often through the finite set M . But this contradicts the definition of C . \square

A natural strategy for constructing up to infinitely many disjoint rays is to inductively construct them in countably many steps. In each step we fix only finitely many finite paths as initial segments instead of whole rays, while extending previously fixed initial segments and ensuring that they can be extended to rays. This strategy is for example used by Halin [30, Satz 1] to prove that the maximum number of disjoint rays in an end is well-defined. Our proof of Theorem 3.1.1 is also based on that strategy. In order to guarantee that the set of rays we construct turns out to devour the end, we also fix an inclusion-maximal set of vertex disjoint rays of our specific end, so a countable set, and an enumeration of the vertices on these rays. Then we try in each step to either contain or separate the least vertex with respect to the enumeration that is not already dealt with from the end with appropriately chosen initial segments if possible. Otherwise, we extend a

finite number of initial segments while still ensuring that all initial segments can be extended to rays. Although it is impossible to always contain or separate the considered vertex with our initial segments while being able to continue with the construction, it will turn out that the rays we obtain as the union of all initial segments actually do this.

For a vertex v and an end ω of a graph G we say that a vertex set $X \subseteq V(G)$ *separates v from ω* if there does not exist any ω -ray that is disjoint from X and contains v .

Furthermore, in addition to the notations for paths introduced in the beginning of this theses, for Q being a v - w path we write $v\bar{Q}$ for the path that is obtained from Q by deleting w .

Proof of Theorem 3.1.1. Let us fix a finite or infinite enumeration $\{R_j \mid j < |\mathcal{R}|\}$ of the rays in \mathcal{R} . Furthermore, let s_j denote the start vertex of R_j for every $j < |\mathcal{R}|$ and define $S := \{s_j \mid j < |\mathcal{R}|\}$.

Next we fix an inclusion-maximal set \mathcal{R}_{\max} of pairwise disjoint ω -rays and an enumeration $\{v_i \mid i \in \mathbb{N}\}$ of the vertices in $W := \bigcup\{V(R) \mid R \in \mathcal{R}_{\max}\}$. This is possible since ω is countable by assumption.

We do an inductive construction such that the following holds for every $i \in \mathbb{N}$:

- (1) P_s^i is a path with start vertex s for every $s \in S$.
- (2) $P_s^i = s$ for all but finitely many $s \in S$.
- (3) $P_s^{i-1} \subseteq P_s^i$ for every $s \in S$.
- (4) For every $s = s_j \in S$ with $j < \min\{i, |\mathcal{R}|\}$ there is a $w_s^i \in W \cap (P_s^i \setminus P_s^{i-1})$.
- (5) P_s^i and $P_{s'}^i$ are disjoint for all $s, s' \in S$ with $s \neq s'$.
- (6) For every $s \in S$ there exists an ω -ray R_s^i with P_s^i as initial segment and s as start vertex such that all rays R_s^i are pairwise disjoint.

If possible and not spoiling any of the properties (1) to (6), we incorporate the following property:

- (*) $\bigcup_{s \in S} P_s^i$ either contains v_{i-1} or separates v_{i-1} from ω if $i > 0$.

We begin the construction for $i = 0$ by defining $P_s^0 := s =: P_s^{-1}$ for every $s \in S$. All conditions are fulfilled as witnessed by $R_{s_j}^0 := R_j$ for every $j < |\mathcal{R}|$.

Now suppose we have done the construction up to some number $i \in \mathbb{N}$. If we can continue with the construction in step $i + 1$ such that properties (1) to (6) together with (*) hold, we do so and define all initial segments P_s^{i+1} and rays R_s^{i+1} accordingly. Otherwise, we set for all $s \in S$

$$\begin{aligned} P_s^{i+1} &:= sR_s^i w_s^i && \text{if } s = s_j \text{ for } j < \min\{i + 1, |\mathcal{R}|\} \text{ and} \\ P_s^{i+1} &:= P_s^i && \text{otherwise,} \end{aligned}$$

where w_s^i denotes the first vertex of W on $R_s^i \setminus P_s^i$ which exist by Lemma 3.2.2. With these definitions properties (1) up to (5) hold for $i + 1$. Witnessed by $R_s^{i+1} := R_s^i$ for every $s \in S$ we immediately satisfy (6) too. This completes the inductive part of the construction.

Using the paths P_s^i we now define the desired ω -rays of \mathcal{R}' . We set $R'_s := \bigcup_{i \in \mathbb{N}} P_s^i$ for every $s \in S$ and $\mathcal{R}' := \{R'_s \mid s \in S\}$. Properties (1), (3) and (4) ensure that R'_s is a ray with start vertex s for every $s \in S$, while we obtain due to property (5) that all rays R'_s are pairwise disjoint. Property (4) also ensures that all rays in \mathcal{R}' are ω -rays by Lemma 3.2.2.

It remains to prove that the set \mathcal{R}' devours the end ω . Suppose for a contradiction that there exists an ω -ray R disjoint from $\bigcup \mathcal{R}'$. By the maximality of our chosen set of ω -rays \mathcal{R}_{\max} , we know that R contains a vertex v_j for some $j \in \mathbb{N}$. In the next paragraph we will show how we could have proceeded in step $j + 1$ to incorporate property (*) as well. For an easier understanding of the technical definitions of that paragraph we refer to Figure 3.2.1.

Without loss of generality, let v_j be the start vertex of R . Let P be an $R \cup \mathcal{R}'$ path among those ones that are disjoint from $\bigcup_{s \in S} s\bar{P}_s^{j+1}$ for which $v_j R p$ is as short as possible where p denotes the common vertex of P and R . Such a path exists, because all rays in $\mathcal{R}' \cup \{R\}$ are equivalent and $\bigcup_{s \in S} s\bar{P}_s^{j+1}$ is finite by property (2). Let $t \in S$ and $q \in V(G)$ be such that $V(P) \cap V(R'_t) = \{q\}$. Furthermore, let R^* be an ω -ray with start vertex $r^* \in R$ such that R^* is disjoint from $\bigcup_{s \in S} R'_s$ and $P(pRr^*) \cap R^* = \{r^*\}$ for which $v_j R r^*$ is as short as possible. Since p and pR are candidates for r^* and R^* , respectively, such a choice is possible. We define

$$\hat{P}_t^{j+1} := (tR'_t q)P(pRr^*) \quad \text{and} \quad \hat{R}_t^{j+1} := \hat{P}_t^{j+1} R^* ;$$

and replace in step $j + 1$ the ray R_t^{j+1} by \hat{R}_t^{j+1} , the path P_t^{j+1} by \hat{P}_t^{j+1} and for all $s \in S \setminus \{t\}$ the ray R_s^{j+1} by R'_s while keeping P_s^{j+1} as it was defined. By this construction properties (1) to (6) are satisfied.




Figure 3.2.1.: Sketch of the situation above. The rays in \mathcal{R}' are drawn vertically, with their fixed initial segments from step $j + 1$. Horizontally drawn is the ray R that is supposed to contradict that \mathcal{R}' devours ω with its start vertex $v_j \in W$. The $R \cup \mathcal{R}'$ path P is chosen with its vertex p on R as close to v_j as possible. The ray R^* is chosen disjoint to the rays in \mathcal{R}' and except from its start vertex r^* on R disjoint from the initial segment of R upto p again with r^* as close to v_j as possible. The ray \hat{R}_t^{j+1} is highlighted in grey with its initial segment fixed up to r^* .

Now we show that (*) holds as well. Suppose for a contradiction that there exists an ω -ray Z disjoint from $(\bigcup_{s \in S \setminus \{t\}} P_s^{j+1}) \cup \hat{P}_t^{j+1}$ with start vertex v_j . First note that Z is disjoint from $r^*Rp \subseteq \hat{P}_t^{j+1}$. Let us now show that Z is also disjoint from $pR \cup \bigcup_{s \in S} R'_s$. Otherwise, let z denote the first vertex along Z that lies in $pR \cup \bigcup_{s \in S} R'_s$. However, z cannot be contained in pR , as this would contradict the choice of r^* , and it cannot be an element of $\bigcup_{s \in S} R'_s$ since this would contradict the choice of p . But now with Z being not only disjoint from $pR \cup \bigcup_{s \in S} R'_s$ but also from r^*Rp , we get again a contradiction to the choice of r^* . Hence, we would have been able to incorporate property (*) without violating any of the properties (1) to (6) in step $j + 1$ of our construction. This, however, is a contradiction since we

always incorporated property $(*)$ under the condition of maintaining properties (1) to (6). So we arrived at a contradiction to the existence of the ray R since by $(*)$ every ray containing v_j meets the initial segments of rays fixed in our construction at step $j + 1$. Therefore, the set \mathcal{R}' devours the end ω . \square

3.3. Ends of uncountable degree

Given an end ω of some graph of uncountable degree, then by reasons of cardinality it cannot be devoured by a set of ω -rays which is strictly smaller than the degree of ω . But, unlike in the countable degree case, the existence of a set of $\deg(\omega)$ many disjoint ω -rays is not sufficient for existence of a set of disjoint ω -rays devouring ω with the same start vertices. We illustrate an obvious obstruction.

A *separation of a graph G* is a tuple (A, B) with $A \cup B = V(G)$ such that there are no edges between $A \setminus B$ and $B \setminus A$. Suppose G contains a separation (A, B) such that both $G[A \setminus B]$ and $G[B \setminus A]$ contain a set of disjoint ω -rays of cardinality more than $|A \cap B|$. At least one of $G[A \setminus B]$ or $G[B \setminus A]$ contains a set \mathcal{R} of $\deg(\omega)$ many disjoint ω -rays, say $G[B \setminus A]$. But no set \mathcal{R}' of disjoint ω -rays with the same start vertices as the rays in \mathcal{R} can devour ω since at most $|A \cap B|$ many rays meet vertices of $A \setminus B$ and hence cannot meet all ω -rays in $G[A \setminus B]$.

For an easy example of this obstruction consider two sets A and B of size $\kappa > \aleph_0$ such that $\aleph_0 \leq |A \cap B| < \kappa$ and let G be the union of the complete graphs on A and B . Then (A, B) a separation where both $G[A \setminus B]$ and $G[B \setminus A]$ contain a set of κ many disjoint rays to the unique end of G .

Hence we can state two necessary conditions for a set $S \subseteq V(G)$ to be a set of start vertices for a set of disjoint ω -rays devouring ω :

- there is a set \mathcal{R} of disjoint ω -rays with S as its start vertices; and
- for each separation (A, B) of G , if $G[A \setminus B]$ contains a set of more than $|A \cap B|$ disjoint ω -rays, then $A \setminus B$ contains a vertex of S .

Problem 3.3.3. *Are these conditions together also sufficient for the existence of a set of disjoint ω -rays devouring ω with S as its start vertices?*

Note that our construction for the proof of Theorem 3.1.1, if continued transfinitely, might face numerous new problems at limit steps.

4. Characterising k -connected sets in infinite graphs

4.1. Introduction

A common aspect of structural graph theory is the study of the duality between connectivity and tree structure. Such type of duality theorems assert that if a graph contains no ‘highly connected part’, then there is some kind of tree structure with certain properties, usually a tree-decomposition of the graph, that, if it exists, clearly precludes the existence of such a ‘highly connected part’. Some of the more well-known examples include the duality between brambles and tree-width, as well as tangles and branch-width.

One of the notions of connectivity, which has been studied in finite graphs, is the one of a so called k -connected set. For $k \in \mathbb{N}$, a set X of at least k vertices of a graph G is called *k -connected in G* , if for all $Z_1, Z_2 \subseteq X$ with $|Z_1| = |Z_2| \leq k$ there are $|Z_1|$ many vertex disjoint paths from Z_1 to Z_2 in G . We often omit stating the graph in which X is k -connected if it is clear from the context.

In finite graphs, k -connected sets have also been studied in connection to tree-width. This connection was first observed by Robertson, Seymour and Thomas [43], and later improved by Diestel, Gorbunov, Jensen and Thomassen [13, Prop. 3], who showed that for any finite graph G and $k \in \mathbb{N}$, if G contains a $(k + 1)$ -connected set of size at least $3k$, then G has tree-width at least k , and conversely if G has no $(k + 1)$ -connected set of size at least $3k$, then G has tree-width less than $4k$.

Recently, Geelen and Joeris [22, 31] studied the duality between k -connected sets and *k -tree-width*, that is the analogue of tree-width when only considering tree-decompositions of adhesion less than k . They showed that the maximum size of a k -connected set is bounded from below by the k -tree-width w and from above by $\binom{w+1}{k-1}(k-1)$ [22, Thm. 1.2].

For infinite graphs there exist different notions of how to cut up a graph in

a tree-like way which extend the notion of tree-decompositions for finite graphs. Robertson, Seymour and Thomas [42] gave a survey of different results characterising the existence of different kinds of these decompositions via forbidden minors. In recent years, one of those decomposition notions, the notion of a *nested set of separations* has been studied in more detail [12, 28]. They correspond to tree-decompositions of finite graphs in a natural way and offer a generalisation for infinite graphs. We define separations and the necessary terms, including the notion of parts for a nested set of separations, which provides some analogue of tree-width, in Section 4.2. These nested separation systems shall allow us to prove the following duality theorem:

Theorem 4.1.1. *Let G be an infinite graph, let $k \in \mathbb{N}$ and let $\kappa \leq |V(G)|$ be an infinite cardinal. Then the following statements are equivalent.*

- (a) *$V(G)$ contains a subset of size κ that is k -connected in G .*
- (b) *There is no nested set of separations of order less than k of G such that every part has size less than κ .*

Our second main result, Theorem 4.1.2, describes how k -connected sets “look like” by characterising their existence with the existence of certain unavoidable (topological) minors*.

It is a well-known and easy-to-prove fact that every large connected finite graph contains a long path or a vertex of high degree. More precisely, for every $m \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that each connected graph with at least n vertices either contains a path P_m of length m or a star $K_{1,m}$ with m leaves as a subgraph (cf. [9, Prop. 9.4.1]). In a sense which can be made precise [9, Thm. 9.4.5], the existence of these ‘unavoidable’ subgraphs characterises connectedness with respect to the subgraph relation in a minimal way: in every infinite collection consisting of graphs with arbitrarily large connected subgraphs we find arbitrarily long paths or arbitrarily large stars as subgraphs; but with only paths or only stars this would not be true, since long paths and large stars do not contain each other.

*Since the non-existence of a large k -connected set in a graph is a property which is closed under both the minor and topological minor relation, the existence of k -connected sets in a graph is a property which is well-suited to be characterised via the existence of certain (topological) minors.

For 2-connected graphs there is an analogous result, which also is folklore: For every $m \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that every 2-connected finite graph with at least n vertices either contains a subdivision of a cycle C_m of length m or a subdivision of a complete bipartite graph $K_{2,m}$ [9, Prop. 9.4.2]. As before, cycles and $K_{2,m}$'s as unavoidable topological minors characterise 2-connectedness with respect to the topological minor relation in a minimal way (cf. [9, Thm. 9.4.5]).

In 2016, Geelen and Joeris [22, 31] generalised these results to arbitrary $k \in \mathbb{N}$. For this, they relaxed ‘ k -connectedness’ to containing a large k -connected set. They introduced certain graphs called generalised wheels (depending on k and m), which together with the complete bipartite graph $K_{k,m}$ are the unavoidable minors: they contain large k -connected sets themselves and they characterise graphs that contain large k -connected sets with respect to the minor relation.

This result encompasses the characterisation for $k = 2$ as mentioned above, as well as earlier results from Oporowski, Oxley and Thomas [37], who proved similar results for $k = 3$ and $k = 4$ (albeit with different notions of ‘ k -connectedness’).

Now let us consider infinite graphs. Again there is a well-known and easy-to-prove fact that each infinite connected graph contains either a *ray*, that is a one-way infinite path, or a vertex of infinite degree. This can also be seen as a characterisation of infinite connected graphs via these two unavoidable subgraphs: a ray and the complete bipartite graph K_{1,\aleph_0} . There is also a more localised version of this result: the Star-Comb Lemma (cf. Lemma 4.2.3). In essence this lemma relates these subgraphs to a given vertex set.

For 2-connected infinite graphs one can easily construct an analogous result. We say a vertex d *dominates* a ray R if they cannot be separated by deleting a finite set of vertices not containing d . Given a ray R and the complete graph on two vertices K_2 , the *one-way infinite ladder* is the graph $R \times K_2$. Now it is a common exercise to prove that the unavoidable (topological) minors for 2-connectivity are the one-way infinite ladder, the union of a ray R with a complete bipartite graph between a single vertex and $V(R)^\dagger$ as well as the complete bipartite graph K_{2,\aleph_0} .

In 1978, Halin [29] studied such a problem for arbitrary $k \in \mathbb{N}$. He showed that every k -connected graph whose set of vertices has size at least κ for some

[†]Note that with the advent of topological infinite graph theory, those results became an even more meaningful extension of the finite result as these unavoidable minors correspond to infinite cycles in locally finite or finitely separable graphs (cf. [9, Section 8.6]) and [11, Section 5]).

uncountable regular cardinal κ contains a subdivision of $K_{k,\kappa}$. Hence for all those cardinals, $K_{k,\kappa}$ is the unique unavoidable topological minor characterising graphs with a subdivision of a k -connected graph of size κ . In a way, this characterisation result is stronger than the results previously discussed, since it obtains a direct equality between the size of a k -connected set and the size of the minors which was not possible for finite graphs. The unavoidable (topological) minors for graphs whose set of vertices has singular cardinality remained undiscovered.

Oporowski, Oxley and Thomas [37] also studied countably infinite graphs for arbitrary $k \in \mathbb{N}$, but again for a different notion of ‘ k -connectedness’. Together with the K_{k,\aleph_0} , the unavoidable minors for countably infinite essentially k -connected[‡] graphs have the following structure. For $\ell, d \in \mathbb{N}$ with $\ell + d = k$, they consist of a set of ℓ disjoint rays, d vertices that dominate one of the rays (or equivalently all of those rays) and infinitely many edges connecting pairs of them in a tree-like way.

This leads to our second main result, Theorem 4.1.2. For $k \in \mathbb{N}$ and an infinite cardinal κ we will define certain graphs with a k -connected set of size κ in Section 4.3, the so called *k -typical graphs*. These graphs will encompass complete bipartite graphs $K_{k,\kappa}$ as well as the graphs described by Oporowski, Oxley and Thomas [37] for $\kappa = \aleph_0$. We will moreover introduce such graphs even for singular cardinals κ . It will turn out that for fixed k and κ there are only finitely many k -typical graphs up to isomorphisms. We shall characterise graphs with a k -connected set of size κ via the existence of a minor of such a k -typical graph with a k -connected set of size κ . In contrast to the finite case, the minimality of the list of these graphs in the characterisation is implied by the fact that it really is a finite list of graphs (for which even if it were not minimal we could pick a minimal sublist), and not a finite list of ‘classes of graphs’, like ‘paths’ and ‘stars’ in the finite case for $k = 1$.

Moreover we will extend the definition of k -typical graphs to so called *generalised k -typical graphs*. As before for fixed k and κ there are only finitely many generalised k -typical graphs up to isomorphisms, and we shall extend the characterisation from before to be with respect to the topological minor relation.

[‡]A graph is *essentially k -connected* if there is a constant $c \in \mathbb{N}$ such that for each separation (A, B) of order less than k one of A or B has size less than c . As before, we will not use this notion in this chapter.

Theorem 4.1.2. *Let G be an infinite graph, let $k \in \mathbb{N}$ and let $\kappa \leq |V(G)|$ be an infinite cardinal. Then the following statements are equivalent.*

- (a) $V(G)$ contains a subset of size κ that is k -connected in G .
- (b) G contains a k -typical graph of size κ as a minor with finite branch sets.
- (c) G contains a subdivision of a generalised k -typical graph of size κ .

In fact, we will prove a slightly stronger result which will require some more notation, Theorem 4.3.7 in Subsection 4.3.3. In the same vein as the Star-Comb Lemma, that result will relate the minors (or subdivisions) with a specific k -connected set in the graph.

After fixing some notation and recalling some basic definitions and simple facts in Section 4.2, we will define the k -typical graphs and generalised k -typical graphs in Section 4.3. In Section 4.4 we will collect some basic facts about k -connected sets and their behaviour with respect to minors or topological minors. Section 4.5 deals with the structure of ends in graphs. Subsection 4.5.1 is dedicated to extend a well-known connection between minimal separators and the degree of an end from locally finite graphs to arbitrary graphs. Afterwards, Subsection 4.5.2 gives a construction of how to find disjoint rays in some end with additional structure between them. Sections 4.6 and 4.7 are dedicated to prove Theorem 4.1.2. The case of κ being a regular cardinal is covered in Section 4.6 and, respectively, the case of κ being a singular cardinal is covered in Section 4.7. In Section 4.8 we will talk about some applications of the characterisation via minors, and in Section 4.9 we shall prove Theorem 4.1.1.

4.2. Preliminaries

Throughout this chapter, let G denote an arbitrary simple and undirected graph.

Given two sets A and B , we denote by $K(A, B)$ the complete bipartite graph between the classes A and B . We also write $K_{\kappa, \lambda}$ for $K(A, B)$ if $|A| = \kappa$ and $|B| = \lambda$ for two cardinals κ and λ .

We shall need the following version of Menger's Theorem for finite parameter k in infinite graphs, which is an easy corollary of Menger's Theorem for finite graphs.

Theorem 4.2.1. [9, Thm. 8.4.1] *Let $k \in \mathbb{N}$ and let $A, B \subseteq V(G)$. If A and B cannot be separated by less than k vertices, then G contains k disjoint $A-B$ paths.*

We shall also need a trivial cardinality version of Menger's theorem, which is easily obtained from Theorem 4.2.1 by noting that the union of less than κ many disjoint $A-B$ paths for an infinite cardinal κ has size less than κ (cf. [9, Section 8.4]).

Theorem 4.2.2. *Let κ be a cardinal and let $A, B \subseteq V(G)$. If A and B cannot be separated by less than κ vertices, then G contains κ disjoint $A-B$ paths. \square*

Recall that a vertex $d \in V(G)$ *dominates* a ray R if d and some tail of R lie in the same component of $G - S$ for every finite set $S \subseteq V(G) \setminus \{d\}$. By Theorem 4.2.2 this is equivalent to the existence of infinitely many $d-R$ paths in G which are disjoint but for d itself. Note that if d dominates an ω -ray, then it also dominates every other ω -ray. Hence we also write that d *dominates* an end $\omega \in \Omega(G)$ if d dominates some ω -ray. Let $\text{Dom}(\omega)$ denote the set of vertices dominating ω and let $\text{dom}(\omega) = |\text{Dom}(\omega)|$. If $\text{dom}(\omega) > 0$, we call ω *dominated*, and if $\text{dom}(\omega) = 0$, we call ω *undominated*.

For an end $\omega \in \Omega(G)$, let $\Delta(\omega)$ denote $\deg(\omega) + \text{dom}(\omega)$, which we call the *combined degree* of ω . Note that the sum of an infinite cardinal with some other cardinal is just the maximum of the two cardinals.

A *comb* C is the union of a ray R together with infinitely many disjoint finite paths each of which has precisely one vertex in common with R , which has to be an endvertex of that path. The ray R is the *spine* of C and the end vertices of the finite paths that are not on R together with the end vertices of the trivial paths are the *teeth* of C . A comb whose spine is in ω is also called an ω -comb. A *star* is the complete bipartite graph $K_{1,\kappa}$ for some cardinal κ , where the vertices of degree 1 are its *leaves* and the vertex of degree κ is its *centre*.

Next we state a version of the Star-Comb lemma in a slightly stronger way than elsewhere in the literature (e.g. [9, Lemma 8.2.2]). We also give a proof for the sake of completeness.

Lemma 4.2.3 (Star-Comb Lemma). *Let $U \subseteq V(G)$ be infinite and let $\kappa \leq |U|$ be a regular cardinal. Then the following statements are equivalent.*

- (a) *There is a subset $U_1 \subseteq U$ with $|U_1| = \kappa$ such that U_1 is 1-connected in G .*

(b) There is a subset $U_2 \subseteq U$ with $|U_2| = \kappa$ such that G either contains a subdivided star whose set of leaves is U_2 or a comb whose set of teeth is U_2 .

(Note that if κ is uncountable, only the former can exist.)

Moreover, if these statements hold, we can choose $U_1 = U_2$.

Proof. Note that a set of vertices is 1-connected, if and only if it belongs to the same component of G . Hence if (b) holds, then U_2 is 1-connected and we can set $U_1 := U_2$ to satisfy (a).

If (a) holds, then we take a tree $T \subseteq G$ containing U_1 such that each edge of T lies on a path between two vertices of U_1 . Such a tree exists by Zorn's Lemma since U_1 is 1-connected in G . We distinguish two cases.

If T has a vertex c of degree κ , then this yields a subdivided star with centre c and a set $U_2 \subseteq U_1$ of leaves with $|U_2| = \kappa$ by extending each incident edge of c to a $c-U_1$ path.

Hence we assume T does not contain a vertex of degree κ . Given some vertex $v_0 \in V$ and $n \in \mathbb{N}$, let D_n denote the vertices of T of distance n to v_0 . Since T is connected, the union $\bigcup\{D_n \mid n \in \mathbb{N}\}$ equals $V(T)$. And because κ is regular, it follows that $\kappa = \aleph_0$, and therefore that T is locally finite. Hence each D_n is finite and, since T is still infinite, each D_n is non-empty. Thus T contains a ray R by König's Infinity Lemma. If R does not already contain infinitely many vertices of U_1 , then by the property of T there are infinitely many edges of T between $V(R)$ and $V(T - R)$. We can extend infinitely many of these edges to a set of disjoint $R-U_1$ paths, ending in an infinite subset $U_2 \subseteq U_1$, yielding the desired comb.

In both cases, U_2 is still 1-connected, and hence serves as a candidate for U_1 as well, yielding the "moreover" part of the claim. \square

The following immediate remark helps to identify when we can obtain stars by an application of the Star-Comb lemma.

Remark 4.2.4. If there is an ω -comb with teeth U and if v dominates ω , then there is also a set $U' \subseteq U$ with $|U'| = |U| = \aleph_0$ such that G contains a subdivided star with leaves U' and centre v . \square

We say that an end ω is in the *closure* of a set $U \subseteq V(G)$, if there is an ω -comb whose teeth are in U . Note that this combinatorial definition of closure coincides

with the topological closure when considering the topological setting of locally finite graphs mentioned in the introduction [9, Section 8.6; 11].

For an end ω of G and an induced subgraph G' of G we write $\omega \upharpoonright G'$ for the set of rays $R \in \omega$ which are also rays of G' . The following remarks are immediate.

Remarks 4.2.5. Let $G' = G - S$ for some finite $S \subseteq V(G)$.

1. $\omega \upharpoonright G'$ is an end of G' for every end $\omega \in \Omega(G)$.
2. For every end $\omega' \in \Omega(G')$ there is an end $\omega \in \Omega(G)$ such that $\omega \upharpoonright G' = \omega'$.
3. The degree of $\omega \in \Omega(G)$ in G is equal to the degree of $\omega \upharpoonright G'$ in G' .
4. $\text{Dom}(\omega) = \text{Dom}(\omega \upharpoonright G') \cup (\text{Dom}(\omega) \cap S)$ for every end $\omega \in \Omega(G)$. □

Let us fix some notations regarding minors. Let G and M be graphs. We say M is a *minor* of G if G contains an *inflated subgraph* $H \subseteq G$ witnessing this, i.e. for each $v \in V(M)$

- there is a non-empty *branch set* $\mathfrak{B}(v) \subseteq V(H)$;
- $H[\mathfrak{B}(v)]$ is connected;
- $\{\mathfrak{B}(v) \mid v \in V(M)\}$ is a partition of $V(H)$; and
- there is an edge between $v, w \in V(M)$ in M if and only if there is an edge between some vertex in $\mathfrak{B}(v)$ and a vertex in $\mathfrak{B}(w)$ in H .

We call M a *finite-branch-set minor* or *fbs-minor* of G if each branch set is finite.

Without loss of generality we may assume that such an inflated subgraph H witnessing that M is a minor of G is minimal with respect to the subgraph relation. Then H has the following properties for all $v, w \in V(M)$:

- $H[\mathfrak{B}(v)]$ is a finite tree T_v ;
- for each $v, w \in V(M)$ there is a unique edge e_{vw} in $E(H)$ between $\mathfrak{B}(v)$ and $\mathfrak{B}(w)$ if $vw \in E(M)$, and no such edge if $vw \notin E(M)$;
- each leaf of T_v is an endvertex of such an edge between two branch sets.

Given a subset $C \subseteq V(M)$ and a subset $A \subseteq V(G)$, we say that M is an *fbs-minor of G with A along C* , if M is an fbs-minor of G such that the map mapping each vertex of the inflated subgraph to the branch set it is contained in induces a bijection between A and the branch sets of C . As before, we assume without loss of generality that an inflated subgraph H witnessing that M is an fbs-minor of G is minimal with respect to the subgraph relation. We obtain the properties as above, but a leaf of T_v could be the unique vertex of A in $\mathfrak{B}(v)$ instead.

For $\ell, k \in \mathbb{N}$, we write $[\ell, k]$ for the *closed integer interval* $\{i \in \mathbb{N} \mid \ell \leq i \leq k\}$ as well as $[k, \ell)$ for the *half open integer interval* $\{i \in \mathbb{N} \mid \ell \leq i < k\}$.

Given some set I , a *family* \mathcal{F} indexed by I is a sequence of the form $(F_i \mid i \in I)$, where the *members* F_i are some not necessarily different sets. For convenience we sometimes use a family and the set of its members with a slight abuse of notation interchangeably, for example with common set operations like $\bigcup \mathcal{F}$. Given some $J \subseteq I$, we denote by $\mathcal{F} \upharpoonright J$ the *subfamily* $(F_j \mid j \in J)$. A set T is a *transversal* of \mathcal{F} , if $|T \cap F_i| = 1$ for all $i \in I$. For a family $(F_i \mid i \in \mathbb{N})$ with index set \mathbb{N} we say some property holds for *eventually all* members, if there is some $N \in \mathbb{N}$ such that the property holds for F_i for all $i \in \mathbb{N}$ with $i \geq N$.

The following lemma is a special case of the famous Delta-Systems Lemma, a common tool of infinite combinatorics.

Lemma 4.2.6. [33, Thm. II.1.6] *Let κ be a regular cardinal, U be a set and $\mathcal{F} = (F_\alpha \subseteq U \mid \alpha \in \kappa)$ a family of finite subsets of U . Then there is a finite set $D \subseteq U$ and a set $I \subseteq \kappa$ with $|I| = \kappa$ such that $F_\alpha \cap F_\beta = D$ for all $\alpha, \beta \in I$ with $\alpha \neq \beta$.*

A *separation* of G is a tuple (A, B) of vertex sets such that $A \cup B = V(G)$ and such that there is no edge of G between $A \setminus B$ and $B \setminus A$. The set $A \cap B$ is the *separator* of (A, B) and the cardinality $|A \cap B|$ is called the *order* of (A, B) . Given $k \in \mathbb{N}$, let $S_k(G)$ denote the set of all separations of G of order less than k .

Two separations (A, B) and (C, D) are *nested* if one of the following conditions hold:

$$\begin{aligned} A \subseteq C \text{ and } D \subseteq B, & \quad \text{or} & \quad B \subseteq C \text{ and } D \subseteq A, & \quad \text{or} \\ A \subseteq D \text{ and } C \subseteq B, & \quad \text{or} & \quad B \subseteq D \text{ and } C \subseteq A. \end{aligned}$$

A set N of separations of G is called a *nested separation system of G* if it is *symmetric*, i.e. $(B, A) \in N$ for each $(A, B) \in N$ and *nested*, i.e. the separations in N are pairwise nested.

An *orientation* O of a nested separation system N is a subset of N that contains precisely one of (A, B) and (B, A) for all $(A, B) \in N$. An orientation O of N is *consistent* if whenever $(A, B) \in O$ and $(C, D) \in N$ with $C \subseteq A$ and $B \subseteq D$, then $(C, D) \in O$. For each consistent orientation O of N we define a *part* P_O of N as the vertex set $\bigcap \{B \mid (A, B) \in O\}$. It is easy to check that the union of all parts cover the vertex set of G . Moreover, we allow the empty set \emptyset as a nested separation system. In this case, we say that $V(G)$ is a part of \emptyset (this can be viewed as the empty intersection of vertex sets of the empty set as an orientation of \emptyset).

A nested separation system N has *adhesion less than k* if all separations it contains have order less than k , i.e. $N \subseteq S_k(G)$.

Note that each oriented edge of the tree of a tree-decomposition of G induces a separation (A, B) where A is the union of the parts on one side of the edge while B is the union of the parts on the other side of the edge. It is easy to check that the set of separations induced by all those edges is a nested separation system. Moreover, properties like adhesion and the size of parts are transferred by this process.

For more information on nested separation systems in a more abstract way and their connection to tree-decompositions we refer the interested reader to [12], as well as Chapter 2 of this thesis.

In Section 4.9 we will make use of the existence of k -lean tree-decompositions for finite graphs to prove our desired duality theorem, which are closely related to k -connected sets. Given $k \in \mathbb{N}$, a tree-decomposition of adhesion less than k is called *k -lean* if for any two (not necessarily distinct) parts V_{t_1}, V_{t_2} of the tree-decomposition and vertex sets $Z_1 \subseteq V_{t_1}, Z_2 \subseteq V_{t_2}$ with $|Z_1| = |Z_2| = \ell \leq k$ there are either ℓ disjoint Z_1 - Z_2 paths in G or there is an edge tt' on the t_1 - t_2 path in the tree inducing a separation of order less than ℓ . In particular, given a k -lean tree-decomposition, each part V_t is $\min\{k, |V_t|\}$ -connected in G .

In [7], the authors noted that the proof given in [5] of a theorem of Thomas [46, Thm. 5] about the existence of lean tree-decompositions witnessing the tree-

width of a finite graph can be adapted to prove the existence of a k -lean tree-decomposition of that graph.

Theorem 4.2.7. [7, Thm. 2.3] *Every finite graph has a k -lean tree-decomposition for any $k \in \mathbb{N}$.*

This definition can easily be lifted to nested separation systems. A nested separation system $N \subseteq S_k(G)$ is called k -lean if given any two (not necessarily distinct) parts P_1, P_2 of N and vertex sets $Z_1 \subseteq P_1, Z_2 \subseteq P_2$ with $|Z_1| = |Z_2| = \ell \leq k$ there are either ℓ disjoint $Z_1 - Z_2$ paths in G or there is a separation (A, B) in N with $P_1 \subseteq A$ and $P_2 \subseteq B$ of order less than ℓ . Here, we specifically allow the empty set as a nested separation system to be k -lean if its part, the whole vertex set of G , is $\min\{k, |V(G)|\}$ -connected. Again, we obtain that each part P of a k -lean nested separation system is $\min\{k, |P|\}$ -connected in G . Moreover, note that the nested separation system that a k -lean tree-decomposition induces is k -lean as well.

4.3. Typical graphs with k -connected sets

Throughout this section, let $k \in \mathbb{N}$ be fixed. Let κ denote an infinite cardinal.

In Subsection 4.3.1 we shall describe an up to isomorphism finite class of graphs each of which contains a designated k -connected set of size κ . We call such a graph a k -typical graph and the designated k -connected set its *core*. These graphs will appear as the minors of Theorem 4.1.2(b).

In Subsection 4.3.2 we shall describe based on these k -typical graphs a more general but still finite class of graphs each of which again contains a designated k -connected set of size κ . We call such a graph a *generalised k -typical graph* and the designated k -connected set its *core*. These graphs will appear as the topological minors of Theorem 4.1.2(c).

4.3.1. k -typical graphs

The most basic graph with a k -connected set of size κ is a complete bipartite graph $K_{k,\kappa} = K([0, k], Z)$ for any infinite cardinal κ and a set Z of size κ disjoint from $[0, k]$. Although in this graph the whole vertex set is k -connected, we only want to consider the infinite side Z as the *core* $C(K_{k,\kappa})$ of $K_{k,\kappa}$, cf. Figure 4.3.1. This is the first instance of a k -typical graph with a core of size κ . For uncountable

regular cardinals κ , this is the only possibility for a k -typical graph with a core of size κ .




Figure 4.3.1.: A stylised version of a $K_{4,\kappa}$, where the large box stands for the core of κ many vertices and the dashed lines from a vertex to the corners of the box represent that this vertex is connected to all vertices in the box.

A k -blueprint \mathcal{B} is a tuple (B, D) such that

- B is a tree of order k ; and
- D is a set of leaves of B with $|D| < |V(B)|$.

Take the ray $\mathfrak{N} := (\mathbb{N}, \{n(n+1) \mid n \in \mathbb{N}\})$ and the Cartesian product $B \times \mathfrak{N}$. For a node $b \in V(B)$ and $n \in \mathbb{N}$ let

- b_n denote the vertex (b, n) ;
- \mathfrak{N}_b denote the ray $(\{b\}, \emptyset) \times \mathfrak{N} \subseteq B \times \mathfrak{N}$; and
- B_n denote the subgraph $B \times (\{n\}, \emptyset) \subseteq B \times \mathfrak{N}$.

Then let $\mathfrak{N}(B/D) := (B \times \mathfrak{N}) / \{\mathfrak{N}_d \mid d \in D\}$ denote the contraction minor of $B \times \mathfrak{N}$ obtained by contracting each ray \mathfrak{N}_d for each $d \in D$ to a single vertex. We denote the vertex of $\mathfrak{N}(B/D)$ corresponding to the contracted ray \mathfrak{N}_d by d for $d \in D$ and call such a vertex *dominating*. Using this abbreviated notation, we call the tree $B_n - D$ the n -th layer of $\mathfrak{N}(B/D)$.

A triple $\mathcal{B} = (B, D, c)$ is called a *regular k -blueprint* if (B, D) is a k -blueprint and $c \in V(B) \setminus D$. We denote by $T_k(\mathcal{B})$ the graph $\mathfrak{N}(B/D)$ and by $C(T_k(\mathcal{B}))$ the vertex set $V(\mathfrak{N}_c)$, which we call the *core* of $T_k(\mathcal{B})$, see Figure 4.3.2 for an example.

Lemma 4.3.1. *For a regular k -blueprint \mathcal{B} the core of $T_k(\mathcal{B})$ is k -connected in $T_k(\mathcal{B})$.*




Figure 4.3.2.: Image of $T_4(P, \{d\}, c)$ where $P = cbd$ is a path of length 3 between nodes c and d . P_0 is represented in gray. The crosses represent its core.

Proof. Let $\mathcal{B} = (B, D, c)$ and let $C = C(T_k(\mathcal{B}))$ denote the core of $T_k(\mathcal{B})$. Let $U, W \subseteq C$ with $|U| = |W| = k' \leq k$. Suppose for a contradiction there is a vertex set S of size less than k' separating U and W . Then there are $m, n \in \mathbb{N}$ with $c_m \in U \setminus S$, $c_n \in W \setminus S$ such that the n -th and m -th layer are both disjoint from S . Moreover there is a $b \in B$ such that \mathfrak{N}_b (or $\{b\}$ if $b \in D$) are disjoint from U and W . Hence we can connect c_m and c_n with the path consisting of the concatenation of the unique $c_m - b_m$ path in B_m , the unique $b_m - b_n$ path in \mathfrak{N}_b and the unique $b_n - c_n$ path in B_n . This path avoids S , contradicting that S is a separator. By Theorem 4.2.1 there are k' disjoint $U - W$ paths, and hence C is k -connected in $T_k(\mathcal{B})$. \square

For any regular k -blueprint $\mathcal{B} = (B, D, c)$ the graph $T_k(\mathcal{B})$ is a k -typical graph with a countable core. Such graphs are besides the complete bipartite graph K_{k, \aleph_0} the only other k -typical graphs with a core of size \aleph_0 .

Note that given two regular k -blueprints $\mathcal{B}_1 = (B_1, D_1, c_1)$ and $\mathcal{B}_2 = (B_2, D_2, c_2)$ such that there is an isomorphism φ between B_1 and B_2 that maps D_1 to D_2 , then $T_k(\mathcal{B}_1)$ and $T_k(\mathcal{B}_2)$ are isomorphic. Moreover, if φ maps c_1 to c_2 , then there is an isomorphism between $T_k(\mathcal{B}_1)$ and $T_k(\mathcal{B}_2)$ that maps the core of $T_k(\mathcal{B}_1)$ to the core of $T_k(\mathcal{B}_2)$. Hence up to isomorphism there are only finitely many k -typical graphs with a core of size \aleph_0 .

Given a singular cardinal κ we have more possibilities for typical graphs with k -connected sets of size κ . We call a sequence $\mathcal{K} = (\kappa_\alpha < \kappa \mid \alpha \in \text{cf } \kappa)$ of infinite cardinals a *good κ -sequence*, if

- it is cofinal, i.e. $\bigcup \mathcal{K} = \kappa$;
- it is strictly ascending, i.e. $\kappa_\alpha < \kappa_\beta$ for all $\alpha < \beta$ with $\alpha, \beta \in \text{cf } \kappa$;
- $\text{cf } \kappa < \kappa_\alpha < \kappa$ for all $\alpha \in \text{cf } \kappa$; and
- κ_α is regular for all $\alpha \in \text{cf } \kappa$.

Note that given any $I \subseteq \text{cf } \kappa$ with $|I| = \text{cf } \kappa$ there is a unique order-preserving bijection between $\text{cf } \kappa$ and I . Hence we can relabel any cofinal subsequence $\mathcal{K} \upharpoonright I$ of a good κ -sequence \mathcal{K} to a good κ -sequence $\overline{\mathcal{K}} \upharpoonright I$. Moreover, note that any cofinal sequence can be made into a good κ -sequence by looking at an strictly ascending subsequence starting above the cofinality of κ , then replacing each element in the sequence by its successor cardinal and relabeling as above. Here we use the fact that each successor cardinal is regular. Hence for every singular cardinal κ there is a good κ -sequence.

Let $\mathcal{K} = (\kappa_\alpha < \kappa \mid \alpha \in \text{cf } \kappa)$ be a good κ -sequence and let $\ell \leq k$ be a non-negative integer. As a generalisation of the graph $K_{k,\kappa}$ we first consider the disjoint union of the complete bipartite graphs K_{k,κ_α} . Then we identify ℓ sets of vertices each consisting of a vertex of the finite side of each graph, and connect the other $k - \ell$ vertices of each with disjoint stars $K_{1,\text{cf } \kappa}$. More formally, let $X = [\ell, k) \times \{0\}$, and for each $\alpha \in \text{cf } \kappa$ let $Y^\alpha = \{\alpha\} \times [0, k) \times \{1\}$ and let $Z^\alpha = \{\alpha\} \times \kappa_\alpha \times \{2\}$. We denote the family $(Y^\alpha \mid \alpha \in \text{cf } \kappa)$ with \mathcal{Y} and the family $(Z^\alpha \mid \alpha \in \text{cf } \kappa)$ with \mathcal{Z} . Then consider the union $\bigcup \{K(Y^\alpha, Z^\alpha) \mid \alpha \in \text{cf } \kappa\}$ of the complete bipartite graphs and let $\ell\text{-}K(k, \mathcal{K})$ denote the graph where for each $i \in [0, \ell)$ we identify the set $\text{cf } \kappa \times \{i\} \times \{1\}$ to one vertex in that union. For this graph we fix some further notation. Let

- x_i denote $(i, 0) \in X$ for $i \in [\ell, k)$;
- $y_i = y_i^\alpha$ for all $\alpha \in \text{cf } \kappa$ denote the vertex corresponding to $\text{cf } \kappa \times \{i\} \times \{1\}$ for $i \in [0, \ell)$; we call such a vertex a *degenerate vertex of $\ell\text{-}K(k, \mathcal{K})$* ;
- y_i^α denote $(\alpha, i, 1)$ for $i \in [\ell, k)$; and
- \mathcal{Y}_i denote $(y_i^\alpha \mid \alpha \in I)$ for $i \in [\ell, k)$.

Note that while the definition of $\ell\text{-}K(k, \mathcal{K})$ formally depends on the choice of a good κ -sequence, the structure of the graph is independent of that choice.

Remark 4.3.2. ℓ - $K(k, \mathcal{K}_0)$ is isomorphic to a subgraph of ℓ - $K(k, \mathcal{K}_1)$, and vice versa, for any two good κ -sequences $\mathcal{K}_0, \mathcal{K}_1$. \square

Given ℓ - $K(k, \mathcal{K})$ as above, let S_i denote the star $K(\{x_i\}, \cup \mathcal{Y}_i)$ for all $i \in [\ell, k]$. Consider the union of ℓ - $K(k, \mathcal{K})$ with $\cup_{i \in [\ell, k]} S_i$. We call this graph ℓ - $FK_{k, \kappa}(\mathcal{K})$, or an ℓ -degenerate frayed $K_{k, \kappa}$ (with respect to \mathcal{K}). As before, any vertex y_i for $i \in [0, \ell)$ is called a *degenerate vertex* of ℓ - $FK_{k, \kappa}(\mathcal{K})$, and any x_i for $i \in [\ell, k]$ is called a *frayed centre* of ℓ - $FK_{k, \kappa}(\mathcal{K})$. The *core* $C(\ell$ - $FK_{k, \kappa}(\mathcal{K}))$ of ℓ - $FK_{k, \kappa}(\mathcal{K})$ is the vertex set $\cup \mathcal{Z}$. As with $K_{k, \kappa}$ it is easy to see that $C(\ell$ - $FK_{k, \kappa}(\mathcal{K}))$ is k -connected in ℓ - $FK_{k, \kappa}(\mathcal{K})$ and of size κ .

Note that Remark 4.3.2 naturally extends to ℓ - $FK_{k, \kappa}(\mathcal{K})$. Hence for each κ we now fix a specific good κ -sequence and write just ℓ - $FK_{k, \kappa}$ when talking about an ℓ -degenerate frayed $K_{k, \kappa}$ regarding that sequence, see Figure 4.3.3 for an example. Further note that k - $FK_{k, \kappa}$ is isomorphic to $K_{k, \kappa}$. We also call a 0-degenerate frayed $K_{k, \kappa}$ just a *frayed* $K_{k, \kappa}$ or $FK_{k, \kappa}$ for short.

For a singular cardinal κ and for any $\ell \in [0, k]$ the graph ℓ - $FK_{k, \kappa}$ is a k -typical graph with a core of size κ . These are besides the complete bipartite graph $K_{k, \kappa}$ the only other k -typical graphs with a core of size κ if κ has uncountable cofinality.




Figure 4.3.3.: Image of 2 - $FK_{4, \kappa}$. The black squares represent the frayed centres and the white squares the degenerate vertices. Its core is represented by the union of the boxes (labelled according to the fixed good κ -sequence) and has size κ as illustrated by the bracket.

Next we will describe the other possibilities of k -typical graphs for singular cardinals with countable cofinality.

A *singular k -blueprint* \mathcal{B} is a 5-tuple (ℓ, f, B, D, σ) such that

- $0 \leq \ell + f < k$;
- (B, D) is a $(k - \ell - f)$ -blueprint with $2 \cdot |D| \leq |V(B)|$; and
- $\sigma : [\ell + f, k) \rightarrow V(B - D) \times \{0, 1\}$ is an injective map.

Let $\mathcal{B} = (\ell, f, B, D, \sigma)$ be a singular k -blueprint and let $\mathcal{K} = (\kappa_\alpha < \kappa \mid \alpha \in \mathbb{N}_0)$ be a good κ -sequence. We construct our desired graph $T_k(\mathcal{B})(\mathcal{K})$ as follows. We start with ℓ -FK $_{k,\kappa}(\mathcal{K})$ with the same notation as above. We remove the set $\{x_i \mid i \in [\ell + f, k)\}$ from the graph we constructed so far. Moreover, we take the disjoint union with $\mathfrak{N}(B/D)$ as above. We identify the vertices $\{y_i^\alpha \mid i \in [\ell + f, k)\}$ with distinct vertices of the $(2\alpha + |V(B)|)$ -th and $(2\alpha + 1 + |V(B)|)$ -th layer for every $\alpha \in \mathbb{N}_0$ as given by the map σ , that is

$$y_i^\alpha \sim \pi_0(\sigma(i))_{2\alpha + \pi_1(\sigma(i)) + |V(B)|}$$

where π_0 and π_1 denote the projection maps for the tuples in the image of σ . For convenience we denote a vertex originated via such an identification by any of its previous names. The *core* of $T_k(\mathcal{B})(\mathcal{K})$ is $C(T_k(\mathcal{B})(\mathcal{K})) := \cup \mathcal{Z}$. For an example we refer to Figure 4.3.4.




Figure 4.3.4.: Image of $T_7(1, 2, P, \{d\}, \sigma)$ for P and $\{d\}$ as in Figure 4.3.2.

As before, the information given by a specific good κ -sequence does not matter for the structure of the graph. Similarly, we get with Remark 4.3.2 that two

graphs $T_k(\mathcal{B})(\mathcal{K}_0)$ and $T_k(\mathcal{B})(\mathcal{K}_1)$ obtained by different good κ -sequences $\mathcal{K}_0, \mathcal{K}_1$ are isomorphic to fbs-minors of each other. Hence when we use the fixed good κ -sequence as before, we call the graph just $T_k(\mathcal{B})$.

Lemma 4.3.3. *For a singular k -blueprint \mathcal{B} , the core of $T_k(\mathcal{B})$ is k -connected in $T_k(\mathcal{B})$.*

Proof. Let $\mathcal{B} = (\ell, f, B, D, \sigma)$ and let C denote the core of $T_k(\mathcal{B})$. Let $U, W \subseteq C$ with $|U| = |W| = k' \leq k$. Suppose for a contradiction there is a vertex set S of size less than k' separating U and W . This separator needs to contain all degenerate vertices as well as block all paths via the frayed centres. Hence there are less than $k' - \ell - f$ many vertices of S on $\mathfrak{N}(B/D)$, and therefore there is either a $b \in V(B) \setminus D$ such that either \mathfrak{N}_b does not contain a vertex of S or a $d \in D \setminus S$. Moreover, there are $m, n \in \mathbb{N}$ such that $u^m \in (U \cap Z^m) \setminus S$ and $w^n \in (W \cap Z^n) \setminus S$. Now $n \neq m$ since S cannot separate two vertices of $Z^n \setminus S$ in $K(Y^n, Z^n) \subseteq T_k(\mathcal{B})$. Since the vertices of $Y^n \cap \mathfrak{N}(B/D)$ lie on at least $(k - \ell - f)/2$ different rays of the form \mathfrak{N}_x for $x \in V(B) \setminus D$, there is a vertex $v^n \in (Y^n \cap \mathfrak{N}(B/D)) \setminus S$ such that the ray \mathfrak{N}_x that contains v^n either has no vertices of S on its tail starting at v^n or on its initial segment up to v^n . Also, there is an $N \in \mathbb{N}$ with $N \geq n$ in the first case and $N \leq n$ in the second case (since $n \geq |V(B)|$) such that B_N does not contain a vertex of S . Hence we can find a path avoiding S starting at w^n and ending on the ray \mathfrak{N}_b or the dominating vertex d . Analogously, we get $v^m \in (Y^m \cap \mathfrak{N}(B/D)) \setminus S$, B_M and a respective path avoiding S . Hence we can connect u^m and w^n via a path avoiding S , contradicting the assumption. \square

For a singular cardinal κ with countable cofinality and for any singular k -blueprint \mathcal{B} the graph $T_k(\mathcal{B})$ is a k -typical graph with core of size κ . These are the only remaining k -typical graphs.

Note that as before there are up to isomorphism only finitely many k -typical graphs with a core of size κ .

In summary we get for each $k \in \mathbb{N}$ and each infinite cardinal κ a finite list of k -typical graphs with a core of size κ :

κ	k -typical graph T	core $C(T)$
$\kappa = \text{cf } \kappa > \aleph_0$	$K_{k,\kappa}$	Z
$\kappa = \text{cf } \kappa = \aleph_0$	$K_{k,\kappa}$	Z
	$T_k(B, D, c)$	$V(\mathfrak{N}_c)$
$\kappa > \text{cf } \kappa > \aleph_0$	$K_{k,\kappa}$	Z
	ℓ - $FK_{k,\kappa}$	$\cup Z$
$\kappa > \text{cf } \kappa = \aleph_0$	$K_{k,\kappa}$	Z
	ℓ - $FK_{k,\kappa}$	$\cup Z$
	$T_k(\ell, f, B, D, \sigma)$	$\cup Z$

Note that for the finiteness of this list we need the fixed good κ -sequence for the singular cardinal κ .

Lemma 4.3.4. *The core of a k -typical graph is k -connected in that graph.* \square

4.3.2. Generalised k -typical graphs

The k -typical graphs cannot serve for a characterisation for the existence of k -connected sets as in Theorem 4.1.2(c) via subdivisions, as the following example illustrates. Consider two disjoint copies of the K_{2,\aleph_0} together with a matching between the infinite sides, see Figure 4.3.5. Now the vertices of the infinite side from one of the copies is a 4-connected set in that graph, but the graph does not contain any subdivision of a 4-typical graph, since it neither contains a path of length greater than 13 (and hence no subdivision of a $T_4(\mathcal{B})$ for some regular k -blueprint \mathcal{B}), nor a subdivision of a K_{4,\aleph_0} .

To solve this problem we introduce *generalised k -typical graphs*, where we ‘blow up’ some of the vertices of our k -typical graph to some finite tree, e.g. an edge in the previous example. This then will allow us to obtain the desired subdivisions for our characterisation.

Let G be a graph, $v \in V(G)$ be a vertex, T be a finite tree and $\gamma : N(v) \rightarrow V(T)$ be a map. We define the (v, T, γ) -blow-up of v in G as the operation where we delete v , add a vertex set $\{v\} \times V(T)$ disjointly and for each $w \in N(v)$ add the edge between w and $(v, \gamma(w))$. We call the resulting graph $G(v, T, \gamma)$.

Given blow-ups (v, T_v, γ_v) and (w, T_w, γ_w) in G , we can apply the blow-up of w in $G(v, T_v, \gamma_v)$ by replacing v in the preimage of γ_w by $(v, \gamma_v(w))$. We call this graph $G(v, T_v, \gamma_v)(w, T_w, \gamma_w)$. Note that no matter in which order we apply the blow-ups




Figure 4.3.5.: A graph with an infinite 4-connected set (marked by the cross vertices) containing no subdivision of a 4-typical graph.

we obtain the same graph, that is $G(v, T_v, \gamma_v)(w, T_w, \gamma_w) = G(w, T_w, \gamma_w)(v, T_v, \gamma_v)$. We analogously define for a set $O = \{(v, T_v, \gamma_v) \mid v \in W\}$ of blow-ups for some $W \subseteq V(G)$ the graph $G(O)$ obtained by successively applying all the blow-ups in O . Note that if W is infinite, then $G(O)$ is still well-defined, since each edge gets each of its endvertices modified at most once.

A *type-1 k -template* \mathcal{T}_1 is a triple (T, γ, c) consisting of a finite tree T , a map $\gamma : [0, k) \rightarrow V(T)$ and a node $c \in V(T)$ such that each node of degree 1 or 2 in T is either c or in the image of γ . Note that for each k there are only finitely many type-1 k -templates up to isomorphisms of the trees, since their trees have order at most $2k + 1$.

Let $\mathcal{T}_1 = (T, \gamma, c)$ be a type-1 k -template and let $O_1 := \{(z, T, \gamma) \mid z \in C(K_{k, \kappa})\}$. We call the graph $K_{k, \kappa}(\mathcal{T}_1) := K_{k, \kappa}(O_1)$ a *generalised $K_{k, \kappa}$* . The *core* $C(K_{k, \kappa}(\mathcal{T}_1))$ is the set $C(K_{k, \kappa}) \times \{c\}$, see Figure 4.3.6 for an example. Note that Figure 4.3.5 is also an example.

Similarly, with \mathcal{T}_1 as above, let $O'_1 := \{(z, T, \gamma_\alpha) \mid \alpha \in \text{cf } \kappa, z \in Z^\alpha\}$, where γ_α denotes the map defined by $y_i^\alpha \mapsto \gamma(i)$. The graph $\ell\text{-}K(k, \mathcal{K})(\mathcal{T}_1) := \ell\text{-}K(k, \mathcal{K})(O'_1)$ is a *generalised $\ell\text{-}K(k, \mathcal{K})$* . We call the vertex set $\cup Z \times \{c\}$ the *precore* of that graph.

Analogously, we obtain a *generalised $\ell\text{-}FK_{k, \kappa}(\mathcal{K})$* for any good κ -sequence \mathcal{K} as $\ell\text{-}FK_{k, \kappa}(\mathcal{K})(\mathcal{T}_1) := \ell\text{-}FK_{k, \kappa}(O'_1)$ with *core* $C(\ell\text{-}FK_{k, \kappa}(\mathcal{K})(\mathcal{T}_1)) := \cup Z \times \{c\}$.

A *type-2 k -template* \mathcal{T}_2 for a k -blueprint $\mathcal{B} = (B, D)$ is a set $\{(b, P_b, \gamma_b) \mid b \in V(B) \setminus D\}$ of blow-ups in B such that for all $b \in V(B) \setminus D$

- P_b is a path of length at most $k + 2$;




Figure 4.3.6.: Image of a generalised $K_{4,\kappa}$ on the left. The crosses represent the core. On the right is how we represent the same graph in a simplified way by labelling the vertices according to their adjacencies.

- the endnodes of P_b are called v_0^b and v_1^b ;
- P_b contains nodes v_\perp^b and v_\top^b ;
- the nodes $v_0^b, v_\perp^b, v_\top^b, v_1^b$ need not be distinct;
- if $v_0^n \neq v_\perp^b$, then $v_0^b v_\perp^b \in E(P_b)$ and if $v_1^n \neq v_\top^b$, then $v_1^b v_\top^b \in E(P_b)$;
- $\gamma_b(N(b)) \subseteq v_\perp^b P_b v_\top^b$;

We say \mathcal{T}_2 is *simple* if $v_0^b = v_\perp^b$ and $v_1^b = v_\top^b$. Note that for each k there are only finitely many type-2 k -templates, up to isomorphisms of the trees in the k -blueprints and the paths for the blow-ups.

Let $\mathcal{T}_2 = \{(b, T_b, \gamma_b) \mid b \in V(B) \setminus D\}$ be a type-2 k -template for a k -blueprint (B, D) . Then $O_2 := \{(b_n, T_b, \gamma_b^n) \mid n \in \mathbb{N}, b \in V(B) \setminus D\}$ is a set of blow-ups in $\mathfrak{N}(B/D)$, where γ_b^n is defined via

$$\gamma_b^n(v) = \begin{cases} \gamma_b(b') & \text{if } v = b'_n \text{ for } b' \in N(b); \\ v_\top^b & \text{if } v = b_{n+1}; \\ v_\perp^b & \text{if } n \geq 1 \text{ and } v = b_{n-1}. \end{cases}$$

Then $\mathfrak{N}(B/D)(\mathcal{T}_2) := \mathfrak{N}(B/D)(O_2)$ is a *generalised* $\mathfrak{N}(B/D)$.

Let $\mathcal{B} = (B, D, c)$ be a regular k -blueprint and $\mathcal{T}_2 = \{(b, T_b, \gamma_b) \mid b \in V(B) \setminus D\}$ be a type-2 k -template for (B, D) . We call $T_k(\mathcal{B})(\mathcal{T}_2) := T_k(\mathcal{B})(O_2)$ a *generalised* $T_k(\mathcal{B})$ with *core* $C(T_k(\mathcal{B})(\mathcal{T}_2)) := V(\mathfrak{N}_c) \times \{v_1^c\}$. For an example that generalises the graph of Figure 4.3.2, see Figure 4.3.7.




Figure 4.3.7.: Image of a generalised $T_4(P, \{d\}, c)$ for $P, \{d\}, c$ as in Figure 4.3.2.

In grey we represent the blow-up of P as given by some type-2 k -template. The crosses represent the core.

A *type-3 k -template* \mathcal{T}_3 for a singular k -blueprint $\mathcal{B} = (\ell, f, B, D, \sigma)$ is a tuple $(\mathcal{T}_1, \mathcal{T}_2)$ consisting of a type-1 $(\ell + f)$ -template \mathcal{T}_1 and a type-2 $(k - \ell - f)$ -template \mathcal{T}_2 . Note that for each k there are only finitely many type-3 k -templates up to isomorphisms as discussed above for \mathcal{T}_1 and \mathcal{T}_2 .

Let $\mathcal{T}_3 = (\mathcal{T}_1, \mathcal{T}_2)$ be a type-3 k -template with $\mathcal{T}_1 = (T, \gamma, c_1)$ for a singular k -blueprint $\mathcal{B} = (\ell, f, B, D, \sigma)$. Then for $(b_n, T_b, \gamma_b^n) \in O_2$ we extend γ_b^n to $\hat{\gamma}_b^n$ via

$$\hat{\gamma}_b^n(v) = \begin{cases} v_1^b & \text{if } v \in \{y_i^n \mid i \in [\ell + f, k]\} \text{ and } n \text{ even;} \\ v_0^b & \text{if } v \in \{y_i^n \mid i \in [\ell + f, k]\} \text{ and } n \text{ odd;} \\ \gamma_b^n(v) & \text{otherwise.} \end{cases}$$

Let $O'_2 := \{(b_n, T_b, \hat{\gamma}_b^n) \mid (b_n, T_b, \gamma_b^n) \in O_2\}$ denote the corresponding set of blow-ups in $T_k(\mathcal{B})$ and let O'_1 be for \mathcal{T}_1 as above. The graph $T_k(\mathcal{B})(\mathcal{T}_3) := T_k(\mathcal{B})(O'_1 \cup O'_2)$ is a *generalised $T_k(\mathcal{B})$* with *core* $C(T_k(\mathcal{B})(\mathcal{T}_3)) := \cup \mathcal{Z} \times \{c_1\}$. For an example that generalises the graph of Figure 4.3.4 see Figure 4.3.8.

We call the graph from which a generalised graph is obtained via this process its *parent*. As before, Remark 4.3.2 and its extensions extend to generalised k -typical graphs as well.

Remark 4.3.5. Every ℓ - $FK_{k,\kappa}(\mathcal{T}_1)(\mathcal{K})$, or $T_k(\mathcal{B})(\mathcal{T}_3)(\mathcal{K})$ respectively, for a singular k -blueprint \mathcal{B} , a type-1 k -template \mathcal{T}_2 , a type-2 k -template \mathcal{T}_3 and a good κ -sequence \mathcal{K} , contains a subdivision of ℓ - $FK_{k,\kappa}(\mathcal{T}_1)$, or $T_k(\mathcal{B})(\mathcal{T}_3)$ respectively.

A *generalised k -typical graph* is either $K_{k,\kappa}(\mathcal{T}_1)$, ℓ - $FK_{k,\kappa}(\mathcal{T}_1)$, $T_k(\mathcal{B})(\mathcal{T}_2)$ or $T_k(\mathcal{B})(\mathcal{T}_3)$ for any type-1 k -template \mathcal{T}_1 , any $\ell \in [0, k)$, any regular k -blueprint \mathcal{B} ,




Figure 4.3.8.: Image of a generalised $T_7(1, 2, P, \{d\}, \sigma)$ for P , $\{d\}$, σ as in Figure 4.3.4.

any type-2 k -template \mathcal{T}_2 for \mathcal{B} , any singular k -blueprint \mathcal{B}' and any type-3 k -template \mathcal{T}_3 for \mathcal{B}' . As with the k -typical graphs we obtain that this list is finite.

Corollary 4.3.6. *The core of a generalised k -typical graph is k -connected in that graph. \square*

4.3.3. Statement of the Main Theorem

Now that we introduced all k -typical and generalised k -typical graphs, let us give the full statement of our main theorem.

Theorem 4.3.7. *Let G be an infinite graph, let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and let $\kappa \leq |A|$ be an infinite cardinal. Then the following statements are equivalent.*

- (a) *There is a subset $A_1 \subseteq A$ with $|A_1| = \kappa$ such that A_1 is k -connected in G .*
- (b) *There is a subset $A_2 \subseteq A$ with $|A_2| = \kappa$ such that there is a k -typical graph which is a minor of G with finite branch sets and with A_2 along its core.*

(c) There is a subset $A_3 \subseteq A$ with $|A_3| = \kappa$ such that there G contains a subdivided generalised k -typical graph with A_3 as its core.

(d) There is no nested separation system $N \subseteq S_k(G)$ such that every part P of N can be separated from A by less than κ vertices.

Moreover, if these statements hold, we can choose $A_1 = A_2 = A_3$.

Note that for $A = V(G)$ we obtain the simple version as in Theorems 4.1.1 and 4.1.2 by forgetting the extra information about the core.

4.4. k -connected sets, minors and topological minors

In this section we will collect a few basic remarks and lemmas on k -connected sets and how they interact with minors and topological minors for future references. We omit some of the trivial proofs.

Remark 4.4.1. If $A \subseteq V(G)$ is k -connected in G , then any $A' \subseteq A$ with $|A'| \geq k$ is k -connected in G as well. \square

Lemma 4.4.2. If M is a minor of G and $A \subseteq V(M)$ is k -connected in M for some $k \in \mathbb{N}$, then any set $A' \subseteq V(G)$ with $|A'| \geq k$ consisting of at most one vertex of each branch set for the vertices of A is k -connected in G . \square

Lemma 4.4.3. For $k \in \mathbb{N}$, if G contains the subdivision of a generalised k -typical graph T with core A , then the parent of T is an fbs-minor with A along its core. \square

A helpful statement for the upcoming inductive constructions would be that for every vertex v of G , every large k -connected set in G contains a large subset which is $(k - 1)$ -connected in $G - v$. But while this is a true statement (cf. Corollary 4.8.2), an elementary proof of it seems to be elusive if v is not itself contained in the original k -connected set. The following lemma is a simplified version of that statement and has an elementary proof.

Lemma 4.4.4. Let $k \in \mathbb{N}$ and let $A \subseteq V(G)$ be infinite and k -connected in G . Then for any finite set $S \subseteq V(G)$ with $|S| < k$ there is a subset $A' \subseteq A$ with $|A'| = |A|$ such that A' is 1-connected in $G - S$.

Proof. Without loss of generality we may assume that A and S are disjoint. Take a sequence $(B_\alpha \mid \alpha \in |A|)$ of disjoint subsets of A with $|B_\alpha| = k$. For every $\alpha \in |A| \setminus \{0\}$ there is at least one path from B_0 to B_α disjoint from S . By the pigeonhole principle there is some $v \in B_0$ such that $|A|$ many of these paths start in v . Now let A' be the set of endvertices of these paths. \square

4.5. Structure within ends

This section studies the structure within an end of a graph.

In Subsection 4.5.1 we will extend to arbitrary infinite graphs a well-known result for locally finite graphs relating end degree with a certain sequence of minimal separators, making use of the combined end degree.

Subsection 4.5.2 is dedicated to the construction of a uniformly connecting structure between disjoint rays in a common end and vertices dominating that end.

4.5.1. End defining sequences and combined end degree

For an end $\omega \in \Omega(G)$ and a finite set $S \subseteq V(G)$ let $C(S, \omega)$ denote the unique component of $G - S$ that contains ω -rays. A sequence $(S_n \mid n \in \mathbb{N})$ of finite vertex sets of G is called an ω -defining sequence if for all $n, m \in \mathbb{N}$ with $n \neq m$ the following hold:

- $C(S_{n+1}, \omega) \subseteq C(S_n, \omega)$;
- $S_n \cap S_m \subseteq \text{Dom}(\omega)$; and
- $\bigcap \{C(S_n, \omega) \mid n \in \mathbb{N}\} = \emptyset$.

Note that for every ω -defining sequence $(S_n \mid n \in \mathbb{N})$ and every finite set $X \subseteq V(G)$ we can find an $N \in \mathbb{N}$ such that $X \subseteq G - C(S_N, \omega)$. Hence we shall also refer to the sets S_n in such a sequence as *separators*. Given $n, m \in \mathbb{N}$ with $n < m$, let $G[S_n, S_m]$ denote $G[(S_n \cup C(S_n, \omega)) \setminus C(S_m, \omega)]$, the *graph between the separators*.

For ends of locally finite graphs there is a characterisation of the end degree given by the existence of certain ω -defining sequences. The degree of an end ω is equal to $k \in \mathbb{N}$, if and only if k is the smallest integer such that there is an

ω -defining sequence of sets of size k , cf. [45, Lemma 3.4.2]. In this subsection we extend this characterisation to arbitrary graphs with respect to the combined degree. Recall the definition of the combined degree, $\Delta(\omega) := \deg(\omega) + \text{dom}(\omega)$.

In arbitrary graphs ω -defining sequences need not necessarily exist, e.g. in K_{\aleph_1} . We start by characterising the ends admitting such a sequence.

Lemma 4.5.1. *Let $\omega \in \Omega(G)$ be an end. Then there is an ω -defining sequence $(S_n \mid n \in \mathbb{N})$ if and only if $\Delta(\omega) \leq \aleph_0$.*

Proof. Note that for all finite $S \subseteq V(G)$, no $d \in \text{Dom}(\omega)$ can lie in a component $C \neq C(S, \omega)$ of $G - S$. Hence for every ω -defining sequence $(S_n \mid n \in \mathbb{N})$ and every $d \in \text{Dom}(\omega)$ there is an $N \in \mathbb{N}$ such that $d \in S_m$ for all $m \geq N$. Therefore, if $\text{dom}(\omega) > \aleph_0$, no ω -defining sequence can exist, since the union of the separators is at most countable. Moreover, note that for every ω -defining sequence every ω -ray meets infinitely many distinct separators. It follows that $\deg(\omega)$ is at most countable as well if an ω -defining sequence exist.

For the converse, suppose $\Delta(\omega) \leq \aleph_0$. Let $\{d_n \mid n < \text{dom}(\omega)\}$ be an enumeration of $\text{Dom}(\omega)$. Recall the definition of an ω -devouring ray from Chapter 3 as a ray that meets every ω -ray. Let $R = r_0 r_1 \dots$ be an ω -devouring ray, which exists by Lemma 3.1.1. We build our desired ω -defining sequence $(S_n \mid n \in \mathbb{N})$ inductively. Set $S_0 := \{r_0\}$. For $n \in \mathbb{N}$ suppose S_n is already constructed as desired. Take a maximal set \mathcal{P}_n of pairwise disjoint $N(S_n \setminus \text{Dom}(\omega)) - R$ paths in $C(S_n, \omega)$. Note that \mathcal{P}_n is finite since otherwise by the pigeonhole principle we would get a vertex $v \in S_n \setminus \text{Dom}(\omega)$ dominating ω . Furthermore, \mathcal{P}_n is not empty as $C(S_n, \omega)$ is connected. Define

$$\begin{aligned} S_{n+1} := & (S_n \cap \text{Dom}(\omega)) \cup \bigcup \mathcal{P}_n \\ & \cup \{r_m \mid m \text{ is minimal with } r_m \in C(S_n, \omega)\} \\ & \cup \{d_m \mid m \text{ is minimal with } d_m \in C(S_n, \omega)\}. \end{aligned}$$

By construction, $S_{n+1} \cap S_i$ contains only vertices dominating ω for $i \leq n$. Let P be any $S_n - C(S_{n+1}, \omega)$ path. We can extend P in $C(S_{n+1}, \omega)$ to an $S_n - R$ path. And since $C(S_{n+1}, \omega) \cap S_{n+1}$ is empty, we obtain $P \cap S_{n+1} \neq \emptyset$ by construction of S_{n+1} . Hence any $S_n - C(S_{n+1}, \omega)$ path meets S_{n+1} . Since for any vertex $v \in C(S_{n+1}, \omega) \setminus C(S_n, \omega)$ there is a path to $C(S_n, \omega) \cap C(S_{n+1}, \omega)$ in $C(S_{n+1}, \omega)$, this path would meet a vertex $w \in S_n$. This vertex would be a trivial $S_n -$

$C(S_{n+1}, \omega)$ path avoiding S_{n+1} , and hence contradicting the existence of such v . Hence $C(S_{n+1}, \omega) \subseteq C(S_n, \omega)$.

Suppose there is a vertex $v \in \bigcap \{C(S_n, \omega) \mid n \in \mathbb{N}\}$. By construction v is neither dominating ω nor is a vertex on R . Note that every $v-R$ path has to contain vertices from infinitely many S_n , hence it has to contain a vertex dominating ω . For each $d \in \text{Dom}(\omega)$ let P_d be either the vertex set of a $v-\text{Dom}(\omega)$ path containing d if it exists, or $P_d = \emptyset$ otherwise. If $X := \bigcup \{P_d \mid d \in \text{Dom}(\omega)\}$ is finite, we can find an $N \in \mathbb{N}$ such that $v \in X \subseteq G - C(S_N, \omega)$, a contradiction. Otherwise apply Lemma 4.2.3 to $X \cap \text{Dom}(\omega)$ in $G[X]$. Note that in $G[X]$ all vertices of $X \cap \text{Dom}(\omega)$ have degree 1 in $G[X]$. Furthermore, we know that $V(R) \cap X \subseteq \text{Dom}(\omega)$, since no P_d contains a vertex of R as an internal vertex. But then the centre of a star would be a vertex dominating ω in $X \setminus \text{Dom}(\omega)$ and the spine of a comb would contain an ω -ray disjoint to R as a tail, again a contradiction. \square

In the proof of the end-degree characterisation via ω -defining sequences we shall need the following fact regarding the relationship of $\text{deg}(\omega)$ and $\text{dom}(\omega)$.

Lemma 4.5.2. *If $\text{deg}(\omega)$ is uncountable for $\omega \in \Omega(G)$, then $\text{dom}(\omega)$ is infinite.*

Proof. Suppose for a contradiction that $\text{dom}(\omega) < \aleph_0$. For $G' := G - \text{Dom}(\omega)$ let \mathcal{R} be a set of disjoint $\omega \upharpoonright G'$ -rays of size \aleph_1 , which exist by Remark 4.2.5. Let T be a transversal of $\{V(R) \mid R \in \mathcal{R}\}$. Applying Lemma 4.2.3 to T yields a subdivided star with centre d and uncountably many leaves in T . Now $d \notin \text{Dom}(\omega)$ dominates $\omega \upharpoonright G'$ in G' and hence ω in G by Remark 4.2.5, a contradiction. \square

Let $\omega \in \Omega(G)$ be an end with $\text{dom}(\omega) = 0$, $(S_n \mid n \in \mathbb{N})$ be an ω -defining sequence and \mathcal{R} be a set of disjoint ω -rays. We call $((S_n \mid n \in \mathbb{N}), \mathcal{R})$ a *degree witnessing pair* for ω , if for all $n \in \mathbb{N}$ and for each $s \in S_n$ there is a ray $R \in \mathcal{R}$ containing s and every ray $R \in \mathcal{R}$ meets S_n at most once for every $n \in \mathbb{N}$. Note that this definition only makes sense for undominated ends, since a ray that contains a dominating vertex meets eventually all separators not only in that vertex.

Lemma 4.5.3. *Let $\omega \in \Omega(G)$ be an end with $\text{dom}(\omega) = 0$. Then there is a degree witnessing pair $((S_n \mid n \in \mathbb{N}), \mathcal{R})$.*

Proof. By Lemmas 4.5.1 and 4.5.2 there is an ω -defining sequence $(S'_n \mid n \in \mathbb{N})$. Since ω is undominated, the separators are pairwise disjoint.

We want to construct an ω -defining sequence $(S_n \mid n \in \mathbb{N})$ with the property, that for all $n \in \mathbb{N}$ and for all $m > n$ there are $|S_n|$ many $S_n - S_m$ paths in $G[S_n, S_m]$.

Let S_0 be an $S'_0 - S'_{f(0)}$ separator for some $f(0) \in \mathbb{N}$ which is of minimum order among all candidates separating S'_0 from S'_m for any $m \in \mathbb{N}$. Suppose we already constructed the sequence up to S_n . Let S_{n+1} be an $S'_{f(n)+1} - S'_{f(n+1)}$ separator for some $f(n+1) > f(n) + 1$ which is of minimum order among all candidates separating $S'_{f(n)+1}$ and S'_m for any $m > f(n) + 1$.

Note that S_m and S_n are disjoint for all $m, n \in \mathbb{N}$ with $n \neq m$ and that $C(S_{n+1}, \omega) \subseteq C(S'_{f(n)+1}, \omega)$ for all $n \in \mathbb{N}$. Hence $(S_n \mid n \in \mathbb{N})$ is an ω -defining sequence. Moreover, note that $|S_n| \leq |S_{n+1}|$ for all $n \in \mathbb{N}$, since S_{n+1} would have been a candidate for S_n as well. In particular, there is no $S_n - S_{n+1}$ separator S of order less than $|S_n|$ for every $n \in \mathbb{N}$, since this would also have been a candidate for S_n . Hence by Theorem 4.2.1 there is a set of $|S_n|$ many disjoint $S_n - S_{n+1}$ paths \mathcal{P}_n in $G[S_n, S_{n+1}]$.

Now the union $\bigcup\{\bigcup \mathcal{P}_n \mid n \in \mathbb{N}\}$ is by construction a union of a set \mathcal{R} of rays, since the union of the paths in \mathcal{P}_n intersect the union of the paths in \mathcal{P}_m in precisely S_{n+1} if $m = n + 1$ and are disjoint if $m > n + 1$. These rays are necessarily ω -rays, meet every separator at most once and every $s \in S_n$ is contained in one of them, proving that $((S_n \mid n \in \mathbb{N}), \mathcal{R})$ is a degree witnessing pair for ω . \square

Corollary 4.5.4. *Let $k \in \mathbb{N}$ and let $\omega \in \Omega(G)$ with $\text{dom}(\omega) = 0$. Then $\text{deg}(\omega) \geq k$ if and only if for every ω -defining sequence $(S_n \mid n \in \mathbb{N})$ the sets S_n eventually have size at least k .*

Proof. Suppose $\text{deg}(\omega) \geq k$. Let $(S_n \mid n \in \mathbb{N})$ be any ω -defining sequence. Then each ray out of a set of k disjoint ω -rays has to go through eventually all S_n . For the other direction take a degree witnessing pair $((S_n \mid n \in \mathbb{N}), \mathcal{R})$. Now $|\mathcal{R}| \geq k$, since eventually all S_n have size at least k . \square

Corollary 4.5.5. *Let $k \in \mathbb{N}$ and let $\omega \in \Omega(G)$ with $\text{dom}(\omega) = 0$. Then $\text{deg}(\omega) = k$ if and only if k is the smallest integer such that there is an ω -defining sequence $(S_n \mid n \in \mathbb{N})$ with $|S_n| = k$ for all $n \in \mathbb{N}$.* \square

We can easily lift these results to ends dominated by finitely many vertices with the following observation based on Remark 4.2.5.

Remarks 4.5.6. Suppose $\text{dom}(\omega) < \aleph_0$. Let G' denote $G - \text{Dom}(\omega)$.

- (a) For every $\omega \upharpoonright G'$ -defining sequence $(S'_n \mid n \in \mathbb{N})$ of G' there is an ω -defining sequence $(S_n \mid n \in \mathbb{N})$ of G with $S'_n = S_n \setminus \text{Dom}(\omega)$ for all $n \in \mathbb{N}$.
- (b) For every ω -defining sequence $(S_n \mid n \in \mathbb{N})$ of G there is an $\omega \upharpoonright G'$ -defining sequence $(S'_n \mid n \in \mathbb{N})$ of G' with $S'_n = S_n \setminus \text{Dom}(\omega)$ for all $n \in \mathbb{N}$. \square

Corollary 4.5.7. *Let $k \in \mathbb{N}$ and let $\omega \in \Omega(G)$. Then $\Delta(\omega) \geq k$ if and only if for every ω -defining sequence $(S_n \mid n \in \mathbb{N})$ the sets S_n eventually have size at least k .*

Proof. As noted before, each vertex dominating ω has to be in eventually all sets of an ω -defining sequence.

Suppose $\Delta(\omega) \geq k$. If $\text{dom}(\omega) \geq \aleph_1$, then there is no ω -defining sequence and there is nothing to show. If $\text{dom}(\omega) = \aleph_0$, then the sets of any ω -defining sequence eventually have all size at least k . If $\text{dom}(\omega) < \aleph_0$, we can delete $\text{Dom}(\omega)$ and apply Corollary 4.5.4 to $G - \text{Dom}(\omega)$ with $k' = \text{deg}(\omega)$. With Remark 4.5.6(b) the claim follows.

If $\Delta(\omega) < k$, we can delete $\text{Dom}(\omega)$ and apply Corollary 4.5.4 with $k' = \text{deg}(\omega)$. With Remark 4.5.6(a) the claim follows. \square

Corollary 4.5.8. *Let $k \in \mathbb{N}$ and let $\omega \in \Omega(G)$. Then $\Delta(\omega) = k$ if and only if k is the smallest integer such that there is an ω -defining sequence $(S_n \mid n \in \mathbb{N})$ with $|S_n| = k$ for all $n \in \mathbb{N}$.*

Proof. As before, we delete $\text{Dom}(\omega)$ and apply Corollary 4.5.5 with $k' = \text{deg}(\omega)$ and Remark 4.5.6. \square

Finally, we state more remarks on the relationship between $\text{deg}(\omega)$ and $\text{dom}(\omega)$ similar to Lemma 4.5.2 without giving the proof.

Remarks 4.5.9. Let κ_1, κ_2 be infinite cardinals and let $k_1, k_2 \in \mathbb{N}$.

- (1) If $\text{dom}(\omega)$ is infinite, then so is $\text{deg}(\omega)$ for every $\omega \in \Omega(G)$.
- (2) If $\Delta(\omega)$ is uncountable, then both $\text{deg}(\omega)$ and $\text{dom}(\omega)$ are infinite for every $\omega \in \Omega(G)$.
- (3) There is a graph with an end ω' such that $\text{deg}(\omega') = \kappa_1$ and $\text{dom}(\omega') = \kappa_2$.
- (4) There is a graph with an end ω' such that $\text{deg}(\omega') = k_1$ and $\text{dom}(\omega') = k_2$.
- (5) There is a graph with an end ω' such that $\text{deg}(\omega') = \aleph_0$ and $\text{dom}(\omega') = k_2$.

4.5.2. Constructing uniformly connected rays

Let $\omega \in \Omega(G)$ be an end of G and let I, J be two disjoint finite sets with $1 \leq |I| \leq \deg(\omega)$ and $0 \leq |J| \leq \text{dom}(\omega)$. Let $\mathcal{R} = (R_i \mid i \in I)$ be a family of disjoint ω -rays and let $\mathcal{D} = (d_j \in \text{Dom}(\omega) \mid j \in J)$ be a family of distinct vertices disjoint from $\bigcup \mathcal{R}$. Let T be a tree on $I \cup J$ such that J is a set of leaves of T . Let $W := \bigcup \mathcal{R} \cup \mathcal{D}$ and $k := |I \cup J|$. We call a finite subgraph $\Gamma \subseteq G$ a (T, \mathcal{T}_2) -connection, if if \mathcal{T}_2 is a simple type-2 k -template for (T, J) , and there is a set \mathcal{P} of internally disjoint W - W paths such that $\Gamma \subseteq \bigcup \mathcal{R} \cup \bigcup \mathcal{P}$ and Γ is isomorphic to a subdivision of $T(\mathcal{T}_2)$. Moreover, the subdivision of $v_{\perp}^i P_i v_{\top}^i$ is the segment $R_i \cap \Gamma$ for all $i \in I$ such that v_{\top}^i corresponds to the top vertex of that segment. Then $(\mathcal{R}, \mathcal{D})$ is called (T, \mathcal{T}_2) -connected if for every finite $X \subseteq V(G) \setminus \mathcal{D}$ there is a (T, \mathcal{T}_2) -connection avoiding X .

Lemma 4.5.10. *Let $\omega \in \Omega(G)$, let $\mathcal{R} = (R_i \mid i \in I)$ be a finite family of disjoint ω -rays with $|I| \geq 1$ and let $\mathcal{D} = (d_j \in \text{Dom}(\omega) \mid j \in J)$ be a finite family of distinct vertices disjoint from $\bigcup \mathcal{R}$ with $I \cap J = \emptyset$. Then there is a tree T on $I \cup J$ and a simple type-2 $|I \cup J|$ -template \mathcal{T} for (T, J) such that $(\mathcal{R}, \mathcal{D})$ is (T, \mathcal{T}) -connected.*

Proof. Let $X \subseteq V(G) \setminus \mathcal{D}$ be any finite set. We extend X to a finite superset X' such that $R_i \cap X'$ is an initial segment of R_i for each $i \in I$, and such that $\mathcal{D} \subseteq X'$. As all rays in \mathcal{R} are ω -rays, we can find finitely many $\bigcup \mathcal{R}$ - $\bigcup \mathcal{R}$ paths avoiding X' which are internally disjoint such that their union with $\bigcup \mathcal{R}$ is a connected subgraph of G . Moreover it is possible to do this with a set \mathcal{P} of $|I| - 1$ many such paths in a tree-like way, i.e. contracting a large enough finite segment avoiding X' of each ray in \mathcal{R} and deleting the rest yields a subdivision Γ'_X of a tree on I whose edges correspond to the paths in \mathcal{P} . For each vertex d_j we can moreover find a d_j - $\bigcup \mathcal{R}$ path avoiding $V(\Gamma'_X) \cup X' \setminus \{d_j\}$ and all paths we fixed so far. This yields a tree T_X on $I \cup J$ and a simple type-2 k -template \mathcal{T}_X for (T_X, J) such that J is a set of leaves and a (T_X, \mathcal{T}_X) -connection Γ_X avoiding X .

Now we iteratively apply this construction to find a family $(\Gamma_i \mid i \in \mathbb{N})$ of (T_i, \mathcal{T}_i) -connections such that $\Gamma_m - \mathcal{D}$ and $\Gamma_n - \mathcal{D}$ are disjoint for all $m, n \in \mathbb{N}$ with $m \neq n$. By the pigeonhole principle we now find a tree T on $I \cup J$, a type-2 $|I \cup J|$ -template \mathcal{T} and an infinite subset $N \subseteq \mathbb{N}$ such that $(T_n, \mathcal{T}_n) = (T, \mathcal{T})$ for all $n \in N$.

Now for each finite set $X \subseteq V(G) \setminus \mathcal{D}$ there is an $n \in N$ such that Γ_n and X are disjoint, hence $(\mathcal{R}, \mathcal{D})$ is (T, \mathcal{T}) -connected. \square

Corollary 4.5.11. *Let $\omega \in \Omega(G)$, let $\mathcal{R} = (R_i \mid i \in I)$ be a finite family of disjoint ω -rays with $|I| \geq 1$ and let $\mathcal{D} = (d_j \in \text{Dom}(\omega) \mid j \in J)$ be a finite family of distinct vertices disjoint from $\bigcup \mathcal{R}$ with $I \cap J = \emptyset$. Then there is a tree T such that G contains a subdivision of a generalised $\mathfrak{N}(T/J)$.*

Proof. By Lemma 4.5.10 there is a tree T and a simple type-2 $|I \cup J|$ -template \mathcal{T} such that $(\mathcal{R}, \mathcal{D})$ is (T, \mathcal{T}) -connected.

For a (T, \mathcal{T}) -connection Γ let r_Γ^i denote the top vertex of $R_i \cap \Gamma$. Then by $\bar{\Gamma}$ we denote the union of Γ with the initial segments $R_i r_\Gamma^i$ for all $i \in I$.

Let $(\Gamma_i \mid i \in \mathbb{N})$ be a family of (T, \mathcal{T}) -connections such that for all $m < n$ the graphs $\Gamma_n - \mathcal{D}$ and $\bar{\Gamma}_m - \mathcal{D}$ are disjoint. Then $H = \bigcup \{\Gamma_n \mid n \in \mathbb{N}\} \cup \bigcup \mathcal{R}$ is the desired subdivision. \square

Finally, this result can be lifted to the minor setting by Lemma 4.4.3.

Corollary 4.5.12. *Let $\omega \in \Omega(G)$, let $\mathcal{R} = (R_i \mid i \in I)$ be a finite family of disjoint ω -rays with $|I| \geq 1$ and let $\mathcal{D} = (d_j \in \text{Dom}(\omega) \mid j \in J)$ be a finite family of distinct vertices disjoint from $\bigcup \mathcal{R}$ with $I \cap J = \emptyset$. Then there is a tree T such that G contains $\mathfrak{N}(T/J)$ as an fbs-minor. \square*

4.6. Minors for regular cardinalities

This section is dedicated to prove the equivalence of (a), (b) and (c) of Theorem 4.3.7 for regular cardinals κ .

4.6.1. Complete bipartite minors

In this subsection we construct the complete bipartite graph $K_{k,\kappa}$ as the desired minor (and a generalised version as the desired subdivision), if possible. The ideas of this construction differ significantly from Halin's construction [29, Thm. 9.1] of a subdivision of $K_{k,\kappa}$ in a k -connected graph of uncountable and regular order κ .

Lemma 4.6.1. *Let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and k -connected in G and let $\kappa \leq |A|$ be a regular cardinal. If*

- either κ is uncountable;
- or there is no end in the closure of A ;
- or there is an end ω in the closure of A with $\text{dom}(\omega) \geq k$;

then there is a subset $A' \subseteq A$ with $|A'| = \kappa$ such that $K_{k,\kappa}$ is an fbs-minor of G with A' along its core.

Moreover, the branch sets for the vertices of the finite side of $K_{k,\kappa}$ are singletons.

Proof. We iteratively construct a sequence of subgraphs H_i for $i \in [0, k]$ witnessing that $K_{i,\kappa}$ is a minor of G . Furthermore, we incorporate that the branch sets for the vertices of the finite side of $K_{i,\kappa}$ are singletons $\{v_j \mid j \in [0, i]\}$ and the branch sets for the vertices of the infinite side induce finite trees on H_i each containing a vertex of A . Moreover, for $i < k - 1$ we will guarantee the existence of a subset $A_i \subseteq A$ with $|A_i| = \kappa$ which is 1-connected in $G_i := G - \{v_j \mid j \in [0, i]\}$ and such that each vertex of A_i is contained in a branch set of H_i and each branch set of H_i contains precisely one vertex of A_i .

Set $G_0 := G$, $A_0 := A$ and $H_0 = G[A]$. For any $i \in [0, k)$ we inductively apply Lemma 4.2.3 (and in the third case also Remark 4.2.4) to A_i in G_i to find a subdivided star S_i with centre v_i and κ many leaves $L_i \subseteq A_i$. Without loss of generality we can assume $v_i \notin V(H_i)$, since otherwise we could just remove the branch set containing v_i and from A_i the vertex contained in that branch set. If $i < k - 1$, then by Lemma 4.4.4 we find a subset $L'_i \subseteq L_i$ with $|L'_i| = \kappa$ which is 1-connected in G_{i+1} . In the case that $i = k - 1$ let $L'_{k-1} = L_{k-1}$.

Again for any $i \in [0, k)$, we first remove from H_i every branch set which corresponds to a vertex of the infinite side of $K_{i,\kappa}$ and does not contain a vertex of L'_i . Now each path in S_i from a neighbour of v_i to L'_i eventually hits a vertex of one of the finite trees induced by one of the remaining branch sets of H_i . Since all these paths are disjoint, only finitely many of them meet the same branch set first. Thus κ many different of the remaining branch sets are met by those paths first. To get H_{i+1} we do the following. First we add $\{v_i\}$ as a new branch set. Then each of the κ many branch sets reached first as described above we extend by the path segment between v_i and that branch set of precisely one of those paths. Finally, we delete all remaining branch sets not connected to $\{v_i\}$. With $A_{i+1} := L'_i \cap V(H_{i+1})$ we now have all the desired properties.

Finally, setting $H := H_k$ and $A' := A_k$ finishes the construction. \square

Let H be an inflated subgraph witnessing that $K_{k,\kappa} = K([0, k], Z)$ is an fbs-minor of G with A along Z for some $A \subseteq V(G)$ where each branch set of $x \in [0, k)$ is a singleton. Given a type-1 k -template $\mathcal{T}_1 = (T, \gamma, c)$ we say H is \mathcal{T}_1 -regular if for each $z \in Z$:

- there is an isomorphism $\varphi_z : T'_z \rightarrow T_z$ between a subdivision T'_z of T and the finite tree $T_z = H[\mathfrak{B}(z)]$;
- $x\varphi_z(\gamma(x)) \in E(H)$ for each $x \in [0, k)$; and
- $A \cap \mathfrak{B}(z) = \{\varphi_z(c)\}$.

We say G contains $K_{k,\kappa}$ as a \mathcal{T}_1 -regular fbs-minor with A along Z if there is such a \mathcal{T}_1 -regular H .

Lemma 4.6.2. *Let $k \in \mathbb{N}$ and κ be a regular cardinal. If $K_{k,\kappa}$ is an fbs-minor of G with A' along its core where each branch set of $x \in [0, k)$ is a singleton, then there is type-1 k -template \mathcal{T}_1 and $A'' \subseteq A'$ with $|A''| = \kappa$ such that G contains $K_{k,\kappa}$ as a \mathcal{T}_1 -regular fbs-minor with A'' along its core.*

Proof. Let H be the inflated subgraph witnessing that $K_{k,\kappa}$ is an fbs-minor as in the statement. Let x also denote the vertex of G in the branch set $\mathfrak{B}(x)$ of $x \in [0, k)$. Let $v_x^z \in \mathfrak{B}(z)$ denote the unique endvertex in $\mathfrak{B}(z)$ of the edge corresponding to $xz \in E(M)$ (cf. Section 4.2). Let T_z denote a subtree of $H[\mathfrak{B}(z)]$ containing $B_z = \{v_x^z \mid x \in [0, k)\} \cup \{a_z\}$ for the unique vertex $a_z \in A \cap \mathfrak{B}(z)$. Without loss of generality assume that each leaf of T_z is in B_z . By suppressing each degree 2 node of T_z that is not in B_z , we obtain a tree suitable for a type-1 k -template where a_z is the node in the third component of the template.

By applying the pigeonhole principle multiple times there is a tree T such that there exist an isomorphism $\varphi_z : T'_z \rightarrow T_z$ for a subdivision T'_z of T for all $z \in Z'$ for some $Z' \subseteq Z$ with $|Z'| = \kappa$, such that $\{\varphi_z(v_x^z) \mid z \in Z'\}$ is a singleton $\{t_x\}$ for all $x \in [0, k)$ as well as $\{\varphi_z(a_z) \mid z \in Z'\}$ is a singleton $\{c\}$.

Therefore with $\gamma : [0, k) \rightarrow V(T)$ defined by $x \mapsto t_x$ and c defined as above, we obtain a type-1 k -template $\mathcal{T}_1 := (T, \gamma, c)$ such that the subgraph H' of H where we delete each branch set for $z \in Z \setminus Z'$ is \mathcal{T}_1 -regular. \square

Hence, we also obtain a subdivision of a generalised $K_{k,\kappa}$.

Corollary 4.6.3. *In the situation of Lemma 4.6.1, there is $A'' \subseteq A'$ with $|A''| = \kappa$ such that G contains a subdivision of a generalised $K_{k,\kappa}$ with core A'' . \square*

4.6.2. Minors for regular k -blueprints

In this subsection we construct the k -typical minors for regular k -blueprints, if possible. While these graphs are essentially the same minors given by Oporowski, Oxley and Thomas [37, Thm. 5.2], we give our own independent proof based on the existence of an infinite k -connected set instead of the graph being essentially k -connected.

The first lemma constructs such a graph along some end of high combined degree.

Lemma 4.6.4. *Let $\omega \in \Omega(G)$ be an end of G with $\Delta(\omega) \geq k \in \mathbb{N}$. Let $A \subseteq V(G)$ be a set with ω in its closure. Then there is a countable subset $A' \subseteq A$ and a regular k -blueprint \mathcal{B} such that G contains a subdivision of a generalised $T_k(\mathcal{B})$ with core A' .*

Proof. Let I and J be any two disjoint sets with $|I| + |J| = k$, $1 \leq |I| \leq \deg(\omega)$ and $|J| \leq \text{dom}(\omega)$. Let $\mathcal{R} = (R_i \mid i \in I)$ be a family of disjoint ω -rays and $\mathcal{D} = (d_j \in \text{Dom}(\omega) \mid j \in J)$ be a family of distinct vertices disjoint from $\bigcup \mathcal{R}$. Applying Lemma 4.5.10 yields a tree B on $I \cup J$ and a type-2 k -template \mathcal{T} for (B, J) such that $(\mathcal{R}, \mathcal{D})$ is (B, \mathcal{T}) -connected. Let $(\Gamma_i \mid i \in \mathbb{N})$ denote the family of (B, \mathcal{T}) -connections as in the proof of Lemma 4.5.10. Moreover, there is an infinite set of disjoint $A - \bigcup \mathcal{R}$ paths by Theorem 4.2.2 since ω is in the closure of A . Now any infinite set of disjoint $A - \bigcup \mathcal{R}$ paths has infinitely many endvertices on one ray R_c for some $c \in I$. Let A'' denote the endvertices in A of such an infinite path system. Next we extend for infinitely many Γ_i the segment of R_c that it contains so that it has the endvertex of such an $A'' - R_c$ path as its top vertex and add that segment together with the path to Γ_i , while keeping them disjoint but for \mathcal{D} . Let A' denote the set of those endvertices of the paths in A'' we used to extend Γ_i for those infinitely many $i \in \mathbb{N}$. Finally, we modify the type-2 k -template accordingly. We obtain the subdivision of the generalised $T_k(B, J, c)$ as in the proof of Corollary 4.5.11. \square

The following lemma allows us to apply Lemma 4.6.4 when Lemma 4.6.1 is not applicable.

Lemma 4.6.5. *Let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and k -connected in G and let $\omega \in \Omega(G)$ be an end in the closure of A . Then $\Delta(\omega) \geq k$.*

Proof. We may assume that $\Delta(\omega)$ is finite since otherwise there is nothing to show. Hence without loss of generality A does not contain any vertices dominating ω . Let $(S_n \mid n \in \mathbb{N})$ be an ω -defining sequence, which exists by Lemma 4.5.1. Take $N \in \mathbb{N}$ such that there is a set $B \subseteq A \setminus C(S_N, \omega)$ of size k . For every $n > N$ let $C_n \subseteq A \cap C(S_n, \omega)$ be a set of size k , which exists since ω is in the closure of A . Since A is k -connected in G , there are k disjoint B - C_n paths in G , each of which contains at least one vertex of S_n . Hence for all $n > N$ we have $|S_n| \geq k$ and by Corollary 4.5.7 we have $\Delta(\omega) \geq k$. \square

We close this subsection with a corollary that is not needed in this chapter, but provides a converse for Lemma 4.6.5 as an interesting observation.

Corollary 4.6.6. *Let $\omega \in \Omega(G)$ be an end of G with $\Delta(\omega) \geq k \in \mathbb{N}$. Then every subset $A \subseteq V(G)$ with ω in the closure of A contains a countable subset $A' \subseteq A$ which is k -connected in G .*

Proof. By Lemma 4.6.4 we obtain a subdivision of a generalised $T_k(\mathcal{B})$ with core A' for some $A' \subseteq A$ in G for a regular k -blueprint \mathcal{B} . Corollary 4.3.6 yields the claim. \square

4.6.3. Characterisation for regular cardinals

Now we have developed all the necessary tools to prove the minor and topological minor part of the characterisation in Theorem 4.3.7 for regular cardinals.

Theorem 4.6.7. *Let G be a graph, let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and let $\kappa \leq |A|$ be a regular cardinal. Then the following statements are equivalent.*

- (a) *There is a subset $A_1 \subseteq A$ with $|A_1| = \kappa$ such that A_1 is k -connected in G .*
- (b) *There is a subset $A_2 \subseteq A$ with $|A_2| = \kappa$ such that*
 - *either $K_{k,\kappa}$ is an fbs-minor of G with A_2 along its core;*

- or $T_k(\mathcal{B})$ is an fbs-minor of G with A_2 along its core for some regular k -blueprint \mathcal{B} .

(c) There is a subset $A_3 \subseteq A$ with $|A_3| = \kappa$ such that

- either G contains a subdivision of a generalised $K_{k,\kappa}$ with core A_3 ;
- or G contains the subdivision of a generalised $T_k(\mathcal{B})$ with core A_3 for some regular k -blueprint \mathcal{B} .

Moreover, if these statements hold, we can choose $A_1 = A_2 = A_3$.

Proof. If (b) holds, then A_2 is k -connected in G by Lemma 4.4.2 with Lemma 4.3.4.

If (a) holds, then we can find a subset $A_3 \subseteq A_1$ with $|A_3| = \kappa$ yielding (c) by either Lemma 4.6.1 and Corollary 4.6.3 or by Lemma 4.6.5 and Lemma 4.6.4.

If (c) holds, then so does (b) by Lemma 4.4.3 with $A_2 := A_3$. Moreover, A_3 is a candidate for both A_2 and A_1 . \square

4.7. Minors for singular cardinalities

In this section we will prove the equivalence of (a), (b) and (c) of Theorem 4.3.7 for singular cardinals κ .

4.7.1. Cofinal sequence of regular bipartite minors with disjoint cores.

In this subsection, given a k -connected set A of size κ , we will construct an ℓ - $K(k, \mathcal{K})$ minor in G for some suitable $\ell \in [0, k]$ and good κ -sequence \mathcal{K} with a suitable subset of A along its precore. This minor is needed as an ingredient for any of the possible k -typical graphs but the $K_{k,\kappa}$ (which we obtain from the following lemma if $\ell = k$). Let $\mathcal{A} = (A^\alpha \subseteq A \mid \alpha \in \text{cf } \kappa)$ be a family of disjoint subsets of A . We say that G contains ℓ - $K(k, \mathcal{K})$ as an fbs-minor with \mathcal{A} along its precore \mathcal{Z} if the map mapping each vertex of the inflated subgraph to its branch set induces a bijection between A^α and Z^α for all $\alpha \in \text{cf } \kappa$.

Lemma 4.7.1. *Let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and k -connected in G and let $\kappa \leq |A|$ be a singular cardinal. Then there is an $\ell \in [0, k]$, a good κ -sequence*

$\mathcal{K} = (\kappa_\alpha < \kappa \mid \alpha \in \text{cf } \kappa)$, and a family $\mathcal{A} = (A^\alpha \subseteq A \mid \alpha \in \text{cf } \kappa)$ of pairwise disjoint subsets of A with $|A^\alpha| = \kappa_\alpha$ such that G contains ℓ - $K(k, \mathcal{K})$ as an fbs-minor with \mathcal{A} along \mathcal{Z} . Moreover, the branch sets for the vertices in $\bigcup \mathcal{Y}$ are singletons.

Proof. We start with any good κ -sequence $\mathcal{K} = (\kappa_\alpha < \kappa \mid \alpha \in \text{cf } \kappa)$. We construct the desired inflated subgraph by iteratively applying Lemma 4.6.1.

For $\alpha \in \text{cf } \kappa$ suppose we have already constructed for each $\beta < \alpha$ an inflated subgraph H^β witnessing that K_{k, κ_β} is an fbs-minor of G with some $A^\beta \subseteq A$ along its core. Furthermore, suppose that the branch sets of the vertices of the finite side are singletons and the branch sets of the vertices of the infinite side are disjoint to all branch sets of H^γ for all $\gamma < \beta$. We apply Lemma 4.6.1 for κ_α to any set $A' \subseteq A \setminus \bigcup_{\beta < \alpha} A^\beta$ of size κ_α to obtain an inflated subgraph for K_{k, κ_α} with the properties as stated in that lemma. If any branch set for a vertex of the infinite side meets any branch set we have constructed so far, we delete it. Since κ_α is regular and $\kappa_\alpha > \text{cf } \kappa$, the union of all inflated subgraphs we constructed so far has order less than κ_α . We obtain that the new inflated subgraph (after the deletions) still witnesses that K_{k, κ_α} is an fbs-minor of G with some $A^\alpha \subseteq A'$ along its core. If a branch set for the finite side meets any branch set of a vertex for the infinite side for some $\beta < \alpha$, we delete that branch set and modify A^β accordingly. As the union of all branch sets for the finite side we will construct in this process has cardinality $\text{cf } \kappa$, each A^β will lose at most $\text{cf } \kappa < \kappa_\beta$ many elements, hence will remain at size κ_β for all $\beta \in \text{cf } \kappa$. We denote the sequence $(A^\alpha \mid \alpha \in \text{cf } \kappa)$ with \mathcal{A} .

By Lemma 4.2.6 there is an $\ell \leq k$ and an $I \subseteq \text{cf } \kappa$ with $|I| = \text{cf } \kappa$ such that H^α and H^β have precisely ℓ branch sets for the vertices of the finite side in common for all $\alpha, \beta \in I$. Hence relabelling the subsequences $\mathcal{K} \upharpoonright I$ and $\mathcal{A} \upharpoonright I$ to $\overline{\mathcal{K}} \upharpoonright I$ and $\overline{\mathcal{A}} \upharpoonright I$ respectively as discussed in Section 4.3 yields the claim, where the union of the respective subgraphs H^α is the witnessing inflated subgraph. \square

4.7.2. Frayed complete bipartite minors

In this subsection we will construct a frayed complete bipartite minor, if possible. We shall use an increasing amount of fixed notation in this subsection based on Lemma 4.7.1, which we will fix as we continue our construction.

Situation 4.7.2. Let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and k -connected in G and let $\kappa \leq |A|$ be a singular cardinal. Let $\ell \leq k$ and let

- $\mathcal{K} = (\kappa_\alpha < \kappa \mid \alpha \in \text{cf } \kappa)$ be a good κ -sequence; and
- $\mathcal{A} = (A^\alpha \mid \alpha \in \text{cf } \kappa)$ be a family of pairwise disjoint subsets of A with $|A^\alpha| = \kappa_\alpha$.

Let H be an inflated subgraph witnessing that G contains ℓ - $K(k, \mathcal{K})$ as an fbs-minor with \mathcal{A} along \mathcal{Z} as in Lemma 4.7.1. To simplify our notation, we denote the unique vertex of H in a branch set of y_i^α also by y_i^α for all $\alpha \in \text{cf } \kappa$ and $i < k$. Similarly, we denote the set $\{y_i^\alpha \in V(H) \mid i \in [0, k]\}$ also with Y^α for all $\alpha \in \text{cf } \kappa$, and denote the family $(Y^\alpha \subseteq V(H) \mid \alpha \in \text{cf } \kappa)$ with \mathcal{Y} . Moreover, let H^α denote the subgraph of H witnessing that $K(Y^\alpha, Z^\alpha)$ is an fbs-minor of G with A^α along Z^α . Finally, let $D_\ell = \{y_i \mid i \in [0, \ell]\} = \bigcap \{V(H^\alpha) \mid \alpha \in \text{cf } \kappa\}$ denote the set of degenerate vertices of ℓ - $K(k, \mathcal{K})$.

For a set $U \subseteq V(G)$ and $\alpha \in \text{cf } \kappa$, we define a Y^α - U bundle P^α to be the union $\bigcup \{P_i^\alpha \mid i \in [0, k]\}$ of k disjoint paths, where $P_i^\alpha \subseteq G$ is a (possibly trivial) Y^α - U path starting in $y_i^\alpha \in Y^\alpha$ and ending in some $u_i^\alpha \in U$. A family $\mathcal{P} = (P^\alpha \mid \alpha \in \text{cf } \kappa)$ of Y^α - U bundles is a \mathcal{Y} - U bundle if $P^\alpha - U$ and $P^\beta - U$ are disjoint for all $\alpha, \beta \in \text{cf } \kappa$ with $\alpha \neq \beta$. Note that if a \mathcal{Y} - U bundle exists, then U contains D_ℓ .

A set $U \subseteq V(G)$ distinguishes \mathcal{Y} if whenever y_i^α and y_j^β are in the same component of $G - U$ for $\alpha, \beta \in \text{cf } \kappa$ and $i, j \in [0, k]$, then $\alpha = \beta$.

Lemma 4.7.3. *If a set $U \subseteq V(G)$ distinguishes \mathcal{Y} , then there is a \mathcal{Y} - U bundle \mathcal{P} .*

Proof. Let $U \subseteq V(G)$ distinguish \mathcal{Y} . By definition every finite set separating Y^α from Y^β in G also has to separate A^α from A^β . Since A is k -connected in G , there are also k disjoint Y^α - Y^β paths in G by Theorem 4.2.1. Hence we fix the initial Y^α - U segments of these paths for each $\alpha \in \text{cf } \kappa$, which are disjoint outside of U by the assumption that U distinguishes \mathcal{Y} . This yields the desired \mathcal{Y} - U bundle. \square

For a cardinal λ , a set $W \subseteq V(G)$ is λ -linked to a set $U \subseteq V(G)$, if for every $w \in W$ and every $u \in U$ there are λ many internally disjoint w - u paths in G .

The following lemma is the main part of the construction.

Lemma 4.7.4. *In Situation 4.7.2, suppose there is a set $U \subseteq V(G)$ such that*

- there is a \mathcal{Y} - U bundle $\mathcal{P} = (P^\alpha \mid \alpha \in \text{cf } \kappa)$; and
- there is a set $W \subseteq U$ with $|W| = k$ such that W is cf κ -linked to U .

Then there is an $I_0 \subseteq \text{cf } \kappa$ with $|I_0| = \text{cf } \kappa$ and a family $\mathcal{A}_0 = (A_0^\alpha \subseteq A^\alpha \mid \alpha \in I_0)$ with $|A_0^\alpha| = \kappa_\alpha$ for all $\alpha \in I_0$ such that ℓ - $FK_{k,\kappa}(\overline{\mathcal{K} \upharpoonright I_0})$ is an fbs-minor of G with $\overline{\mathcal{A}_0}$ along $\overline{\mathcal{Z} \upharpoonright I_0}$.

Proof. Let U , \mathcal{P} and W be as above. By Lemma 4.2.6 there is a $j \in [0, k]$ and a subset $I' \subseteq \text{cf } \kappa$ with $|I'| = \text{cf } \kappa$ such that (after possibly relabelling the sets Y^α for all $\alpha \in I'$ simultaneously) for every $\alpha, \beta \in I'$ with $\alpha \neq \beta$

- $y_i = u_i^\alpha = u_i^\beta$ for all $i \in [0, \ell)$;
- $x_i := u_i^\alpha = u_i^\beta$ for all $i \in [\ell, \ell + j)$; and
- $u_{i_0}^\alpha \neq u_{i_1}^\beta$ for all $i_0, i_1 \in [\ell + j, k)$.

Furthermore, after deleting at most j more elements from I' we obtain I'' such that

- $u_i^\alpha \neq y_i^\alpha$ for all $i \in [\ell, \ell + j)$ and all $\alpha \in I''$.

Note that if $|U| < \text{cf } \kappa$, then $\ell + j = k$ and we set $I_0 := I''$ and $L := \emptyset$.

Otherwise we construct subdivided stars with distinct centres in W . We start with a $k - \ell - j$ element subset $W' = \{w_i \mid i \in [\ell + j, k)\} \subseteq W$ disjoint from both D_ℓ as well as $\{x_i \mid i \in [\ell, \ell + j)\}$. A subgraph L of G is a *partial star-link* if there is a set $I(L) \subseteq I''$ such that L is the disjoint union of subdivided stars S_i for all $i \in [\ell + j, k)$ with centre w_i and leaves u_i^α , and L is disjoint to $P^\alpha - \{u_i^\alpha \mid i \in [\ell + j, k)\}$ for all $\alpha \in I(L)$. A partial star-link L is a *star-link* if $|I(L)| = \text{cf } \kappa$. Note that the union of a chain of partial star-links (ordered by the subgraph relation) yields another partial star-link. Hence by Zorn's Lemma there is a maximal partial star-link M . Assume for a contradiction that M is not a star-link. Then the set $N = V(M) \cup \bigcup_{\alpha \in I(M)} V(P^\alpha)$ has size less than $\text{cf } \kappa$. Take some $\beta \in I' \setminus I(M)$ such that M is disjoint to P^β . Since W is cf κ -linked to U , we can find $k - \ell - j$ disjoint $W' - \{u_i^\beta \mid i \in [\ell + j, k)\}$ paths disjoint from $N \setminus W'$, contradicting the maximality of M (after possibly relabelling). Hence there is a star-link L , and we set $I_0 := I(L)$.

Let H_{I_0} denote the subgraph of H containing only the branch sets for vertices in $Y^\alpha \cup Z^\alpha$ for $\alpha \in I_0$. Since $L \cup \bigcup_{\alpha \in I_0} P^\alpha$ has size cf $\kappa < \kappa_\alpha$ for all $\alpha \in I_0$, we can remove every branch set for some $z \in Z^\alpha$ meeting $L \cup \bigcup_{\alpha \in I_0} P^\alpha$ and obtain $A_0^\alpha \subseteq A^\alpha$ with $|A_0^\alpha| = \kappa_\alpha$. The union of the resulting subgraph with L and $\bigcup_{\alpha \in I_0} P^\alpha$ witnesses that $\ell\text{-FK}_{k,\kappa}(\overline{\mathcal{K} \upharpoonright I_0})$ is an fbs-minor of G with $\overline{\mathcal{A}_0} := (\overline{A_0^\alpha} \mid \alpha \in I_0)$ along $\overline{\mathcal{Z} \upharpoonright I_0}$. \square

As before, the previous lemma can be translated to find a desired subdivision of a generalised $\ell\text{-FK}_{k,\kappa}$.

Lemma 4.7.5. *In the situation of Lemma 4.7.4, there is an $I_1 \subseteq I_0$ with $|I_1| = \text{cf } \kappa$ and a family $\mathcal{A}_1 = (A_1^\alpha \subseteq A_0^\alpha \mid \alpha \in I_1)$ with $|A_1^\alpha| = \kappa_\alpha$ for all $\alpha \in I_1$ such that G contains a subdivision of a generalised $\ell\text{-FK}_{k,\kappa}(\overline{\mathcal{K} \upharpoonright I_1})$ with core $\bigcup \mathcal{A}_1$.*

Proof. Let H be the inflated subgraph witnessing that $\ell\text{-FK}_{k,\kappa}(\overline{\mathcal{K} \upharpoonright I_0})$ is an fbs-minor of G with $\overline{\mathcal{A}_0}$ along its core. Let $H^\alpha \subseteq H$ be the subgraph corresponding to the subgraph $K(Y^\alpha, Z^\alpha)$ of $\ell\text{-FK}_{k,\kappa}(\overline{\mathcal{K} \upharpoonright I_0})$ for each $\alpha \in I_0$. For each $\alpha \in I_0$ we apply Lemma 4.6.2 to H^α . By the pigeonhole principle there is a set $I_1 \subseteq I_0$ with $|I_1| = \text{cf } \kappa$ such that the type-1 k -template we got is the same for each $\alpha \in I_1$. This yields the desired subdivision as for Corollary 4.6.3. \square

The remainder of this subsection is dedicated to identify when we can apply Lemma 4.7.4.

Lemma 4.7.6. *In Situation 4.7.2, if either cf κ is uncountable or there is no end in the closure of some transversal T of \mathcal{A} , then there is a set $U \subseteq V(G)$ with the properties needed for Lemma 4.7.4.*

Proof. We start with a transversal T of \mathcal{A} (whose closure does not contain any end if cf κ is countable). We apply Lemma 4.6.1 to T to obtain an inflated subgraph witnessing that $K_{k,\text{cf } \kappa}$ is an fbs-minor of G with $T_0 \subseteq T$ along its core. We call the union of the singleton branch sets for the vertices of the finite side $W =: U_0$. By construction W is cf κ -linked to U_0 . Let I_0 denote the set $\{\alpha \in \text{cf } \kappa \mid |T_0 \cap A^\alpha| = 1\}$. We construct U inductively.

For some ordinal α we assume we already constructed a strictly \subseteq -ascending sequence $(U_\beta \mid \beta < \alpha)$ such that W is cf κ -linked to U_β for all $\beta < \alpha$. If there is a subset $I \subseteq I_0$ with $|I| = \text{cf } \kappa$ such that $U' := \bigcup_{\beta < \alpha} U_\beta$ distinguishes $\overline{\mathcal{Y} \upharpoonright I}$, then we are done by Lemma 4.7.3 since by construction W is still cf κ -linked to U' .

Otherwise there is a component of $G - U'$ containing a transversal T_α of $\mathcal{Y}|I_\alpha$ for some $I_\alpha \subseteq I_0$ with $|I_\alpha| = \text{cf } \kappa$. Applying Lemma 4.2.3 to T_α yields a subdivided star with centre u_α and $\text{cf } \kappa$ many leaves $L_\alpha \subseteq T_\alpha$. We then set $U_\alpha := U' \cup \{u_\alpha\}$. By Theorem 4.2.2 there are $\text{cf } \kappa$ many internally disjoint $w - u_\alpha$ paths for all $w \in W$, since no set of size less than $\text{cf } \kappa$ could separate u_α from L_α , L_α from T_0 , or any subset of size $\text{cf } \kappa$ of T_0 from w . Hence W is $\text{cf } \kappa$ -linked to U_α and we can continue the construction. This construction terminates at the latest if $U' = V(G)$. \square

If $\text{cf } \kappa$ is countable and there is an end in the closure of some transversal of \mathcal{A} , then there is still a chance to obtain an ℓ - $FK_{k,\kappa}$ minor. We just need to check whether G contains a \mathcal{Y} - $\text{Dom}(\omega)$ bundle, since we have the following lemma.

Lemma 4.7.7. *For every end $\omega \in \Omega(G)$, the set $\text{Dom}(\omega)$ is \aleph_0 -linked to itself.*

Proof. Suppose there are $u, v \in \text{Dom}(\omega)$ with only finitely many internally disjoint $u - v$ paths. Hence there is a finite separator $S \subseteq V(G)$ such that u and v are in different components of $G - S$. Then at least one of them is in a different component than $C(S, \omega)$, a contradiction. \square

Hence, we obtain the final corollary of this subsection.

Corollary 4.7.8. *In Situation 4.7.2, suppose $\text{cf } \kappa$ is countable and there is an end ω in the closure of some transversal of \mathcal{A} with $\text{dom } \omega \geq k$ such that $\text{Dom}(\omega)$ distinguishes \mathcal{Y} . Then $\text{Dom}(\omega)$ satisfies the properties needed for Lemma 4.7.4. \square*

4.7.3. Minors for singular k -blueprints

This subsection builds differently upon Situation 4.7.2 in the case where we do not obtain the frayed complete bipartite minor. We incorporate new assumptions and notation, establishing a new situation, which we will further modify according to some assumptions that we can make without loss of generality during this subsection.

Situation 4.7.9. Building upon Situation 4.7.2, suppose $\text{cf } \kappa$ is countable and there is an end ω in the closure of some transversal of \mathcal{A} , i.e. an ω -comb whose teeth are a transversal T of $\{A^i \mid i \in J\}$ for some infinite $J \subseteq \mathbb{N}$. Suppose that

- (*) there is no $\overline{\mathcal{Y}|I}$ - $\text{Dom}(\omega)$ bundle for any infinite $I \subseteq \mathbb{N}$.

In particular $\text{Dom}(\omega)$ does not distinguish $\overline{\mathcal{Y} \upharpoonright J}$ by Lemma 4.7.3. Hence there is a component C of $G - \text{Dom}(\omega)$ containing a comb with teeth in $\mathcal{Y} \upharpoonright J$, since a subdivided star would yield a vertex dominating ω outside $\text{Dom}(\omega)$. This comb is an ω -comb since its teeth cannot be separated from T by a finite vertex set. Without loss of generality we may assume that $J = \mathbb{N}$ by redefining \mathcal{K} , \mathcal{Y} and \mathcal{A} as $\overline{\mathcal{K} \upharpoonright J}$, $\overline{\mathcal{Y} \upharpoonright J}$ and $\overline{\mathcal{A} \upharpoonright J}$ respectively.

Let $G' := G[C]$ and let ω' be the end of G' containing the spine of the aforementioned ω -comb in G' . Let $\mathcal{S} = (S^n \mid n \in \mathbb{N})$ be an ω' -defining sequence in G' and let \mathcal{R} be a family of disjoint ω' -rays in G' such that $(\mathcal{S}, \mathcal{R})$ witnesses the degree of the undominated end ω' of G' , which exist by Lemma 4.5.3. Moreover, we will modify this situation with some assumptions that we can make without loss of generality. We will fix them in some of the following lemmas and corollaries.

Lemma 4.7.10. *In Situation 4.7.9, we may assume without loss of generality that for all $n \in \mathbb{N}$ the following hold:*

- $S^n \cap \bigcup \mathcal{Y} = \emptyset$; and
- S^n is contained in a component of $G'[S^n, S^{n+1}]$.

Hence we include these assumptions into Situation 4.7.9.

Proof. Given $x, y \in \mathbb{N}$ we can choose $n \in \mathbb{N}$ with $n \geq y$ and $m \in \mathbb{N}$ with $m > n$ such that S^n is contained in a component of $G'[S^n, S^m]$ and Y^x is disjoint to $S^n \cup S^m$. Note that it is possible to incorporate the first property since $(\mathcal{S}, \mathcal{R})$ is degree-witnessing in G' . Iteratively applying this observation yields subsequences of \mathcal{S} and \mathcal{Y} . Taking the respective subsequences of \mathcal{K} and \mathcal{A} and relabelling all of them accordingly as before yields the claim. \square

Lemma 4.7.11. *In Situation 4.7.9, we may assume without loss of generality that $\emptyset \neq Y^n \setminus \text{Dom}(\omega) \subseteq V(G'[S^n, S^{n+1}])$ for all $n \in \mathbb{N}$. Hence we include this assumption into Situation 4.7.9.*

Proof. Note that $Y^n \setminus \text{Dom}(\omega) = \emptyset$ for only finitely many $n \in \mathbb{N}$ by (*). Moreover, for all but finitely many $n \in \mathbb{N}$ there is an $x_n \in \mathbb{N}$ such that $Y^{x_n} \setminus \text{Dom}(\omega)$ meets $V(G'[S^n, S^{n+1}])$ since ω is in the closure of \mathcal{Y} . Suppose that $Y^{x_n} \setminus \text{Dom}(\omega)$ is not contained in $V(G'[S^n, S^{n+1}])$ for some $n \in \mathbb{N}$. Since for any $i, j \in [0, k]$ with $i \neq j$ there are κ_x many disjoint $y_i^x - y_j^x$ paths in H^x , all but finitely many of them have

to traverse $\text{Dom}(\omega)$. In particular, there is an $Y^x - \text{Dom}(\omega)$ bundle in H^x . Such a bundle trivially also exists if $Y^x \subseteq \text{Dom}(\omega)$. If this happens for all x in some infinite $I \subseteq \mathbb{N}$, then there is a $\overline{\mathcal{Y}} \overline{I} - \text{Dom}(\omega)$ bundle in G , contradicting $(*)$. Hence this happens at most finitely often. Again, relabelling and taking subsequences yields the claim. \square

The following lemma allows some control on how we can find a set of disjoint paths from Y^n to the rays in \mathcal{R} and has two important corollaries.

Lemma 4.7.12. *In Situation 4.7.9, let $\mathcal{R}' \subseteq \mathcal{R}$ with $|\mathcal{R}'| = \min(\deg(\omega'), k)$. Then for all $n > 2k$ there is an $M > n$ such that for all $m \geq M$ there exists an $Y^n - (\text{Dom}(\omega) \cup (S^m \cap \bigcup \mathcal{R}'))$ bundle $P^{n,m}$ with $P^{n,m} - \text{Dom}(\omega) \subseteq G'[S^{n-2k}, S^m]$.*

Proof. Let $n > 2k$ be fixed. As in the proof of Lemma 4.7.3 for each $x > 0$ there are k disjoint $Y^n - Y^{x-1}$ paths in G , whose union contains a $Y^n - (\text{Dom}(\omega) \cup S^x)$ bundle Q^x in $G[C \cup \text{Dom}(\omega)]$. Considering Q^{n-2k} , let $M \in \mathbb{N}$ be large enough such that $Q^{n-2k} - \text{Dom}(\omega) \subseteq G'[S^{n-2k}, S^M]$, and let $m \geq M$.

Suppose for a contradiction that there is a vertex set S of size less than k separating Y^n from $\text{Dom}(\omega) \cup (S^m \cap \bigcup \mathcal{R}')$ in $G[V(G'[S^{n-2k}, S^m]) \cup \text{Dom}(\omega)]$. Then for at least one $i \in [n - 2k, n)$ the graph $G'[S^i, S^{i+1}]$ does not contain a vertex of S . We distinguish two cases.

Suppose $\deg(\omega') \geq k$. Then S contains a vertex from every path of Q^{n-2k} ending in $\text{Dom}(\omega)$, but does not contain a vertex from every path of Q^{n-2k} . Let Q be such a $Y^n - S^{n-2k}$ path avoiding S . Now Q meets S^i by construction. There is at least one ray $R \in \mathcal{R}'$ that does not contain a vertex of S . Since S^i is contained in a component of $G'[S^i, S^{i+1}]$ and $R \cap S^i \neq \emptyset$, we can connect Q with R and hence with $S^m \cap R$ in $G'[S^i, S^{i+1}]$ avoiding S , which contradicts the assumption.

Suppose $\deg(\omega') < k$, then $\mathcal{R}' = \mathcal{R}$ and hence $S^m \cap \bigcup \mathcal{R}' = S^m$. As before, there is a $Y^n - S^m$ path Q in Q^m not containing a vertex of S . This path being contained in $G'[S^{n-2k}, S^m]$ would contradict the assumption. Hence we may assume the path meets S^j for every $j \in [n - 2k, m]$ and in particular S^i . Let $Q_1 \subseteq Q$ denote $Y^n - S^i$ path in $G'[S^i, S^m]$, and let $Q_2 \subseteq Q$ denote $S^i - S^m$ path in $G'[S^i, S^m]$. As before, we can connect Q_1 and Q_2 in $G'[S^i, S^{i+1}]$ avoiding S , which again contradicts the assumption. \square

Corollary 4.7.13. *In Situation 4.7.9, let $\mathcal{R}' \subseteq \mathcal{R}$ with $|\mathcal{R}'| = \min(\deg(\omega'), k)$. Without loss of generality for all $n \in \mathbb{N}$ there is a $Y^n - (\text{Dom}(\omega) \cup (S^{n+1} \cap \bigcup \mathcal{R}'))$*

bundle P^n such that $P^n - \text{Dom}(\omega) \subseteq G'[S^n, S^{n+1}] - S^n$. Hence we include this assumption into Situation 4.7.9.

Proof. We successively apply Lemma 4.7.12 to obtain suitable subsequences. Relabelling them yields the claim. \square

Corollary 4.7.14. *Situation 4.7.9 implies $\text{dom}(\omega) < k$.*

Proof. Suppose $\text{dom}(\omega) \geq k$. Then for every $n > 2k$ there is no $Y^n - \text{Dom}(\omega)$ separator S of size less than k by Lemma 4.7.12, since m can be chosen such that $S \cap C(S^m, \omega') = \emptyset$. Hence we can extend a path of the bundle in $C(S^m, \omega')$. Therefore, for each $n > 2k$ there is an $m > n$ such that we can find a $Y^n - \text{Dom}(\omega)$ bundle $P^{n,m}$ such that $P^{n,m} - \text{Dom}(\omega) \subseteq G'[S^{n-2k}, S^m]$, and consequently an infinite subset $I' \subseteq \mathbb{N}$ such that $\overline{(P^{n,m} \mid n \in I')}$ is an $\overline{\mathcal{Y} \upharpoonright I'} - \text{Dom}(\omega)$ bundle, contradicting the assumption (*) in Situation 4.7.9. \square

This last corollary is quite impactful. From this point onwards, we know that $\omega' = \omega \upharpoonright (G - \text{Dom}(\omega))$ by Remark 4.2.5.

Lemma 4.7.15. *In Situation 4.7.9, we may assume without loss of generality that for all $n \in \mathbb{N}$ the following hold:*

- $H^n - \text{Dom}(\omega) \subseteq G'[S^n, S^{n+1}] - (S^n \cup S^{n+1})$;
- $H^n \cap \text{Dom}(\omega) = D_\ell \subseteq Y^n$.

Hence we include these assumptions into Situation 4.7.9.

Proof. Note that $H^n \cap G'[S^n, S^{n+1}] - (S^n \cup S^{n+1}) \neq \emptyset$ by Lemmas 4.7.10 and 4.7.11. We delete the finitely many branch sets of vertices corresponding to the infinite side of K_{k, κ_n} in H^n containing a vertex of $\text{Dom}(\omega)$, S^n or S^{n+1} . Since the remaining inflated subgraph is connected, no branch set of the infinite side meets a vertex outside of $G'[S^n, S^{n+1}]$. Moreover, for all but finitely many $n \in \mathbb{N}$ the branch sets of vertices corresponding to the finite side of H^n that meet $\text{Dom}(\omega)$ are precisely the singletons of the elements in D_ℓ by Corollary 4.7.14. Deleting the exceptions and relabelling accordingly yields the claims. \square

The next lemma reroutes some rays to find a bundle from Y^n to those new rays and dominating vertices with some specific properties.

Lemma 4.7.16. *In Situation 4.7.9, there is a set \mathcal{R}'' of $|\mathcal{R}'|$ disjoint ω' -rays in G' and a $Y^n - (\cup \mathcal{R}'' \cup \text{Dom}(\omega))$ bundle Q^n for each $n \in \mathbb{N}$ such that for every $R'' \in \mathcal{R}''$*

- *there is an $R' \in \mathcal{R}'$ with $V(R'') \cap \cup \mathcal{S} = V(R') \cap \cup \mathcal{S}$; and*
- *$|Q^n \cap R''| \leq 2$ for every $n \in \mathbb{N}$.*

Hence we include references to these objects into Situation 4.7.9.

Proof. Given $n \in \mathbb{N}$, let P^n be as in Corollary 4.7.13. Let \mathcal{P} be a set of $|\mathcal{R}'|$ disjoint $S^n - S^{n+1}$ paths in $G'[S^n, S^{n+1}]$ each with end vertices $R' \cap (S^n \cup S^{n+1})$ for some $R' \in \mathcal{R}'$. We call such a set \mathcal{P} *feasible*. For a feasible \mathcal{P} , let $P^n(\mathcal{P})$ denote the $Y^n - (\text{Dom}(\omega) \cup \cup \mathcal{P})$ bundle contained in P^n and let $p^n(\mathcal{P})$ denote the finite parameter $|P^n - P^n(\mathcal{P}) - \cup \mathcal{P}|$. Note that $\{R' \cap G'[S^n, S^{n+1}] \mid R' \in \mathcal{R}'\}$ is a feasible set. Now choose a feasible \mathcal{P}^n such that $p^n(\mathcal{P}^n)$ is minimal and let $Q^n := P^n(\mathcal{P}^n)$.

Assume for a contradiction that there is a path $P \in \mathcal{P}^n$ with $|Q^n \cap P| > 2$. Let v_0, v_1 and v_2 denote vertices in this intersection such that $v_1 \in V(v_0 P v_2)$. We replace the segment $v_0 P v_2$ by the path consisting of the paths Q_i^n and Q_j^n that contain v_0 and v_2 respectively, as well as any $y_i^n - y_j^n$ path in H^n avoiding the finite set $\text{Dom}(\omega) \cup Q^n \cup S^n \cup S^{n+1}$. The resulting set \mathcal{P} is again feasible and the parameter $p^n(\mathcal{P})$ is strictly smaller than $p^n(\mathcal{P}^n)$, contradicting the choice of \mathcal{P}^n .

Now let \mathcal{R}'' be the set of components in the union $\cup \{P^n \mid n \in \mathbb{N}\}$. Indeed, this is a set of ω' -rays that together with the bundles Q^n satisfy the desired properties. \square

For $m, n \in \mathbb{N}$, we say Q^m and Q^n *follow the same pattern*, if for all $i, j \in [0, k)$

- Q_i^m and Q_i^n either meet the same ray in \mathcal{R}'' or the same vertex in $\text{Dom}(\omega)$;
- if Q_i^m and Q_j^m both meet some $R \in \mathcal{R}''$ and Q_i^m meets R closer to the start vertex of R than Q_j^m , then Q_i^n meets R closer to the start vertex of R than Q_j^n .

Lemma 4.7.17. *In Situation 4.7.9, we may assume without loss of generality that*

- *there are $k_0, k_1, f \in \mathbb{N}$ with $1 \leq k_0 \leq \deg(\omega')$, $0 \leq \ell + f + k_1 \leq \text{dom}(\omega)$ and $\ell + f + k_0 + k_1 = k$;*

- there is a subset $\mathcal{R}_0 \subseteq \mathcal{R}''$ with $|\mathcal{R}_0| = k_0$; and
- there are disjoint $D_f, D_1 \subseteq \text{Dom}(\omega) \setminus D_\ell$ with $|D_f| = f$ and $|D_1| = k_1$;

such that for all $m, n \in \mathbb{N}$

- (a) Q^n is a $Y^n - (\cup \mathcal{R}_0 \cup D_\ell \cup D_f)$ bundle;
- (b) $Q^n \cap \text{Dom}(\omega) = D_f \cup D_\ell$; and
- (c) Q^m and Q^n follow the same pattern.

Hence we include these assumptions and references to the existing objects into Situation 4.7.9.

Proof. Using the fact that $\text{Dom}(\omega)$ is finite, we apply the pigeonhole principle to find a set $D_f \subseteq \text{Dom}(\omega) \setminus D_\ell$ and an infinite subset $I \subseteq \mathbb{N}$ such that (b) hold for all $n \in I$. Set $f := |D_f|$. Applying it multiple times again, we find an infinite subset $I' \subseteq I$ such that (c) holds for all $m, n \in I'$. If $|\mathcal{R}''| \geq k - \ell - f$, then set \mathcal{R}_0 to be any subset of \mathcal{R}'' of size $k - \ell - f$ containing each ray that meets Q^n for any $n \in I'$. Otherwise set $\mathcal{R}_0 := \mathcal{R}''$ and set $k_0 := |\mathcal{R}_0| = \deg(\omega') = \deg(\omega)$. Now (a) holds by the choices of D_f and \mathcal{R}_0 . Since $\Delta(\omega) \geq k$ by Lemma 4.6.5, there is a set $D_1 \subseteq \text{Dom}(\omega) \setminus (D_\ell \cup D_f)$ of size $k_1 := k - \ell - f - k_0$, completing the proof. \square

Finally, we construct the subdivision of a generalised k -typical graph for some singular k -blueprint.

Lemma 4.7.18. *In Situation 4.7.9, there is a singular k -blueprint $\mathcal{B} = (\ell, f, B, D)$ for a tree B of order $k_0 + k_1$ with $|D| = k_0$, such that G contains a subdivision of a generalised $T_k(\mathcal{B})(\mathcal{K})$ with core $\cup \mathcal{A}$.*

Proof. We apply Lemma 4.5.10 to \mathcal{R}_0 and D_1 to obtain a simple type-2 k -template \mathcal{T}_2 , a tree B of order $k_0 + k_1$ and a set $D \subseteq V(B)$ with $|D| = k_1$ such that (\mathcal{R}_0, D_1) is (B, \mathcal{T}_2) -connected. For each $n \in \mathbb{N}$ let Γ_n denote a (B, \mathcal{T}_2) -connection avoiding S^n , $\text{Dom}(\omega) \setminus D_1$ as well as for each $R \in \mathcal{R}_0$ its initial segment R_s for $s \in (S^n \cap V(R))$.

Note that there is an $m > n$ such that $\Gamma_n - D_1 \subseteq G'[S^n, S^m] - S^m$. Hence Γ_n and Γ_{m+1} are disjoint to $G'[S^m, S^{m+1}] \supseteq Q^m$. For rays $R \in \mathcal{R}_0$ with $|Q^m \cap R| \geq 1$,

we extend Γ_n on that ray to include precisely one vertex in the intersection as well as with the corresponding path in Q^m to Y^m . If furthermore $|Q^m \cap R| = 2$, we also extend Γ_{m+1} on that ray to include the other vertex of the intersection and with the corresponding path in Q^m to Y^m . Since Q^m and Q^n follow the same pattern for all $m, n \in \mathbb{N}$ by Lemma 4.7.17, we can modify \mathcal{T}_2 to \mathcal{T}'_2 accordingly to have infinitely many (B, \mathcal{T}'_2) -connections which pairwise meet only in D_1 and contain Y^n for each $n \in I$ for some infinite subset $I \subseteq \mathbb{N}$. After relabelling and setting $\mathcal{B} := (\ell, f, B, D)$, we obtain the subdivision of $T_k(\mathcal{B})(\mathcal{T}'_2)$ as in the proof of Corollary 4.5.11. \square

4.7.4. Characterisation for singular cardinals

Now we have developed all the necessary tools to prove the minor and topological minor part of the characterisation in Theorem 4.3.7 for singular cardinals.

Theorem 4.7.19. *Let G be a graph, let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and let $\kappa \leq |A|$ be a singular cardinal. Then the following statements are equivalent.*

- (a) *There is a subset $A_1 \subseteq A$ with $|A_1| = \kappa$ such that A_1 is k -connected in G .*
- (b) *There is a subset $A_2 \subseteq A$ with $|A_2| = \kappa$ such that*
 - *either G contains an ℓ -degenerate frayed $K_{k,\kappa}$ as an fbs-minor with A_2 along its core for some $0 \leq \ell \leq k$;*
 - *or $T_k(\mathcal{B})$ is an fbs-minor of G with A_2 along its core for a singular k -blueprint \mathcal{B} .*
- (c) *There is a subset $A_3 \subseteq A$ with $|A_3| = \kappa$ such that*
 - *either G contains a subdivision of a generalised ℓ -FK $_{k,\kappa}$ with core A_3 for some $0 \leq \ell \leq k$;*
 - *or G contains the subdivision of a generalised $T_k(\mathcal{B})$ with core A_3 for some singular k -blueprint \mathcal{B} .*

Moreover, if these statements hold, we can choose $A_1 = A_2 = A_3$.

Proof. If (b) holds, then A_2 is k -connected in G by Lemma 4.4.2 with Lemma 4.3.4.

Suppose (a) holds. Either we can find a subset $A_3 \subseteq A_1$ with $|A_3| = \kappa$ and a subdivision of ℓ -FK $_{k,\kappa}(\mathcal{K})$ with core A_3 for some good κ -sequence \mathcal{K} by Lemma 4.7.4

and either Lemma 4.7.6 or Corollary 4.7.8. Otherwise, we can apply Lemma 4.7.18 to obtain $A_3 \subseteq A_1$ with $|A_3| = \kappa$ and a subdivision of $T_k(\mathcal{B})(\mathcal{K})$ with core A_3 for some singular k -blueprint \mathcal{B} and a good κ -sequence \mathcal{K} . With Remark 4.3.5 we obtain the subdivision of the respective generalised k -typical graph with respect to the fixed good κ -sequence.

If (c) holds, then so does (b) by Lemma 4.4.3 with $A_2 := A_3$. Moreover, A_3 is a candidate for both A_2 and A_1 . \square

4.8. Applications of the minor-characterisation

In this section we will present some applications of the minor-characterisation of k -connected sets.

As a first corollary we just restate the theorem for $k = 1$, giving us a version of the Star-Comb Lemma for singular cardinalities. For this, given a singular cardinal κ , we call the graph $FK_{1,\kappa}$ a *frayed star*, whose *centre* is the vertex x_0 of degree κ and whose *leaves* are the vertices $\cup \mathcal{Z}$. Moreover, we call the 1-typical graph obtained from the single vertex tree (i.e. $T_1(0, 0, (\{c\}, \emptyset), \emptyset, 0 \mapsto (c, 0))$) a *frayed comb* with *spine* \mathfrak{N}_c and *teeth* $\cup \mathcal{Z}$. Note that each generalised frayed star or generalised frayed comb contains a subdivision of the frayed star or frayed comb respectively.

Corollary 4.8.1 (Frayed-Star-Comb Lemma). *Let $U \subseteq V(G)$ be infinite and let $\kappa \leq |U|$ be a singular cardinal. Then the following statements are equivalent.*

- (a) *There is a subset $U_1 \subseteq U$ with $|U_1| = \kappa$ such that U_1 is 1-connected in G .*
- (b) *There is a subset $U_2 \subseteq U$ with $|U_2| = \kappa$ such that G either contains a subdivided star or frayed star whose set of leaves is U_2 , or a subdivided frayed comb whose set of teeth is U_2 .*

(Note that if κ is uncountable, only one of the former two can exist.)

Moreover, if these statements hold, we can choose $U_1 = U_2$. \square

Even though this Frayed-Star-Comb Lemma has a much more elementary proof, we state it here only as a corollary of our main theorem.

Now Theorems 4.6.7 and 4.7.19 give us the tools to prove the statement we originally wanted to prove instead of Lemma 4.4.4.

Corollary 4.8.2. *Let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and k -connected in G and let $\kappa \leq |A|$ be an infinite cardinal. Then for every $v \in V(G)$ there is a subset $A' \subseteq A$ with $|A'| = \kappa$ such that A' is $(k - 1)$ -connected in $G - v$.*

Proof. First we apply Theorem 4.6.7 or Theorem 4.7.19 to A to get a k -typical graph T and an inflated subgraph H witnessing that T is an fbs-minor of G with some $A'' \subseteq A$ along its core such that $|A''| = \kappa$. Let us call a vertex of T *essential*, if either

- it is a vertex of the finite side of $K_{k,\kappa}$ if $T = K_{k,\kappa}$;
- it is a degenerate vertex or frayed centre of ℓ - $FK_{k,\kappa}$ if $T = \ell$ - $FK_{k,\kappa}$ for some $\ell \in [0, k]$; or
- it is a dominating vertex, a degenerate vertex or a frayed centre of $T_k(\mathcal{B})$ if $T = T_k(\mathcal{B})$ for some regular or singular k -blueprint \mathcal{B} .

We distinguish four cases.

If $v \notin V(H)$, then $H \subseteq G - v$ still witnesses that T is an fbs-minor of $G - v$ with $A' := A''$ along its core.

If v belongs to a branch set of a vertex c of the core, then the inflated subgraph obtained by deleting that branch set still yields a witness that T is an fbs-minor of $G - v$ with $A' := A'' \setminus \{c\}$ along its core.

If v belongs to a branch set of an essential vertex $w \in V(T)$, then the inflated subgraph where we delete this branch set from H witnesses that the obvious $(k - 1)$ -typical subgraph of $T - w$ is an fbs-minor of $G - v$ with $A' := A''$ along its core.

If v belongs to a branch set of a vertex $w \in V(\mathfrak{N}(B/D) - D)$, then we delete the branch sets of the layers (not including D) up to the layer containing w and relabelling accordingly (and modifying the κ -sequence if necessary). This yields a supergraph of an inflated subgraph witnessing that T is an fbs-minor of $G - v$ with A' along its core for some $A' \subseteq A''$ with $|A'| = \kappa$. Similar arguments yield the statement if v belongs to a branch set of a neighbour of a frayed centre.

In any case, with the other direction of Theorem 4.6.7 or Theorem 4.7.19 we get that A' is $(k - 1)$ -connected in $G - v$. \square

As another corollary we prove that we are able to find k -connected sets of size κ in sets which cannot be separated by less than κ many vertices from

another k -connected set. This will be an important tool for our last part of the characterisation in the main theorem.

Corollary 4.8.3. *Let $k \in \mathbb{N}$, let $A, B \subseteq V(G)$ be infinite and let $\kappa \leq |A|$ be an infinite cardinal. If B is k -connected in G and A cannot be separated from B by less than κ vertices, then there is an $A' \subseteq A$ with $|A'| = \kappa$ which is k -connected in G .*

Proof. Let \mathcal{P} be a set of κ many disjoint A – B paths as given by Theorem 4.2.2. Let B' denote $B \cap \bigcup \mathcal{P}$. Let $H \subseteq G$ be an inflated subgraph witnessing that a k -typical graph is an fbs-minor of G with B'' along its core for some $B'' \subseteq B'$ with $|B''| = \kappa$ as given by Theorem 4.6.7 or Theorem 4.7.19. Let \mathcal{P}' denote the set of the A – H subpaths of the A – B'' paths in \mathcal{P} . We distinguish two cases.

If the k -typical graph is a $T_k(\mathcal{B})$ for some regular k -blueprint $\mathcal{B} = (T, D, c)$, then (since each branch set in H is finite) there is an infinite subset $\mathcal{P}'' \subseteq \mathcal{P}'$ and a node $c' \in V(T \setminus D)$ such that each branch set in H of vertices in $V(\mathfrak{N}_{c'})$ meets $\bigcup \mathcal{P}''$ at most once and no other branch set meets $\bigcup \mathcal{P}''$. Let $A' := \bigcup \mathcal{P}'' \cap A$. We extend each of these branch sets with the path from \mathcal{P}'' meeting it. This yields a subgraph H' witnessing that $T_k(T, D, c')$ is an fbs-minor of G with some A' along its core with $A' \subseteq A''$.

Otherwise, since each branch set in H is finite, there is a subset $\mathcal{P}'' \subseteq \mathcal{P}'$ of size κ such that each branch set in H of vertices corresponding to the core meets $\bigcup \mathcal{P}''$ at most once and no other branch set meets $\bigcup \mathcal{P}''$. Let $A' := \bigcup \mathcal{P}'' \cap A$. Again, we extend each of these branch sets with the path from \mathcal{P}'' meeting it. This yields a subgraph H' witnessing that the same k -typical graph is an fbs-minor of G with some A'' along its core such that $A' \subseteq A'' \subseteq A' \cup B''$.

Applying Theorem 4.6.7 or Theorem 4.7.19 again together with Remark 4.4.1 yields the claim. \square

4.9. Nested separation systems

This section will finish the proof of Theorem 4.3.7 by proving the duality theorem and hence providing the last equivalence of the characterisation.

Recall that a nested separation system $N \subseteq S_k(G)$ is called k -lean if given any two (not necessarily distinct) parts P_1, P_2 of N and vertex sets $Z_1 \subseteq P_1, Z_2 \subseteq P_2$

with $|Z_1| = |Z_2| = \ell \leq k$ there are either ℓ disjoint $Z_1 - Z_2$ paths in G or there is a separation (A, B) in N with $P_1 \subseteq A$ and $P_2 \subseteq B$ of order less than ℓ .

For a subset $X \subseteq V(G)$ consider the induced subgraph $G[X]$. Every separation of $G[X]$ is of the form $(A \cap X, B \cap X)$ for some separation (A, B) of G . We denote this separation also as $(A, B) \upharpoonright X$. Given a set S of separations of G we write $S \upharpoonright X$ for the set consisting of all separations $(A, B) \upharpoonright X$ for $(A, B) \in S$.

Consider the directed partially ordered set \mathcal{F} of finite subsets of $V(G)$ ordered by inclusion, as well as the directed inverse system $(S_k(G[X]) \mid X \in \mathcal{F})$.

Observation 4.9.1. Every separation in $S_k(G)$ is determined by all its restrictions to finite subsets of $V(G)$.

More precisely, on the one hand for each element $((A_X, B_X) \in S_k(G[X]) \mid X \in \mathcal{F})$ of the inverse limit the separation $(\cup\{A_X \mid X \in \mathcal{F}\}, \cup\{B_X \mid X \in \mathcal{F}\})$ is the unique separation in $S_k(G)$ inducing (A_X, B_X) on $G[X]$ for each $X \in \mathcal{F}$. On the other hand, $((A, B) \upharpoonright X \mid X \in \mathcal{F})$ is an element of the inverse limit for each separation $(A, B) \in S_k(G)$. For more information on this approach, see [12, 14].

The following theorem lifts the existence of k -lean nested separation systems for finite graphs as in Theorem 4.2.7 to infinite graphs via the Generalised Infinity Lemma.

Theorem 4.9.2. *For every graph G and every $k \in \mathbb{N}$ there is a nested separation system $N \subseteq S_k(G)$ such that N is k -lean.*

Proof. As above, consider the directed partially ordered set \mathcal{F} of finite subsets of $V(G)$ ordered by inclusion. For every $X \in \mathcal{F}$ let $\mathcal{N}(X)$ denote the set of nested separation systems $N \upharpoonright X$ of $G[X]$ such that there is a nested separation system $N \subseteq S_k(G[Z])$ that is k -lean for a $Z \in \mathcal{F}$ containing X . Note that $\mathcal{N}(X)$ is not empty by Theorem 4.2.7 for every $X \in \mathcal{F}$. Moreover, for every $Y \subseteq X \in \mathcal{F}$ there is a natural map $f_{X,Y} : \mathcal{N}(X) \rightarrow \mathcal{N}(Y)$ defined by $f_{X,Y}(N) := N \upharpoonright Y$. It is easy to check that this defines a directed inverse system of finite sets. By the Generalised Infinity Lemma the inverse limit of that system is non-empty and contains an element $(N_X \mid X \in \mathcal{F})$.

Let $N := \{(A, B) \in S_k(G) \mid (A, B) \upharpoonright X \in N_X \text{ for all } X \in \mathcal{F}\}$. Note that N is non-empty and contains for each $(A, B) \in N_X$ at least one separation inducing (A, B) on $G[X]$ by Observation 4.9.1. It is easy to check that N is a nested

separation system since the fact that two separations are crossing is witnessed in a finite set of vertices.

For two (not necessarily distinct) parts P_1, P_2 of N and vertex sets $Z_1 \subseteq P_1$ and $Z_2 \subseteq P_2$ with $|Z_1| = |Z_2| = \ell \leq k$, we consider a largest possible set \mathcal{P} of disjoint Z_1 – Z_2 paths in G . We may assume that $|\mathcal{P}| < \ell$, since otherwise there is nothing to show. For every $X \in \mathcal{F}$ containing $Z := Z_1 \cup Z_2 \cup V(\cup \mathcal{P})$, note that N_X is the restriction of some k -lean tree set of a finite supergraph $G[X']$ of $G[X]$ to X . Hence there is a separation (A, B) of $G[X']$ of order $|\mathcal{P}|$ separating Z_1 and Z_2 , whose restriction $(A, B)|_X$ is non-trivial and in N_X by the choice of X . For each finite $X, Y \subseteq V(G)$ with $Z \subseteq Y \subseteq X$, each such separation (C, D) induces a separation $(C, D)|_Y \in N_Y$ of order $|\mathcal{P}|$ separating Z_1 and Z_2 . Applying the Generalised Infinity Lemma again yields an element $((A_X, B_X) \in N_X \mid X \in \mathcal{F})$ from the inverse limit. By Observation 4.9.1, this element corresponds to a separation of order $|\mathcal{P}|$ of G which by construction separates Z_1 and Z_2 and is an element of N . Hence N is k -lean. \square

Now we are able to prove the duality theorem and hence the remaining equivalence of our main theorem.

Theorem 4.9.3. *Let G be an infinite graph, let $k \in \mathbb{N}$, let $A \subseteq V(G)$ be infinite and let $\kappa \leq |A|$ be an infinite cardinal. Then the following statements are equivalent.*

- (a) *There is a subset $A_1 \subseteq A$ with $|A_1| = \kappa$ such that A_1 is k -connected in G .*
- (d) *There is no nested separation system $N \subseteq S_k(G)$ such that every part P of N can be separated from A by less than κ vertices.*

Proof. Assume that (a) does not hold. Let N be a k -lean nested separation system as obtained from Theorem 4.9.2. Suppose for a contradiction that there exists a part P of N that cannot be separated from A by less than κ vertices. Then P is k -connected in G and has size at least κ . By Corollary 4.8.3, there is a subset $A_1 \subseteq A$ of size κ which is k -connected in G , a contradiction. Hence every part of N can be separated from A by less than κ vertices, so (d) does not hold.

If (a) holds, let $N \subseteq S_k(G)$ be any nested separation system and let H be an inflated subgraph witnessing that a k -typical graph T is an fbs-minor of G with some $A' \subseteq A$ along its core for $|A'| = \kappa$ as in Theorem 4.6.7 or Theorem 4.7.19.

If $T = T_k(B, D, c)$ for some regular k -blueprint (B, D, c) , then since T contains k disjoint paths between $B_i \cup D$ and $B_j \cup D$ for all $i, j \in \mathbb{N}$, no separation of G of order less than k can separate the unions of the branch sets corresponding to the vertices of the layers $B_i \cup D$ and $B_j \cup D$. Hence there is a part of N containing at least one vertex in a branch set corresponding to some vertex of every layer of T .

In every other case T contains k internally disjoint paths between any two core vertices. Hence there cannot exist a separation of G of order less than k that separates two distinct branch sets containing vertices of the core, and therefore there is a part of N containing at least one vertex from each branch set corresponding to the core of T .

In any case, this part has to have size at least κ , and the disjoint paths in each branch set from a vertex of A' to the part witness by Theorem 4.2.2 that A cannot be separated by less than κ vertices from that part. Since N was arbitrarily chosen, (d) holds. \square

Let us finish this section with an open problem regarding the question when it is possible to extend this duality theorem to tree-decompositions.

Problem 4.9.4. *For which class of infinite graphs is the existence of a k -connected set of size κ equivalent to the non-existence of a tree-decomposition of adhesion less than k where every part has size less than κ ?*

We suspect that the class of locally finite connected graphs should be a solution for Problem 4.9.4, where κ is necessarily equal to \aleph_0 , since locally finite connected graphs are countable.

Part II.

Directed graphs

5. An analogue of Edmonds' branching theorem for infinite digraphs

5.1. Introduction

Studying how to force spanning structures in finite graphs is a basic task. The most fundamental spanning structure is a spanning tree, whose existence is already characterised by the connectedness of the graph. Moving on and characterising the existence of a given number of edge-disjoint spanning trees via an immediately necessary condition, Nash-Williams [36] and Tutte [51] independently proved the following famous theorem.

Theorem 5.1.1. [36, 51], [9, Theorem 2.4.1] *A finite multigraph G has $k \in \mathbb{N}$ edge-disjoint spanning trees if and only if for every partition \mathcal{P} of $V(G)$ there are at least $k(|\mathcal{P}| - 1)$ edges in G whose endvertices lie in different partition classes.*

Later, Edmonds [19] generalised Theorem 5.1.1 to finite digraphs, also involving a condition which is immediately seen to be necessary for the existence of the spanning structures. In his theorem, Edmonds considers as spanning structures *out-arborescences* rooted in a vertex r , i.e. spanning trees whose edges are directed away from the root r . His theorem immediately implies a corresponding result for *in-arborescences* rooted in r , i.e. spanning trees directed towards r , via reversing every edge in the digraph. For this reason we shall focus in this chapter only on out-arborescences and denote them just by *arborescences*.

Theorem 5.1.2. [19], [4, Theorem 9.5.1] *A finite digraph G with a vertex $r \in V(G)$ has $k \in \mathbb{N}$ edge-disjoint spanning arborescences rooted in r if and only if there are at least k edges from X to Y for every bipartition (X, Y) of $V(G)$ with $r \in X$.*

One of the main results of this chapter is to extend Theorem 5.1.2 to a certain class of infinite digraphs. There has already been work in this area. In order to

mention two important results about this let us call a one-way infinite path all of whose edges are directed away from the unique vertex incident with only one edge a *forwards directed ray*. Similarly, we call the digraph obtained by reversing all edges of a forwards directed ray a *backwards directed ray*. Thomassen [48] extended Theorem 5.1.2 to infinite digraphs that do not contain a backwards directed ray, while Joó [32] obtained an extension for infinite digraphs without forwards directed rays using different methods. In contrast to these two results we shall demand a local property for our digraphs by considering *locally finite digraphs*, i.e. digraphs in which every vertex has finite in- and out-degree. Similarly, undirected multigraphs are called *locally finite* if every vertex has finite degree.

When trying to extend Theorem 5.1.2 to infinite digraphs it is important to know that a complete extension is not possible. The reason for this is that Oxley [38, Example 2] constructed a locally finite graph without two edge-disjoint spanning trees but fulfilling the condition in Theorem 5.1.1. Following up, Aharoni and Thomassen [3, Theorem] gave a construction for further counterexamples to Theorem 5.1.2, which are all locally finite and can even be made $2k$ -edge-connected for arbitrary $k \in \mathbb{N}$. Hence, using ordinary spanning trees for an extension of Theorem 5.1.1 to locally finite graphs does not work. This immediately implies that extending Theorem 5.1.2 to locally finite digraphs fails as well if ordinary arborescences are used. While Thomassen and Joó could overcome this problem by forbidding certain one-way infinite paths, for us it is necessary to additionally change the notion of arborescence since the counterexamples to direct extensions of Theorem 5.1.1 and Theorem 5.1.2 to infinite (di-)graphs are locally finite.

For undirected locally finite (connected) multigraphs G the problem of how to extend Theorem 5.1.1 has successfully been overcome. The key was to not just consider G but the Freudenthal compactification $|G|$ [9, 11] of the 1-complex of G . Instead of ordinary spanning trees, now packings of topological spanning trees of G are considered. We call a connected subspace of $|G|$ which is the closure of a set of edges of G , contains all vertices of G but contains no homeomorphic image of the unit circle $S^1 \subseteq \mathbb{R}^2$, a *topological spanning tree* of G . There is an equivalent but more combinatorial, and in particular finitary, way of defining topological spanning trees of G . They are precisely the closures in $|G|$ of the minimal edge sets that meet every finite cut of G [9]. As already observed by Tutte, this finitary condition can be used to obtain the following packing theorem for disjoint edge

sets each meeting every finite cut, via the compactness principle.

Theorem 5.1.3. [51] *A locally finite multigraph G has $k \in \mathbb{N}$ disjoint edge sets each meeting every finite cut of G if and only if for every finite partition \mathcal{P} of $V(G)$ there are at least $k(|\mathcal{P}| - 1)$ edges in G whose endvertices lie in different partition classes.*

By the equivalence noted above, Theorem 5.1.3 implies a packing result for topological spanning trees:

Theorem 5.1.4. [9, Theorem 8.5.7] *A locally finite multigraph G has $k \in \mathbb{N}$ edge-disjoint topological spanning trees if and only if for every finite partition \mathcal{P} of $V(G)$ there are at least $k(|\mathcal{P}| - 1)$ edges in G whose endvertices lie in different partition classes.*

In the spirit of Tutte's approach, we prove the following packing theorem generalising Theorem 5.1.2 to locally finite digraphs for what we call spanning pseudo-arborescence rooted in some vertex r . For a locally finite weakly connected digraph G and $r \in V(G)$ we define a *spanning pseudo-arborescence rooted in r* as a minimal edge set $F \subseteq E(G)$ such that F contains, for every bipartition (X, Y) of $V(G)$ with $r \in X$ and finitely many edges between X and Y in either direction, an edge directed from X to Y .

Theorem 5.1.5. *A locally finite weakly connected digraph G with $r \in V(G)$ has $k \in \mathbb{N}$ edge-disjoint spanning pseudo-arborescences rooted in r if and only if for every bipartition (X, Y) of $V(G)$ with $r \in X$ and finitely many edges between X and Y in either direction there are at least k edges from X to Y .*

In fact we shall prove a slightly stronger version of this theorem, Theorem 5.4.3, which requires more notation.

While minimal edges sets meeting every finite cut in an undirected multigraph turn out to be topological extensions of finite trees, there is no analogous topological interpretation of spanning pseudo-arborescences in terms of the Freudenthal compactification of the underlying multigraph. In Section 5.5 we give an example of a digraph G with underlying multigraph H for which the closure in $|H|$ of the underlying undirected edges of any spanning pseudo-arborescence of G contains a homeomorphic image of S^1 . We shall be able to extend to pseudo-arborescences,

in a suitable topological setting, the property of finite arborescences of being edge-minimal such that each vertex is still reachable by a directed path from the root. While in finite arborescences such directed paths are unique, however, their analogues in pseudo-arborescences are not in general unique. This will be illustrated by an example given in Section 5.5.

Finally, we prove the following structural characterisation for spanning pseudo-arborescences.

Theorem 5.1.6. *Let G be a locally finite weakly connected digraph and $r \in V(G)$. Then the following statements are equivalent for an edge set $F \subseteq E(G)$ containing, for every bipartition (X, Y) of $V(G)$ with $r \in X$ and finitely many edges between X and Y in either direction, an edge from X to Y .*

- (i) *F is a spanning pseudo-arborescence rooted in r .*
- (ii) *For every vertex $v \neq r$ of G there is a unique edge in F whose head is v , and no edge in F has r as its head.*
- (iii) *For every weak component T of $G[F]$ the following holds: If $r \in V(T)$, then T is an arborescence rooted in r . Otherwise, the underlying multigraph of T is a tree, T contains a backwards directed ray and all other edges of T are directed away from that ray.*

We prove a slightly more general version of Theorem 5.1.6 in Section 5.5 (cf. Theorem 5.5.3).

The structure of this chapter is as follows. In Section 5.2 we give basic definitions and fix our notation for directed and undirected (multi-)graphs. We in particular refer to the topology we consider on locally finite (weakly) connected digraphs and (undirected) multigraphs, and state some basic lemmas that we shall need for our main results. In Section 5.3 we extend some lemmas about directed walks and paths in finite digraphs to locally finite (weakly) connected digraphs. Section 5.4 is dedicated to the proof of Theorem 5.1.5. We complete the chapter in Section 5.5 with the proof of Theorem 5.1.6 and a discussion about how much pseudo-arborescences resemble finite arborescences or topological trees.

5.2. Preliminaries

In this chapter we consider both digraphs and multigraphs.

Recall that for a multigraph or digraph G we call the edge set $E(X, Y)$ a *cut* if (X, Y) is a bipartition of $V(G)$. If we introduce a cut $E(X, Y)$, then we implicitly want (X, Y) to be the corresponding bipartition of $V(G)$ defining the cut.

We define a *finite directed walk* as a tuple $(\mathcal{W}, <_{\mathcal{W}})$ with the following properties:

1. \mathcal{W} is a non-empty weakly connected graph with edge set $\{e_1, e_2, \dots, e_n\}$ for some $n \in \mathbb{N}$ such that the head of e_{i-1} is the tail of e_i for every $i \in \mathbb{N}$ satisfying $2 \leq i \leq n$.
2. $<_{\mathcal{W}}$ is a linear order on $E(\mathcal{W})$ stating that $e_i <_{\mathcal{W}} e_j$ if and only if $i < j$ for all $i, j \in \{1, \dots, n\}$.

Note that the second condition implies that the edges e_1, \dots, e_n are all distinct, i.e. the walk traverses its edges only once. We call a directed walk without edges *trivial* and call its unique vertex its *endvertex*. Otherwise, we call the head of e_1 the *start vertex* of $(\mathcal{W}, <_{\mathcal{W}})$ and the tail of e_n the *endvertex* of $(\mathcal{W}, <_{\mathcal{W}})$. If the start vertex and the endvertex of finite directed walk are equal, we call it *closed*. Lastly, we call $(\mathcal{W}, <_{\mathcal{W}})$ a *finite directed s - t walk* for two vertices $s, t \in V(\mathcal{W})$ if s is the start vertex of $(\mathcal{W}, <_{\mathcal{W}})$ and t is the endvertex of $(\mathcal{W}, <_{\mathcal{W}})$. We might call a finite graph \mathcal{W} a finite directed walk and implicitly assume that there exists a linear order $<_{\mathcal{W}}$, which we then also fix, such that $(\mathcal{W}, <_{\mathcal{W}})$ is a finite directed walk. In particular, we will say that a finite directed walk $(\mathcal{W}, <_{\mathcal{W}})$ is contained in a graph G' if \mathcal{W} is a subgraph of G' . Note that directed paths are directed walks when equipped with the obviously suitable linear order.

We call a digraph A an *out-arborescence rooted in r* if $r \in V(A) \cup \Omega(A)$ and the underlying multigraph of A is a tree such that $d^-(v) = 1$ holds for every vertex $v \in V(A) \setminus \{r\}$ and additionally $d^-(r) = 0$ in the case that $r \in V(A)$, while we demand that r contains a backwards directed ray if $r \in \Omega(A)$.

Note that if $r \in V(A)$, then A does not contain a backwards directed ray. In the case where $r \in \Omega(A)$, then r is the unique end of A containing a backwards directed ray, since a second one would yield a vertex with in-degree bigger than 1 by using that the underlying multigraph of A is a tree. Also note that if A is a finite digraph, the condition $d^-(r) = 0$ for $r \in V(A)$ in the definition of an out-

arborescence rooted in r is redundant, because it is implied by the tree structure of A .

Similarly, an *in-arborescence rooted in r* is defined with d^- replaced by d^+ . Corresponding results about in-arborescences are immediate by reversing the orientations of all edges. For both types of arborescences we call r the *root* of the arborescence. In this chapter we shall only work with out-arborescences. Hence, we shall drop the prefix ‘out’ and just write arborescence from now on.

For a vertex set X in a locally finite connected multigraph G we define its *combinatorial closure* $\overline{X} \subseteq V(G) \cup \Omega(G)$ as the set X together with all ends of G that contain a ray which we cannot separate from X by finitely many vertices. Note that for a finite cut $E(X, Y)$ of G we obtain that $(\overline{X}, \overline{Y})$ is a bipartition of $V(G) \cup \Omega(G)$, because every end in \overline{X} can be separated from Y by the finitely many vertices of X that are incident with edges of $E(X, Y)$, and, furthermore, each ray contains a subray that is either completely contained in X or in Y since $E(X, Y)$ is finite. The *combinatorial closure* of a vertex set in a digraph is just defined as the combinatorial closure of that set in the underlying undirected multigraph.

Let G be a locally finite digraph and $Z \subseteq V(G) \setminus \{r\}$ with $r \in V(G) \cup \Omega(G)$. An edge set $F \subseteq E(G)$ is called *r -reachable for Z* if $|F \cap \overrightarrow{E}(X, Y)| \geq 1$ holds for every finite cut $E(X, Y)$ of G with $r \in \overline{X}$ and $Y \cap Z \neq \emptyset$. Furthermore, if F is an r -reachable set for $V(G) \setminus \{r\}$, we call F a *spanning r -reachable set*. Note that a spanning r -reachable set spans $V(G)$ as an edge set. We continue with a very basic remark about spanning r -reachable sets.

Remark 5.2.1. Let G be a locally finite digraph with a spanning r -reachable set F with $r \in V(G) \cup \Omega(G)$. Then $|F \cap \overrightarrow{E}(V(G) \setminus M, M)| \geq 1$ holds for every non-empty finite set $M \subseteq V(G)$ with $r \notin M$.

Proof. Since G is locally finite and M is finite, the cut $E(V(G) \setminus M, M)$ is finite as well. The assumption $r \notin M$ ensures that $r \in \overline{V(G) \setminus M}$. Using that F is a spanning r -reachable set and that M , as a non-empty set, contains a vertex different from r , we get the desired inequality $|F \cap \overrightarrow{E}(V(G) \setminus M, M)| \geq 1$ by the definition of spanning r -reachable sets. \square

Note that for a locally finite digraph G with a spanning r -reachable set F the digraph $G[F]$ is spanning. This follows by applying Remark 5.2.1 to the

set $M := \{v\}$ for every vertex $v \in V(G)$. Furthermore, note that if G is finite, the subgraph induced by a spanning r -reachable set contains a spanning arborescence rooted in $r \in V(G)$.

We conclude this section with a last definition. We call an inclusion-wise minimal r -reachable set F for a set $Z \subseteq V(G) \setminus \{r\}$ a *pseudo-arborescence for Z rooted in r* . Moreover, if F is spanning, i.e. $Z = V(G) \setminus \{r\}$, we call it a *spanning pseudo-arborescence rooted in r* .

5.2.1. Topological notions for undirected multigraphs

For this subsection let $G = (V, E)$ denote a locally finite connected multigraph. We can endow G together with its ends with a topology which yields the topological space $|G|$. A precise definition of $|G|$ for locally finite connected simple graphs can be found in [9, Chapter 8.5]. However, this concept and definition directly extends to locally finite connected multigraphs. For a better understanding we should point out here that a ray of G converges in $|G|$ to the end of G that it is contained in. An equivalent way of describing $|G|$ is by first endowing G with the topology of a 1-complex and then compactifying this space using the Freudenthal compactification [15].

For an edge $e \in E$ let \mathring{e} denote the set of points in $|G|$ that correspond to inner points of the edge e . For an edge set $F \subseteq E$ we define $\mathring{F} = \bigcup \{\mathring{e} \mid e \in F\} \subseteq |G|$. Given a point set X in $|G|$, we denote the *closure* of X in $|G|$ by \overline{X} . To ease notation we shall also use this notation when X denotes an edge set or a subgraph of G , meaning that we apply the closure operator to the set of all points in $|G|$ that correspond to X . Note that for a vertex set its closure coincides with its combinatorial closure in locally finite connected multigraphs. Hence, we shall use the same notation for these two operators. Furthermore we call a subspace $Z \subseteq |G|$ *standard* if $Z = \overline{H}$ for some subgraph H of G .

Let $W \subseteq |G|$ and $<_W$ be a linear order on $\mathring{E} \cap W$. We call the tuple $(W, <_W)$ a *topological walk* in $|G|$ if there exists a continuous map $\sigma : [0, 1] \rightarrow |G|$ such that the following hold:

1. W is the image of σ ,
2. each point $p \in \mathring{E} \cap W$ has precisely one preimage under σ , and

3. the linear order $<_W$ equals the linear order $<_\sigma$ on $\dot{E} \cap W$ defined via $p <_\sigma q$ if and only if $\sigma^{-1}(p) <_{\mathbb{R}} \sigma^{-1}(q)$, where $<_{\mathbb{R}}$ denotes the natural linear order of the reals.

We call such a map σ a *witness* of $(W, <_W)$. When we talk about a topological walk $(W, <_W)$ we shall often omit stating its linear order $<_W$ explicitly and just refer to the topological walk by writing W . In particular, we might say that a topological walk $(W, <_W)$ is contained in some subspace X of $|G|$ if $W \subseteq X$ holds. Furthermore, we call a point x of $|G|$ an *endpoint* of W if 0 or 1 is mapped to x by a witness of W . Note that this definition is independent of the particular witness. Similar to finite walks in graphs we call an endpoint x of W an *endvertex* of W if x corresponds to a vertex of G . Furthermore, we denote W as an x - y topological walk, if x and y are endpoints of W . If W has just one endpoint, which then has to be an end or a vertex by definition, we call it *closed*. Note that an x - y topological walk is a standard subspace for any $x, y \in V \cup \Omega(G)$. We say that a witness σ of a topological walk W *pauses* at a vertex $v \in V$ if the preimage of v under σ is a disjoint union of closed nontrivial intervals.

We define an *arc* in $|G|$ as the image of a homeomorphism mapping into $|G|$ and with the closed real unit interval $[0, 1] \subseteq \mathbb{R}$ as its domain. Note that arcs in $|G|$ are also topological walks in $|G|$ if we equip them with a suitable linear order, of which there exist only two. Since the choice of such a linear order does not change the set of endpoints of the arc if we then consider it as a topological walk, we shall use the notion of endpoints and endvertices also for arcs. Furthermore, note that finite paths of G which contain at least one edge correspond to arcs in $|G|$, but again there might be infinite subgraphs, for example rays, whose closures form arcs in $|G|$. We now call a subspace X of $|G|$ *arc-connected* if there exists an x - y arc in X for any two points $x, y \in X$.

Lastly, we define a *circle* in $|G|$ as the image of a homeomorphism mapping into $|G|$ and with the unit circle $S^1 \subseteq \mathbb{R}^2$ as its domain. We might also consider any circle as a closed topological walk if we equip it with a suitable linear order, which, however, depends on the point on the circle that we choose as the endpoint for the closed topological walk, and on choosing one of the two possible orientations of S^1 . Similarly as for finite paths, note that finite cycles in G correspond to circles in $|G|$, but there might be infinite subgraphs of G whose closures are circles in $|G|$.

as well.

Using these definitions we can now formulate a topological extension of the notion of trees. We define a *topological tree* in $|G|$ as an arc-connected standard subspace of $|G|$ that does not contain any circle. Note that in a topological tree there is a unique arc between any two points of the topological tree, which resembles a property of finite trees with respect to vertices and finite paths. Furthermore, we denote by a *topological spanning tree of G* a topological tree in $|G|$ that contains all vertices of G . Since topological spanning trees are closed subspaces of $|G|$, they need to contain all ends of G as well.

5.2.2. Topological notions for digraphs

In this subsection we extend some of the notions of the previous subsection to directed graphs. Throughout this subsection let G denote a locally finite weakly connected digraph and let H denote its underlying multigraph. We define the topological space $|G|$ as $|H|$. Additionally, every edge $e = uv \in E(G)$ defines a certain linear order $<_e$ on $\overline{\{e\}} \subseteq |G|$ via its direction. For the definition of $<_e$ we first take any homeomorphism $\varphi_e : [0, 1] \rightarrow \overline{\{e\}} \subseteq |G|$ with $\varphi_e(0) = u$ and $\varphi_e(1) = v$. Now we set $p <_e q$ for arbitrary $p, q \in \overline{\{e\}}$ if $\varphi_e^{-1}(p) <_{\mathbb{R}} \varphi_e^{-1}(q)$ where $<_{\mathbb{R}}$ is the natural linear order on the real numbers. Note that the definition of $<_e$ does not depend on the choice of the homeomorphism φ_e .

Let $(W, <_W)$ be a topological walk in $|G|$ with witness σ . We call $(W, <_W)$ *directed* if $<_e \upharpoonright \dot{e}$ equals $<_W \upharpoonright \dot{e}$ for every edge $e \in E(G)$ with $\dot{e} \cap W \neq \emptyset$. If $(W, <_W)$ is directed and $\sigma(0) = s \neq t = \sigma(1)$ for $s, t \in |G|$, then there is no linear order $<'_W$ such that $(W, <'_W)$ is a directed topological walk with a witness σ' satisfying $\sigma'(0) = t$ and $\sigma'(1) = s$, because every topological s - t walk uses inner points of some edge. Hence, if we consider a directed topological s - t walk $(W, <_W)$ for $s, t \in |G|$, we implicitly assume that $\sigma(0) = s \neq t = \sigma(1)$ holds for every witness σ of $(W, <_W)$.

As arcs and circles can be seen as special instances of topological walks, *directed arcs* and *directed circles* are analogously defined. Note that if we can equip an arc with a suitable linear order such that it becomes a directed topological walk, then this linear order is unique. Hence, when we call an arc directed we implicitly associate this unique linear order with it.

5.2.3. Basic lemmas

We shall heavily work with the topological space $|G|$ of a locally finite multigraph G appearing as the underlying graph of digraphs we consider. Therefore, we shall make use of some basic statements and properties of the space $|G|$, in particular those involving connectivity. Although the following lemmas are only stated for locally finite graphs, their proofs immediately extend to locally finite multigraphs.

Proposition 5.2.2. [9, Lemma 8.5.1] *If G is a locally finite connected multigraph, then $|G|$ is a compact Hausdorff space.*

The next lemma is essential for decoding the topological property of arc-connectedness of standard subspaces of $|G|$ into a combinatorial one.

Lemma 5.2.3. [9, Lemma 8.5.3] *Let G be a locally finite connected multigraph and $F \subseteq E(G)$ be a cut with sides V_1 and V_2 .*

(i) *If F is finite, then $\overline{V_1} \cap \overline{V_2} = \emptyset$, and there is no arc in $|G| \setminus \overset{\circ}{F}$ with one endpoint in V_1 and the other in V_2 .*

(ii) *If F is infinite, then $\overline{V_1} \cap \overline{V_2} \neq \emptyset$, and there may be such an arc.*

Note that for a finite cut $E(X, Y)$ of G we obtain that $(\overline{X}, \overline{Y})$ is a bipartition of $V(G) \cup \Omega(G)$.

The following lemma captures the equivalence of arc-connectedness and connectedness for standard subspaces of $|G|$.

Lemma 5.2.4. [9, Lemma 8.5.4] *If G is a locally finite connected multigraph, then every connected standard subspace of $|G|$ is arc-connected.*

We conclude with a convenient lemma which combines the essences of the previous two.

Lemma 5.2.5. [9, Lemma 8.5.5] *If G is a locally finite connected multigraph, then a standard subspace of $|G|$ is connected if and only if it contains an edge from every finite cut of G of which it meets both sides.*

5.3. Fundamental statements about topological directed walks in locally finite digraphs

In this section we lift several facts about topological walks and arcs to their directed counterparts. Most of the involved techniques and proof ideas are similar to the ones used in undirected locally finite connected multigraphs. Nevertheless, because of overlying directed structure on the multigraph, some adjustments and additional arguments are needed in the proofs. We start with a statement that combinatorially characterises the existence of directed topological walks in a standard subspace via finite cuts.

Lemma 5.3.1. *Let G be a locally finite weakly connected digraph, let $s, t \in V(G) \cup \Omega(G)$ with $s \neq t$ and let $F \subseteq E(G)$. Then the following statements are equivalent.*

- (i) \overline{F} contains a directed topological s - t walk.
- (ii) $|F \cap \overrightarrow{E}(X, Y)| \geq 1$ for every finite cut $E(X, Y)$ of G with $s \in \overline{X}$ and $t \in \overline{Y}$.
- (iii) There is a subset $W \subseteq F$ such that $|W \cap \overrightarrow{E}(X, Y)| = |W \cap \overrightarrow{E}(Y, X)| + 1$ for every finite cut $E(X, Y)$ of G with $s \in \overline{X}$ and $t \in \overline{Y}$.

Proof. First we prove the implication from (i) to (iii). Let $E(X, Y)$ be any finite cut of G with $s \in \overline{X}$ and $t \in \overline{Y}$. Since \overline{F} contains a directed topological s - t walk $(\overline{W}, <_{\overline{W}})$ for an edge set $W \subseteq E(G)$, we know that $F \cap E(X, Y) \neq \emptyset$ by Lemma 5.2.5. Note furthermore that $\overline{X} \cap \overline{Y} = \emptyset$ by Lemma 5.2.3. As \overline{X} and \overline{Y} are closed and $|G|$ is compact by Proposition 5.2.2, we get that \overline{X} and \overline{Y} are compact too. Now let φ be a witness of \overline{W} . Since \overline{Y} is compact and φ is continuous, there exists a smallest number $q \in [0, 1]$ such that $\varphi(q) \in \overline{Y}$. Furthermore, there is a biggest number $p \in [0, q]$ such that $\varphi(p) \in \overline{X}$. Note that $p \neq q$ since $\overline{X} \cap \overline{Y} = \emptyset$. Now let $M := \{\varphi(r) \in |G| \mid p < r < q\}$. Obviously, M contains only inner points of edges in $E(X, Y)$. Since M is connected, we obtain $M = \mathring{e}$ for some edge $e \in E(X, Y)$. Using that $<_{\overline{W}} \upharpoonright \mathring{e}$ equals $<_e \upharpoonright \mathring{e}$ because $(\overline{W}, <_{\overline{W}})$ is a directed s - t walk, we see that $e \in W \cap \overrightarrow{E}(X, Y)$. By the continuity of φ , we get that $\varphi(q) = y$ for some vertex $y \in Y$. If $|W \cap \overrightarrow{E}(Y, X)| = 0$, we know that $|W \cap E(X, Y)| = 1$ since $\varphi \upharpoonright [q, 1]$ is connected and hence a subset of \overline{Y} , and we are done. Otherwise, consider $\varphi \upharpoonright [q, 1]$, which is a witness for $(\overline{Q}, <_{\overline{Q}})$ being a directed y - t walk where

$Q = \{e \in W ; \forall a \in \dot{e} : \varphi^{-1}(a) > q\}$. Note that since $|W \cap \vec{E}(Y, X)| > 0$ we get that $|Q \cap \vec{E}(Y, X)| > 0$ as well by the choice of q . Therefore, \overline{Q} also contains an element of \overline{X} . Similarly as before, let $p' \in [q, 1]$ denote the smallest number such that $\varphi(p') \in \overline{X}$ and $q' \in [q, p']$ denote the biggest number such that $\varphi(q') \in \overline{Y}$. Now considering the set $M' := \{\varphi(r) \in |G| \mid q' < r < p'\}$ we obtain as before that $M' = \mathring{f}$ for some edge $f \in E(X, Y)$. More precisely, since Q is a directed x - t walk, we get that $f \in Q \cap \vec{E}(Y, X)$. Finally, we consider the directed $\varphi(p')$ - t walk $(\overline{P}, <_{\overline{P}})$ with witness $\varphi \upharpoonright [p', 1]$ where $P = \{e \in W \mid \forall a \in \dot{e} : \varphi^{-1}(a) > p'\}$. By the previous observations we know that $|P \cap \vec{E}(X, Y)| = |W \cap \vec{E}(X, Y)| - 1$ and $|P \cap \vec{E}(Y, X)| = |W \cap \vec{E}(Y, X)| - 1$ hold. Using that $E(X, Y)$ contains only finitely many edges, we inductively get that $|W \cap \vec{E}(X, Y)| = |W \cap \vec{E}(Y, X)| + 1$ is true.

The implication from (iii) to (ii) is immediate.

It remains to show that (ii) implies (i). For this we first fix a sequence $(S_n)_{n \in \mathbb{N}}$ of finite sets $S_n \subseteq V(G)$ such that $S_n \subsetneq S_{n+1}$ for every $n \in \mathbb{N}$ and $\bigcup_{n \in \mathbb{N}} S_n = V(G)$. For every $n \in \mathbb{N}$ let G_n denote the digraph which arises by contracting $E(G - S_n)$ in G . Since G is locally finite, we know that each G_n is a finite digraph. We call the vertices of G_n that are not contained in S_n *dummy vertices*. Note that each dummy vertex of G_n corresponds to a unique weak component of $G - S_n$.

If some $v \in V(G) \cup \Omega(G)$ is not contained in S_n , there exists a unique component C_n of $G - S_n$ such that $v \in \overline{C_n}$. This is obviously true if v is a vertex of G , but also holds if v is an end of G . To see the latter statement suppose $v \in \Omega(G)$ is contained in $\overline{C_n}$ for a component C_n of $G - S_n$. Then the cut $E(V(C_n), V(G) \setminus V(C_n))$ is finite as S_n is finite and G is locally finite. Hence $\overline{V(C_n)} \cap (\overline{V(G)} \setminus \overline{V(C_n)}) = \emptyset$ by Lemma 5.2.3, which means that v cannot lie in the closure of another component of $G - S_n$. With a slight abuse of notation, we refer to the dummy vertex of G_n corresponding to C_n as v .

Since for each $n \in \mathbb{N}$ every cut of G_n corresponds to a finite cut of G , we obtain by Theorem 5.1.2 that $F \cap E(G_n)$ contains the edge set of a finite directed s - t walk in the digraph G_n . Moreover, any finite directed s - t walk $(\mathcal{W}_{n+1}, <_{\mathcal{W}_{n+1}})$ in G_{n+1} induces a finite directed s - t walk $(\mathcal{W}_n, <_{\mathcal{W}_n})$ in G_n via $E(\mathcal{W}_n) := E(\mathcal{W}_{n+1}) \cap E(G_n)$ and defining $<_{\mathcal{W}_n}$ as $<_{\mathcal{W}_{n+1}} \upharpoonright E(\mathcal{W}_n)$. Note that each maximal interval with respect to $<_{\mathcal{W}_{n+1}}$ of $E(\mathcal{W}_{n+1}) \setminus E(\mathcal{W}_n)$ corresponds to some v - w walk where v and w are the same dummy vertex of G_n . Hence each time a dummy vertex of G_n

appears as the head of some edge $e \in E(\mathcal{W}_n)$ there is a corresponding, possibly trivial, walk \mathcal{W}_{n+1}^e using edges of of such a maximal interval with the induced order $<_{\mathcal{W}_{n+1}} \upharpoonright E(\mathcal{W}_{n+1}^e)$.

For every $n \in \mathbb{N}$ let V_n denote the set of all finite directed s - t walks in G_n that use only edges from F . Obviously, each set V_n is finite as G_n is a finite digraph. By the previously given arguments, none of the sets V_n is empty and each element of V_{n+1} induces one of V_n . Hence, we get a sequence $((\mathcal{W}_n, <_{\mathcal{W}_n}))_{n \in \mathbb{N}}$ of finite directed s - t walks with $(\mathcal{W}_n, <_{\mathcal{W}_n}) \in V_n$ such that $E(\mathcal{W}_{n+1}) \cap E(\mathcal{W}_n) = E(\mathcal{W}_n)$ and $<_{\mathcal{W}_{n+1}} \upharpoonright E(\mathcal{W}_n)$ equals $<_{\mathcal{W}_n}$ for every $n \in \mathbb{N}$ by Königs Infinity Lemma. We define $W_n := E(\mathcal{W}_n)$ for every $n \in \mathbb{N}$, and set $W := \bigcup_{n \in \mathbb{N}} W_n$ and $<_W := \bigcup_{n \in \mathbb{N}} <_{\mathcal{W}_n}$. Furthermore, we define a linear order $<_{\overline{W}}$ on \overline{W} as follows for $p, q \in \overline{W}$ with $p \neq q$:

$$p <_{\overline{W}} q \text{ if } \begin{cases} p \in \mathring{e} \text{ and } q \in \mathring{f} \text{ with } e <_W f \text{ for some } e, f \in W \text{ with } e \neq f, \text{ or} \\ p, q \in \mathring{e} \text{ and } p <_e q \text{ for some } e \in W. \end{cases}$$

Now we claim that $(\overline{W}, <_{\overline{W}})$ is a directed topological s - t walk in $|G|$. In order to show this we first have to define a witness φ for $(\overline{W}, <_{\overline{W}})$. We shall obtain φ as a limit of countably many certain witnesses φ_n of directed topological walks $(\overline{W}_n, <_{\overline{W}_n})$ in $|G_n|$ that we define inductively, where $<_{\overline{W}_n}$ is defined analogously as $<_{\overline{W}}$ but with respect to W_n .

We start with a witness φ_0 of the directed topological s - t walk $(\overline{W}_0, <_{\overline{W}_0})$ in $|G_0|$ which pauses at every dummy vertex of G_0 contained in \overline{W}_0 .

Now suppose that the witness φ_n of $(\overline{W}_n, <_{\overline{W}_n})$ has already been defined such that it pauses at every dummy vertex of G_n that is contained in \overline{W}_n . Then we define φ_{n+1} as some witness of $(\overline{W}_{n+1}, <_{\overline{W}_{n+1}})$ as follows. For every edge $e \in W_n$ whose head is a dummy vertex of G_n , let W_{n+1}^e be the edge set of the walk \mathcal{W}_{n+1}^e as above and let φ_{n+1}^e be a witness that \overline{W}_{n+1}^e is the corresponding directed topological walk that pauses at every dummy vertex of G_{n+1} that is contained in \overline{W}_{n+1}^e . Starting with φ_n , each time we enter some dummy vertex d of G_n by an edge e , we replace the image of the interval that is mapped to d with a rescaled version of φ_{n+1}^e .

Using the maps φ_n we are able to define φ as follows: For every $q \in [0, 1]$ for which there exists an $n \in \mathbb{N}$ such that $\varphi_n(q) \in |G[S_n]| \subseteq |G_n|$, we set $\varphi(q) := \varphi_n(q)$. Otherwise, $\varphi_n(q)$ corresponds to a contracted component C_n of $G - S_n$ for every $n \in \mathbb{N}$. Since $S_n \subsetneq S_{n+1}$ for every $n \in \mathbb{N}$ and $\bigcup_{n \in \mathbb{N}} S_n = V(G)$, it is easy to

check that $\bigcap_{n \in \mathbb{N}} \overline{C_n} = \{\omega\}$ for some end ω of G . In this case, we define $\varphi(q) := \omega$. This completes the definition of φ . It is straightforward to verify that φ is continuous and also onto \overline{W} because each φ_n is onto $\overline{W_n}$ and $W := \bigcup_{n \in \mathbb{N}} W_n$. This ensures that it is a witness of $(\overline{W}, <_{\overline{W}})$ being a topological s - t walk. Note that the linear order $<_{\overline{W}} \upharpoonright e$ equals $<_e \upharpoonright e$ for each edge $e \in W$ since each linear order $<_{\overline{W_n}}$ has this property. Hence, φ witnesses that $(\overline{W}, <_{\overline{W}})$ is a directed topological s - t walk in $|G|$ with $W \subseteq F$. \square

We proceed with a lemma which gives a combinatorial description for a standard subspace to be a directed arc.

Lemma 5.3.2. *Let G be a locally finite weakly connected digraph, let $s, t \in V(G) \cup \Omega(G)$ with $s \neq t$ and let $A \subseteq E(G)$. Then the following statements are equivalent:*

- (i) \overline{A} is a directed s - t arc.
- (ii) A is inclusion-wise minimal such that $|A \cap \overrightarrow{E}(X, Y)| \geq 1$ holds for every finite cut $E(X, Y)$ of G with $s \in \overline{X}$ and $t \in \overline{Y}$.
- (iii) A is inclusion-wise minimal such that $|A \cap \overrightarrow{E}(X, Y)| = |A \cap \overrightarrow{E}(Y, X)| + 1$ holds for every finite cut $E(X, Y)$ of G with $s \in \overline{X}$ and $t \in \overline{Y}$.

Proof. First we show the implication from (i) to (iii). As \overline{A} is a directed s - t arc, it is also a directed topological s - t walk. So by Lemma 5.3.1, we only need to check the minimality of A for property (iii). Since \overline{A} is an s - t arc, we know that s and t are in different topological components of $\overline{A \setminus \{e\}}$ for any edge $e \in A$. So no proper subset of A has the property that its closure in $|G|$ contains a directed topological s - t walk. Again by Lemma 5.3.1 we know that no proper subset of A satisfies statement (iii) of Lemma 5.3.1. This proves the minimality of A and hence statement (iii).

Next let us verify that (iii) implies (ii). Assume for a contradiction that statement (iii) holds, but (ii) does not. Then there must exist a proper subset $A' \subsetneq A$ meeting $\overrightarrow{E}(X, Y)$ for every finite cut $E(X, Y)$ of G with $s \in \overline{X}$ and $t \in \overline{Y}$. By Lemma 5.3.1 we get that A' satisfies also statement (iii) of Lemma 5.3.1. This contradicts the minimality of A .

It remains to prove the implication from (ii) to (i). By assuming (ii) we know from Lemma 5.3.1 that \overline{A} contains a directed topological s - t walk and by the minimality

of A we know that \overline{A} is in fact a directed topological s - t walk, say witnessed by $\varphi : [0, 1] \rightarrow |G|$. Now suppose for a contradiction that \overline{A} is not a directed s - t arc. Then there exists a point $a \in V(G) \cup \Omega(G)$ that spoils injectivity for φ . Note that \overline{A} is compact because it is a closed set in $|G|$, which is a compact space by Proposition 5.2.2. Since φ is continuous and \overline{A} is compact, there exists a smallest number $x \in [0, 1]$ and a largest number $y \in [0, 1]$ such that $\varphi(x) = \varphi(y) = a$. We obtain from this that the image of $\varphi|_{[0, x]}$ is a directed topological s - a walk and the image of $\varphi|_{[y, 1]}$ is a directed topological a - t walk. Concatenating these two walks yields another directed topological s - t walk, which is the closure in $|G|$ of some edge set $A' \subseteq A$. Knowing that $x \neq y$, we get that $A' \subsetneq A$ since the image of $\varphi|_{[x, y]}$ contains points that correspond to inner points of edges. This is a contradiction to the minimality of A . \square

We conclude this section with the following corollary which allows us to extract a directed s - t arc from a directed topological s - t walk for distinct points s, t of $|G|$.

Corollary 5.3.3. *Let $s, t \in V(G) \cup \Omega(G)$ with $s \neq t$ for some locally finite weakly connected digraph G . Then every directed topological s - t walk in $|G|$ contains a directed s - t arc.*

Proof. Let \overline{W} be a directed topological s - t walk with $W \subseteq E(G)$. So W has property (ii) of Lemma 5.3.1. Now consider the set \mathcal{W} of all subsets of W that also have property (ii) of Lemma 5.3.1. This set is ordered by inclusion and not empty since $W \in \mathcal{W}$. Next let us check that every decreasing chain $C \subseteq \mathcal{W}$ is bounded from below by $\bigcap C$, which is an element of \mathcal{W} . Obviously, $\bigcap C \subseteq c$ holds for every $c \in C$. To see that $\bigcap C$ is an element of \mathcal{W} note that for every finite cut $E(X, Y)$ of G with $s \in \overline{X}$ and $t \in \overline{Y}$ there exists a final segment C' of the decreasing chain C such that all $c \in C'$ contain the same edges from $E(X, Y)$. As every $c \in C$ has also at least one edge from $E(X, Y)$, we know that the same is true for $\bigcap C$, which shows that $\bigcap C \in \mathcal{W}$ holds. Now Zorn's Lemma implies that \mathcal{W} has a minimal element, which is a directed s - t arc by Lemma 5.3.2. \square

5.4. Packing pseudo arborescences

We begin this section with a lemma characterising when a packing of $k \in \mathbb{N}$ many edge-disjoint spanning r -reachable sets is possible in a locally finite weakly

connected digraph G with $r \in V(G) \cup \Omega(G)$. This lemma is the main ingredient to prove our first main result. The proof is mainly based on a compactness argument.

Lemma 5.4.1. *A locally finite weakly connected digraph G with $r \in V(G) \cup \Omega(G)$ has $k \in \mathbb{N}$ edge-disjoint spanning r -reachable sets if and only if every bipartition (X, Y) of $V(G)$ with $r \in \bar{X}$ and $|E(X, Y)| < \infty$ satisfies $d^-(Y) \geq k$.*

Proof. The condition that every bipartition (X, Y) of $V(G)$ with $r \in \bar{X}$ and $|E(X, Y)| < \infty$ satisfies $d^-(Y) \geq k$ is obviously necessary for the existence of k edge-disjoint spanning r -reachable sets.

Let us now prove the converse. First we fix a sequence $(S_n)_{n \in \mathbb{N}}$ of finite vertex sets $S_n \subseteq V(G)$ such that $\bigcup_{n \in \mathbb{N}} S_n = V(G)$. For every $n \in \mathbb{N}$ let G_n denote the digraph which arises by contracting, inside of G , each weak component of $G - S_n$ to a single vertex. Here we keep multiple edges, but delete loops that arise. Since G is locally finite, we know that each G_n is a finite digraph.

Note that, as in the proof of Lemma 5.3.1, if $r \notin S_n$, there exists a unique component C_n of $G - S_n$ such that $r \in \bar{C}_n$ and we refer to the vertex of G_n corresponding to C_n as r .

Now we define V_n as the set of all k -tuples consisting of k edge-disjoint spanning r -reachable sets of G_n . As every cut of G_n is finite and also corresponds to a cut of G , our labelling with r ensures that each G_n has k edge-disjoint arborescences rooted in r by Theorem 5.1.2. So none of the V_n is empty. Furthermore, each V_n is finite as G_n is a finite digraph.

Next we show that every spanning r -reachable set F_{n+1} of G_{n+1} induces one for G_n via $F_n := F_{n+1} \cap E(G_n)$. So let F_{n+1} be given and consider a cut $E(X_n, Y_n)$ of G_n with $r \in X_n$. As each component of $G - S_{n+1}$ is contained in a component of $G - S_n$, we can find a cut $E(X_{n+1}, Y_{n+1})$ of G_{n+1} with $r \in X_{n+1}$ such that $\vec{E}(X_n, Y_n) = \vec{E}(X_{n+1}, Y_{n+1})$ (and in fact also $\vec{E}(Y_n, X_n) = \vec{E}(Y_{n+1}, X_{n+1})$). Since F_{n+1} is a spanning r -reachable set of G_{n+1} , we obtain that F_n is one of G_n .

Now we can apply König's Infinity Lemma to the graph defined on the vertex set $\bigcup_{n \in \mathbb{N}} V_n$ where two vertices $v_{n+1} \in V_{n+1}$ and $v_n \in V_n$ are adjacent if the i -th spanning

r -reachable set in v_n is induced by the i -th one of v_{n+1} for every i with $1 \leq i \leq k$. We obtain a ray $r_0 r_1 \dots$ with $r_n \in V_n$ and set $\mathcal{F} := (F^1, \dots, F^k)$ where $F^i := \bigcup_{n \in \mathbb{N}} r_n^i$ and r_n^i denotes the i -th entry of the k -tuple r_n for every i with $1 \leq i \leq k$. Let us

now check that each F^i is a spanning r -reachable set of G . As $\bigcup_{n \in \mathbb{N}} S_n = V(G)$ holds, we can find for every finite cut $E(X, Y)$ of G with $r \in \overline{X}$ an $n \in \mathbb{N}$ such that all endvertices of edges of $E(X, Y)$ are contained in S_n . Hence, there exists a cut $E(X_n, Y_n)$ of G_n with $r \in X_n$ such that $\vec{E}(X_n, Y_n) = \vec{E}(X, Y)$ and $\vec{E}(Y_n, X_n) = \vec{E}(Y, X)$. Since each F^i contains the edges of r_n^i , which is a spanning r -reachable set of G_n and, therefore, contains an edge of $\vec{E}(X_n, Y_n)$, we know that each F^i is a spanning r -reachable set of G . Finally, we get that all the F^i are pairwise edge-disjoint since for every $n \in \mathbb{N}$ the r_n^i are pairwise edge-disjoint. \square

The next lemma ensures the existence of pseudo-arborescences for a set $Z \subseteq V(G) \setminus \{r\}$ in the sense that every r -reachable set for Z contains one. The proof of this lemma works by an application of Zorn's Lemma and is very similar to the proof of Corollary 5.3.3. Therefore, we omit stating its proof.

Lemma 5.4.2. *Let G be a locally finite weakly connected digraph, $Z \subseteq V(G) \setminus \{r\}$ with $r \in V(G) \cup \Omega(G)$. Then every r -reachable set for Z in G contains a pseudo-arborescence for Z rooted in r .* \square

Combining Lemma 5.4.1 and Lemma 5.4.2 with $Z = V(G) \setminus \{r\}$ we now obtain one of our main results, Theorem 5.1.5, which we now state in a slightly stronger version.

Theorem 5.4.3. *A locally finite weakly connected digraph G with $r \in V(G) \cup \Omega(G)$ has $k \in \mathbb{N}$ edge-disjoint spanning pseudo-arborescences rooted in r if and only if every bipartition (X, Y) of $V(G)$ with $r \in \overline{X}$ and $|E(X, Y)| < \infty$ satisfies $d^-(Y) \geq k$.* \square

5.5. Structure of pseudo-arborescences

The following lemma characterises r -reachable sets in terms of directed arcs. Additionally, it justifies the naming of r -reachable sets.

Lemma 5.5.1. *Let G be a locally finite weakly connected digraph, let $F \subseteq E(G)$ and $Z \subseteq V(G) \setminus \{r\}$ and let $r \in V(G) \cup \Omega(G)$. Then F is an r -reachable set for Z in G if and only if there exists a directed r - z arc inside \overline{F} for every $z \in \overline{Z}$.*

Proof. Let us first assume that F is an r -reachable set for Z in G . We fix some $z \in \overline{Z}$ and prove that $|F \cap \vec{E}(X, Y)| \geq 1$ holds for each finite cut $E(X, Y)$

with $r \in \overline{X}$ and $z \in \overline{Y}$. If z is a vertex, this follows immediately from the definition of an r -reachable set for Z . In the case that $z \in \Omega(G)$, we also get that some vertex of Z lies in Y . This follows, because z is contained in the closed and, therefore, compact set \overline{Z} , which implies the existence of a sequence S of vertices in Z converging to z . Since $E(X, Y)$ is a finite cut and $z \in \overline{Y}$, the set $O := |G| \setminus \overline{X \cup E(X, Y)}$ is open, contained in $\overline{G[Y]}$ and contains z by Lemma 5.2.3. Now O must contain infinitely many vertices of S and hence Y must do so as well. Therefore, the desired inequality follows again by the definition of an r -reachable set for Z .

Now we are able to use Lemma 5.3.1, which yields that \overline{F} contains a directed topological r - z walk. We complete the argument by applying Corollary 5.3.3 telling us that \overline{F} contains also a directed r - z arc.

Conversely, we consider any finite cut $E(X, Y)$ with $r \in \overline{X}$ and $Y \cap Z \neq \emptyset$, say $z \in Y \cap Z$. The assumption ensures the existence of a directed r - z arc in \overline{F} . By Lemma 5.3.2 we obtain that $|F \cap \overrightarrow{E}(X, Y)| \geq 1$ holds as desired. \square

Now let us turn our attention towards spanning pseudo-arborescences rooted in some vertex or end in a locally finite weakly connected digraph. The question arises how similarly these objects behave compared to spanning arborescences rooted in some vertex in a finite graph. A basic property of finite arborescences is the existence of a unique directed path in the arborescence from the root to any other vertex of the graph. Closely related is the absence of any cycle, directed or undirected, in a finite arborescence since its underlying graph is a tree. Although we know by Lemma 5.5.1 that the closure of a spanning pseudo-arborescences contains a directed arc from the root to any other vertex (or even end) of the graph, we shall see in the following example that we can neither guarantee the uniqueness of such arcs nor avoid infinite circles (directed or undirected ones).

Example 5.5.2. Consider the graph depicted in Figure 5.5.1. This graph contains spanning r -reachable sets, for example the bold black edges together with the bold grey edges. However, every spanning r -reachable set of this graph must contain all bold black edges because for any head of such an edge there is no other edge of which it is a head. As this graph has only one end, namely ω , we see that there are directed and undirected infinite circles containing only bold black edges. This shows already that, in general, it is not possible to find spanning r -reachable sets that do not contain directed or undirected infinite circles. So there does not exist

a stronger version of Theorem 5.4.3 in the sense that the edges of the underlying multigraph of every spanning pseudo-arborescences form a topological spanning tree in the Freudenthal compactification of the underlying multigraph.




Figure 5.5.1.: An example of a graph with a marked vertex r where the closure of any spanning r -reachable set contains an infinite circle and multiple arcs to the end ω and certain vertices.

The graph in Figure 5.5.1 shows furthermore that, in general, we cannot find spanning r -reachable sets F such that there exists a unique directed arc from r to every vertex and every end of the graph inside \bar{F} . In the example we have two different directed arcs from r to the end ω that contain only bold black edges and are therefore in every spanning r -reachable set of this graph. Hence, we also get two different directed arcs from r to every vertex on the infinite directed circle that consists only of bold black edges.

Although, in general, spanning pseudo-arborescences do not behave like trees in the sense that their underlying graphs correspond to topological spanning trees, they do so in a local sense. We conclude this section with our second main result, Theorem 5.1.6, characterising those spanning r -reachable sets that are inclusion-wise minimal via some local tree-like properties. In particular, we obtain the absence of finite cycles (directed or undirected ones) in any spanning

pseudo-arborescence. As mentioned before, we will prove a slightly stronger version of the theorem.

Theorem 5.5.3. *Let G be a locally finite weakly connected digraph and further let $r \in V(G) \cup \Omega(G)$. Then the following statements are equivalent for a spanning r -reachable set F of G :*

- (i) F is a spanning pseudo-arborescence rooted in r .
- (ii) For every vertex $v \neq r$ of G there is a unique edge in F whose head is v , and no edge in F has r as its head.
- (iii) For every weak component T of $G[F]$ the following holds: If $r \in V(T)$, then T is an arborescence rooted in r . Otherwise, T is an arborescence rooted in some end of T .

Proof. We start by proving the implication from (i) to (ii). Let us first suppose for a contradiction that F contains an edge e whose head is r . Obviously, there is no finite cut $E(X, Y)$ of G such that $r \in \bar{X}$ and $e \in \vec{E}(X, Y)$. Hence, $F \setminus \{e\}$ is a smaller spanning r -reachable set of G contradicting the minimality of F .

Next let us consider an arbitrary vertex $v \neq r$ of G . We know by Remark 5.2.1 that F contains at least one edge of $\vec{E}(V(G) \setminus \{v\}, \{v\})$. So F contains at least one edge whose head is v .

Now suppose for a contradiction that there exists some vertex $v \neq r$ of G which is the head of at least two edges of F , say e and f . We know by Lemma 5.5.1 that \bar{F} contains a directed r - v arc A . Since the cut $E(V(G) \setminus \{v\}, \{v\})$ is finite and A is a directed r - v arc, we get that A must contain precisely one edge of $\vec{E}(V(G) \setminus \{v\}, \{v\})$. Hence, one of the edges e, f is not contained in A , say e . By the minimality of F , we obtain that $F \setminus \{e\}$ cannot be a spanning r -reachable set of G . So there must exist a finite cut $E(X, Y)$ of G with $r \in \bar{X}$ such that e is the only edge in $F \cap \vec{E}(X, Y)$. Now we have a contradiction since the head of e is v and lies in Y , which means that the directed arc A contains at least one edge of $\vec{E}(X, Y)$ by Lemma 5.3.2, but such an edge is different from e . Therefore, e was not the only edge in $F \cap \vec{E}(X, Y)$.

We continue with the proof that statement (ii) implies statement (iii). For this let us fix an arbitrary weak component T of $G[F]$. We now show that T is a tree. Suppose for a contradiction that T contains a directed or undirected cycle C .

If C is a directed cycle, each vertex on C would already be a head of some edge of the cycle. Hence, r cannot be a vertex on C . Applying Remark 5.2.1 with the finite set $V(C)$, we obtain that there needs to be an edge uv of F with $v \in V(C)$ and $u \in V(G) \setminus V(C)$. So v is the head of two edges of F , which contradicts statement (ii).

In the case that C is a cycle, but not a directed one, take a maximal directed path on C . Its endvertex is the head of two edges of C . So we get again a contradiction to statement (ii). We can conclude that T is a tree.

If r is a vertex of T , then it is immediate from statement (ii) that T is an arborescence rooted in r . Otherwise, there needs to be a backwards directed ray R in T as each vertex different from r is the head of a unique edge of F . Let ω be the end of T which contains R . Hence, T is an arborescence rooted in ω , completing the proof of this implication.

It remains to show the implication from (iii) to (i). For this we assume statement (iii) and suppose for a contradiction that F is not minimal with respect to inclusion. Hence, $F' = F \setminus \{e\}$ is a spanning r -reachable set as well for some $e = uv \in F$. Let T be the weak component of $G[F]$ which contains v . As T is an arborescence rooted in r or some end of T , we get that no edge of F' has v as its head. Note that $r \neq v$ because of the edge $uv \in F$. Now we get a contradiction by applying Remark 5.2.1 with F' and the set $\{v\}$, which tells us that F' needs contains an edge whose head is v . \square

The question might arise whether we can be more specific in statement (iii) of Theorem 5.5.3 in the case when r is an end of G . Unfortunately, it is not true that there has to exist a weak component of $G[F]$ whose unique backwards directed ray lies in r . The reason for this is that the end r might be an accumulation point of a sequence of infinitely many different weak components of $G[F]$ in $|G|$ each of which contains a backwards directed ray to a different end of G . It is not difficult to construct an example for this situation and so we omit such a description here. On the other hand if the end $r \in \Omega(G)$ is not an accumulation point of different ends of G , then there exists at least one weak component of $G[F]$ whose backwards directed ray is contained in r . To see this fix an arbitrary directed r - v arc A inside \bar{F} for some vertex v . Since F is a spanning r -reachable set of G , we can find such an arc. If among all of the weak components of $G[F]$ which are met

by A , there is a first one with respect to the linear order of A , then a backwards directed ray of this component is an initial segment of A and, therefore, contained in r . Note for the other case that tails of the backwards directed rays of each component of $G[F]$ that is met by A must be contained in A . Since A is an arc, all these backwards directed rays must be contained in different ends of G . These ends, however, would then have r as an accumulation point in $|G|$ contradicting the assumption on r .

6. On the infinite Lucchesi-Younger conjecture

6.1. Introduction

In finite structural graph theory there are a lot of theorems which illustrate the dual nature of certain objects by relating the maximum number of disjoint objects of a certain type in a graph with the minimal size of an object of another type in that graph. More precisely, the size of the latter object trivially bounds the number of disjoint objects of the first type existing in the graph.

This duality aspect of such packing and covering results is closely related to the duality of linear programs appearing in combinatorial optimisation. However, a purely graph theoretic interpretation requires integral solutions of both linear programs, which are hard to detect, if they even exist.

Probably the most well-known example of such a min-max result is Menger's theorem for finite undirected graphs. It states that for any two vertex sets A, B in a finite graph the maximum number of disjoint paths between A and B equals the minimum size of a vertex set separating A from B . In fact, there is a structural reformulation of this quantitative description of this dual nature of connectivity: for any two vertex sets A, B in a finite graph there exists a set of disjoint paths between A and B together with a vertex set separating A from B that consists of precisely one vertex from each of the paths.

While for finite graphs this is an easy corollary from the quantitative version, in infinite graphs it turns out that such a structural version is much more meaningful. While Erdős observed that a version of Menger's theorem based on the equality of infinite cardinals is quite trivial, he conjectured that the analogue of the structural version is the better way to interpret this dual nature of connectivity. Such a version has been established by Aharoni and Berger [2]. Their theorem restored many of the uses of connectivity duality that the trivial cardinality version could not

provide, and hence it influenced much of the development of infinite connectivity theory and matching theory.

Another such min-max theorem was established by Lucchesi and Younger [35] for directed graphs. To state that theorem we have to give some definitions first.

In a weakly connected directed graph D we call a cut of D *directed*, or a *dicut* of D , if all of its edges have their head in a common side of the cut. We call a set of edges a *dijoin* of D if it meets every non-empty dicut of D . Now we can state the mentioned theorem.

Theorem 6.1.1. [35, Theorem] *In every weakly connected finite digraph, the maximum number of disjoint dicuts equals the minimum size of a dijoin.*

Beside the original proof of Theorem 6.1.1 due to Lucchesi and Younger [35, Theorem], further ones appeared. Among them are an inductive proof by Lovász [34, Theorem 2] and an algorithmic proof of Frank [20, Section 9.7.2]. As for Menger's Theorem, we now state a structural reformulation of Theorem 6.1.1, which for finite digraphs is easily seen to be equivalent.

Theorem 6.1.2. *Let D be a finite weakly connected digraph. Then there exists a tuple (F, \mathcal{B}) such that the following statements hold.*

- (i) \mathcal{B} is a set of disjoint dicuts of D .
- (ii) $F \subseteq E(D)$ is a dijoin of D .
- (iii) $F \subseteq \bigcup \mathcal{B}$.
- (iv) $|F \cap B| = 1$ for every $B \in \mathcal{B}$.

In this chapter we consider the question whether Theorem 6.1.2 extends to infinite digraphs. Let us first show that a direct extension of this formulation to arbitrary infinite digraphs fails. Now consider the digraph depicted in Figure 6.1.1. Its underlying graph is the Cartesian product of a double ray with an edge. Then we consistently orient all edges corresponding to one copy of the double ray in one direction and all edges of the other copy in the different direction. Finally, we direct all remaining edges such that they have their tail in the same copy of the double ray. This digraph contains no finite dicut, but it does contain infinite ones. Note that every dicut of this digraph contains at most one horizontal edge, which

corresponds to an oriented one of some copy of the double ray, and all vertical edges to the left of some vertical edge. Hence, we cannot even find two disjoint dicuts. However, a dijoin of this digraph cannot be finite, as we can easily find a dicut avoiding any finite set of edges by considering a horizontal edge to the left of the finite set. So we obtain that each dijoin hits every dicut infinitely often in this digraph. Therefore, neither the statement of Theorem 6.1.2 nor the statement of Theorem 6.1.1 remain true if we consider arbitrary dicuts in infinite digraphs.




Figure 6.1.1.: A counterexample to an extension of Theorem 6.1.2 to infinite digraphs where infinite dicuts are considered too.

Another counterexample for these naive extensions is the countably infinite transitive tournament without a sink, i.e. an orientation of the countably infinite clique without any directed cycles or sinks. We leave the verification of this fact to the reader.

In order to overcome the problem of this example let us again consider the situation in Menger’s theorem. There, even in the infinite version, we are only considering finite paths for those objects that we want to pack. Together with the example in Figure 6.1.1, this suggests that we might need to restrict our attention to finite dicuts when extending Theorem 6.1.2 to infinite digraphs. Hence, we make the following definitions.

In a weakly connected digraph D we call an edge set $F \subseteq E(D)$ a *finitary dijoin* of D if it intersects every non-empty finite dicut of D . Building up on this definition, we call a tuple (F, \mathcal{B}) as in Theorem 6.1.2 but where F is now a finitary dijoin and \mathcal{B} a set of disjoint finite dicuts of D , an *optimal pair* for D .

Not in contradiction to the example given above, we now state the following conjecture raised by Heuer, which we call the Infinite Lucchesi-Younger Conjecture.

Conjecture 6.1.3. *There exists an optimal pair for every weakly connected digraph.*

Apparently, an extension of Theorem 6.1.1 as in Conjecture 6.1.3 turns out to be very similar to a more general problem about infinite hypergraphs independently

raised by Aharoni [1, Prob. 6.7]. We will discuss this connection further in Section 6.6.

The three mentioned proofs [35, Theorem] [34, Theorem 2] [20, Theorem 9.7.2] of Theorem 6.1.1 even show a slightly stronger result. We call an optimal pair *nested* if the elements of \mathcal{B} are pairwise nested, i.e. any two finite dicuts $E(X_1, X_2), E(Y_1, Y_2) \in \mathcal{B}$ either satisfy one of the following conditions: $X_1 \subseteq Y_1$, $Y_1 \subseteq X_1$, $X_1 \subseteq Y_2$, or $Y_2 \subseteq X_1$.

Theorem 6.1.4. [35, Theorem] *There exists a nested optimal pair for every weakly connected finite digraph.*

Hence, we also make the following conjecture.

Conjecture 6.1.5. *There exists a nested optimal pair for every weakly connected digraph.*

In weakly connected infinite digraphs there are indications that, in contrast to the finite case, Conjecture 6.1.5 may be strictly stronger than Conjecture 6.1.3. In Section 6.3 we will illustrate examples of digraphs with a finitary disjoint which is part of an optimal pair, but not of any nested one.

One of the main results of this chapter is the reduction of Conjectures 6.1.3 and 6.1.5 to countable digraphs with a certain separability property and whose underlying multigraphs are 2-connected. We call a digraph D *finitely diseparable* if for any two vertices $v, w \in V(D)$ there is a finite dicut of D such that v and w lie in different sides of that finite dicut.

Theorem 6.1.6. *If Conjecture 6.1.3 (or Conjecture 6.1.5, respectively) holds for all countable finitely diseparable digraphs whose underlying multigraphs are 2-connected, then Conjecture 6.1.3 (or Conjecture 6.1.5, respectively) holds for all weakly connected digraphs.*

Moreover, we verify Conjecture 6.1.5 for several classes of digraphs. We gather all these results in the following theorem. Before we can state the theorem we have to give some further definitions. We call a minimal non-empty dicut of a digraph a *dibond*. We say a digraph is *rayless* if its underlying multigraph does not contain a ray. We define the *ends of a digraph* as the ends of the underlying multigraph.

Theorem 6.1.7. *Conjecture 6.1.5 holds for a weakly connected digraph D if it has any of the following properties:*

- (i) *There exists a finitary dijoin of D of finite size.*
- (ii) *The maximal number of disjoint finite dicuts of D is finite.*
- (iii) *The maximal number of disjoint and pairwise nested finite dicuts of D is finite.*
- (iv) *Every edge of D lies in only finitely many finite dibonds of D .*
- (v) *D has no infinite dibond.*
- (vi) *D is rayless.*
- (vii) *D is finitely diseparable, contains only finitely many sources, sinks and ends, and contains no backwards directed ray.*

The structure of this chapter is as follows. In Section 6.2 we introduce our needed notation and prove some basic tools that we will need throughout the chapter. In Section 6.3 we will discuss some examples which shall illustrate the difficulties of relating Conjecture 6.1.3 to Conjecture 6.1.5. Section 6.4 is dedicated to the proof of Theorem 6.1.6. In Section 6.5 we shall deduce several items of Theorem 6.1.7 via several lemmas by lifting Theorem 6.1.4 to infinite digraphs via the compactness principle. Section 6.6 is dedicated to a short discussion of the connection between Conjecture 6.1.3 and the more general problem from Aharoni about matchings in infinite hypergraphs. We will extend several parts of the algorithmic proof of Theorem 6.1.1 of Frank [20, Section 9.7.2] in Section 6.7 and introduce sufficient conditions for when this proof yields a positive answer to Conjecture 6.1.5. Finally, in Section 6.8 we will use the results from Section 6.7 to deduce several items of Theorem 6.1.7.

6.2. Basic notions and tools

In general, we allow our digraphs to have parallel edges, but no loops unless we explicitly mention them. Similarly, all undirected multigraphs we consider do not have loops if nothing else is explicitly stated.

Throughout this section let D denote a digraph with vertex set $V(D)$ and edge set $E(D)$.

6.2.1. Cuts and dicuts

Let D be a weakly connected digraph.

Recall that for two vertex set $X, Y \subseteq V(D)$, if $X \cup Y = V(D)$ and $X \cap Y = \emptyset$, we call $E(X, Y)$ a *cut* of D and refer to X and Y as the *sides* of the cut. Moreover, by writing $E(M, N)$ and calling it a cut of D we implicitly assume M and N to be the sides of that cut, and by calling an edge set B a cut we implicitly assume that B is of the form $E(M, N)$ for suitable sets M and N .

We call two cuts $E(X_1, Y_1)$ and $E(X_2, Y_2)$ of D *nested* if one of X_1, Y_1 is \subseteq -comparable with one of X_2, Y_2 . Moreover, we call a set or sequence of cuts of D nested if its elements are pairwise nested. If two cuts of D are not nested, we call them *crossing* (or say that they *cross*).

A cut is said to *separate* two vertices $v, w \in V$ if v and w lie on different sides of that cut.

A minimal non-empty cut is called a *bond*. Note that a cut $E(X, Y)$ is a bond, if and only if the induced subdigraphs $D[X]$ and $D[Y]$ are weakly connected digraphs.

We call a cut $E(X, Y)$ *directed*, or briefly a *dicut*, if all edges of $E(X, Y)$ have their head in one common side of the cut. A bond that is also a dicut is called a *dibond*.

We call D *finitely separable* if for any two different vertices $v, w \in V$ there exists a finite cut of D such that v and w are separated by that cut. Note that if two vertices are separated by some finite cut, then they are separated by some finite bond as well. If furthermore any two different vertices $v, w \in V(D)$ can even be separated by a finite dicut, or equivalently a finite dibond, of D , we call D *finitely diseparable*.

Given a dicut $B = \vec{E}(X, Y)$ we call Y the *in-shore* of B and X the *out-shore* of B . We shall also write $\text{in}(B)$ for the in-shore of the dicut B and $\text{out}(B)$ for the out-shore of B .

For undirected multigraphs *cuts*, *bonds*, *sides*, the notion of being *nested* and the notion of *separating* two vertices are analogously defined. Hence, we call an

undirected multigraph *finitely separable* if any two vertices can be separated by a finite cut of the multigraph.

Given a set $\mathcal{B} = \{B_i \mid i \in I\}$ of dicuts of D , we write

- $\bigwedge \mathcal{B} := \delta^- (\bigcap \{\text{in}(B) \mid B \in \mathcal{B}\})$, or simply $B_1 \wedge B_2$ for $\bigwedge \{B_1, B_2\}$; and
- $\bigvee \mathcal{B} := \delta^- (\bigcup \{\text{in}(B) \mid B \in \mathcal{B}\})$, or simply $B_1 \vee B_2$ for $\bigvee \{B_1, B_2\}$.

Note that since D is weakly-connected, $\bigwedge \mathcal{B}$ is empty if and only if the set $\bigcap \{\text{in}(B) \mid B \in \mathcal{B}\}$ is empty, and $\bigvee \mathcal{B}$ is empty if and only if $\bigcup \{\text{in}(B) \mid B \in \mathcal{B}\}$ equals $V(D)$.

Remark 6.2.1. Let \mathcal{B} be a set of dicuts of D .

1. $\bigwedge \mathcal{B}$ is either empty, or a dicut of D .
2. $\bigvee \mathcal{B}$ is either empty, or dicut of D . □

Note that $\bigwedge \mathcal{B}$ and $\bigvee \mathcal{B}$ might be infinite dicuts of D , even if each $B \in \mathcal{B}$ is finite. Furthermore, note that if B_1 and B_2 are dibonds then $B_1 \wedge B_2$ does not need to be a dibond, even if it is non-empty.

A simple double-counting argument yields the following remark.

Remark 6.2.2. Let B_1 and B_2 be dicuts of D , and let $F \subseteq E(D)$. Then

1. $(B_1 \cap F) \cup (B_2 \cap F) = ((B_1 \wedge B_2) \cap F) \cup ((B_1 \vee B_2) \cap F)$; and
2. $|B_1 \cap F| + |B_2 \cap F| = |(B_1 \wedge B_2) \cap F| + |(B_1 \vee B_2) \cap F|$.

Moreover, if B_1 and B_2 are disjoint, then $B_1 \wedge B_2$ and $B_1 \vee B_2$ are disjoint as well. □

Let B be a dicut. We call a set $\mathcal{B} = \{B_i \mid i \in I\}$ a *decomposition of B* if for each $i, j \in I$

- $B_i \subseteq B$ is a dicut or $B_i = \emptyset$;
- $B_i \cap B_j = \emptyset$ for $i \neq j$;
- $\bigcup_{k \in I} B_k = B$.

We write $B = \bigoplus \mathcal{B}$ if \mathcal{B} is a decomposition of B .

Remark 6.2.3. Let \mathcal{B} be a set of dicuts of D and for each $B \in \mathcal{B}$ let \mathcal{B}_B be a set of dicuts of D such that $B = \bigoplus \mathcal{B}_B$. Let \mathcal{F} be the set of functions from \mathcal{B} to $\bigcup \{\mathcal{B}_B \mid B \in \mathcal{B}\}$ such that each $B \in \mathcal{B}$ is mapped to an element of \mathcal{B}_B . Then

$$\bigwedge \mathcal{B} = \bigoplus \left\{ \bigwedge \{f(B) \mid B \in \mathcal{B}\} \mid f \in \mathcal{F} \right\}$$

and

$$\bigvee \mathcal{B} = \bigoplus \left\{ \bigvee \{f(B) \mid B \in \mathcal{B}\} \mid f \in \mathcal{F} \right\}.$$

6.2.2. Dijoints and optimal pairs for classes of finite dibonds

Let D be a weakly connected digraph.

We call an edge set $F \subseteq E$ a *dijoin* of D if $F \cap B \neq \emptyset$ holds for every dicut B of D . Similarly, we call an edge set $F \subseteq E$ a *finitary dijoin* of D if $F \cap B \neq \emptyset$ holds for every finite dicut B of D . Note that an edge set $F \subseteq E$ is already a (finitary) dijoin if $F \cap B \neq \emptyset$ holds for every (finite) dibond of D since every (finite) dicut is a disjoint union of (finite) dibonds.

Let \mathfrak{B} be a class of finite dibonds of D . Then we call an edge set $F \subseteq E(D)$ a \mathfrak{B} -*dijoin* of D if $F \cap B \neq \emptyset$ holds for every $B \in \mathfrak{B}$. Note that for the class $\mathfrak{B}_{\text{fin}}$ of finite dibonds of D we immediately get that the finitary dijoints of D are precisely the $\mathfrak{B}_{\text{fin}}$ -dijoints of D .

We call a tuple (F, \mathcal{B}) a \mathfrak{B} -*optimal pair* for D if

- (i) $F \subseteq E(D)$ is a \mathfrak{B} -dijoin of D ;
- (ii) $\mathcal{B} \subseteq \mathfrak{B}$ is a set of disjoint dibonds in \mathfrak{B} ;
- (iii) $F \subseteq \bigcup \mathcal{B}$; and
- (iv) $|F \cap B| = 1$ for every $B \in \mathcal{B}$.

We call a \mathfrak{B} -optimal pair (F, \mathcal{B}) for D *nested* if the elements of \mathcal{B} are pairwise nested.

Note that the (nested) optimal pairs as defined in the introduction are precisely the (nested) $\mathfrak{B}_{\text{fin}}$ -optimal pairs.

The main topic of study in this chapter is the following question.

Question 6.2.4. *For which weakly connected digraphs and classes \mathfrak{B} of finite dibonds is there a (nested) \mathfrak{B} -optimal pair?*

Let \mathfrak{B} be a class of dicuts of D . Then let \mathfrak{B}^\oplus denote the class of dicuts B of D which have a partition $B = \bigoplus \mathcal{B}$ for some $\mathcal{B} \subseteq \mathfrak{B}$.

We say \mathfrak{B} is *finite-corner-closed* if

1. If $B_1, B_2 \in \mathfrak{B}$ then either $B_1 \wedge B_2 = \emptyset$ or $B_1 \wedge B_2 \in \mathfrak{B}^\oplus$.
2. If $B_1, B_2 \in \mathfrak{B}$ then either $B_1 \vee B_2 = \emptyset$ or $B_1 \vee B_2 \in \mathfrak{B}^\oplus$.

Note that $(\mathfrak{B}^\oplus)^\oplus = \mathfrak{B}^\oplus$, and that by Remark 6.2.3, if \mathfrak{B} is finite-corner-closed, then so is \mathfrak{B}^\oplus .

Throughout this chapter we will mostly consider classes of finite dibonds of D which are finite-corner-closed, for example $\mathfrak{B}_{\text{fin}}$, the class of finite dibonds of D .

6.2.3. Stars and combs

In this subsection, we first recall the Star-Comb Lemma, which is featured in Chapter 4 of this dissertation and some of the relevant definitions.

Recall that we call an undirected graph a *star* if it is isomorphic to the complete bipartite graph $K_{1,\kappa}$ for some cardinal κ , where the vertices of degree 1 are its *leaves* and the vertex of degree κ is its *centre*.

An undirected multigraph that does not contain a ray is called *rayless*. We define the *ends of a digraph D* as the ends of $\text{Un}(D)$. If ω is an end of an undirected multigraph G (resp. of a digraph D), we call any ray that is contained in ω an *ω -ray*.

Recall that a *comb* C is an undirected graph that is the union of a ray R together with infinitely many disjoint undirected finite paths each of which has precisely one vertex in common with R , which has to be an endvertex of that path. The ray R is called the *spine* of C . The endvertices of the finite paths that are not on R together with the endvertices of the trivial paths are the *teeth* of C .

The following lemma, the Star-Comb Lemma, is a basic tool in infinite graph theory, cf. Lemma 4.2.3. We shall only apply it for vertex sets of cardinality \aleph_0 and \aleph_1 in this chapter.

Lemma 6.2.5. *Let G be an infinite connected undirected multigraph and let $U \subseteq V(G)$ be such that $|U| = \kappa$ for some regular cardinal κ . Then there exists a set $U' \subseteq U$ with $|U'| = |U|$ such that G either contains a comb whose set of teeth is U' or a subdivided star whose set of leaves is U' .*

Later in Section 6.8 we shall need two lemmas about digraphs, which are similar to Lemma 6.2.5 for undirected infinite graphs. In order to state these lemmas we have to give some definitions:

We call a digraph S a *subdivided out-star* (resp. *subdivided in-star*) if $\text{Un}(S)$ is a subdivided star, precisely one vertex c of S has in-degree 0 (resp. out-degree 0) and all other vertices of S have in-degree 1 (resp. out-degree 1). We call the vertex c of S the *centre of S* . Furthermore, we call a vertex v of S a *leaf of S* if v is a leaf of $\text{Un}(S)$. Note that the centre of S coincides with the centre of $\text{Un}(S)$.

Next we call a digraph C a *weak forward* (resp. *weak backward*) *comb* if $\text{Un}(C)$ is a comb and C orients the spine of $\text{Un}(C)$ such that it is a forwards (resp. backwards) directed ray. We call a vertex v of C a *tooth of C* if v is a tooth of $\text{Un}(S)$. A weak forward comb C with set of teeth T and spine S is called a *forward out-comb* (resp. *forward in-comb*) if each S - T path in $\text{Un}(C)$ is a directed one from S to T (resp. from T to S) in C . Analogously, a weak backward comb C with set of teeth T and spine S is called a *backward out-comb* (resp. *backward in-comb*) if each S - T path in $\text{Un}(C)$ is a directed one from S to T (resp. from T to S) in C .

In a digraph D with a vertex $v \in V(D)$ let $N_{\infty}^{+}(v)$ denote the set of all vertices of D that can be reached in D by some directed path starting at v . If Z is a subgraph of D or some subset of the vertices of D , we define $N_{\infty}^{+}(Z) = \bigcup_{z \in V(Z)} N_{\infty}^{+}(z)$. The notation $N_{\infty}^{-}(v)$ and $N_{\infty}^{-}(Z)$ is analogously defined.

Let D be a digraph, $v \in V(D)$ and ω be an end of D . Now we call ω *reachable from v* if there exists an ω -ray R in D that is a forwards directed ray whose start vertex is v . Similarly, we call v *reachable from ω* if there exists an ω -ray R in D that is a backwards directed ray whose start vertex is v . Similarly as for vertices, we let $N_{\infty}^{+}(\omega)$ denote the set of all vertices in D that can be reached from ω and we denote by $N_{\infty}^{-}(\omega)$ the set of all vertices in D that reach ω .

Recall that we call a digraph A an *out-arborescence rooted in $r \in V(A)$* if

- $\text{Un}(A)$ is a tree;
- $d_A^{-}(v) = 1$ for all $v \in V(D) \setminus \{r\}$; and

- $d_A^-(r) = 0$.

A straight-forward transfinite construction yields the following remark:

Remark 6.2.6. Let D be a weakly connected digraph and $U \subseteq V(D)$. Let $U \subseteq N_\infty^+(v)$ for some vertex v of D . Then D contains an out-arborescence rooted in v that contains U .

Now we are able to state and prove two lemmas about digraphs, which have a certain resemblance with Lemma 6.2.5.

Lemma 6.2.7. *Let D be a weakly connected digraph and $U \subseteq V(D)$ an infinite vertex set. If $U \subseteq N_\infty^+(v)$ (resp. $U \subseteq N_\infty^-(v)$) holds for some vertex v of D , then there exists an infinite subset $U' \subseteq U$ such that one of the following assertions is true:*

- there exists a subdivided out-star (resp. subdivided in-star) in D whose set of leaves is U' ; or*
- there exists a forward out-comb (resp. backward in-comb) in D whose set of teeth is U' .*

Proof. We assume that $U \subseteq N_\infty^+(v)$ holds for some vertex v of D . The case that $U \subseteq N_\infty^-(v)$ follows from the first case by reversing the orientation of each edge in D .

Let $A \subseteq D$ be an out-arborescence rooted in v that contains U as in Remark 6.2.6. Applying Lemma 6.2.5 to $\text{Un}(A)$ and U yields either subdivided star S with leaves U' or a comb C with teeth U' for some infinite $U' \subseteq U$.

In the first case, let S' be the subdigraph of D such that $\text{Un}(S') = S$. Without loss of generality we may assume that the centre c of S' has in-degree 0 in S' . So since no vertex has in-degree 2, every path from a c to a leaf u is a directed c - u path. Hence S' is the desired subdivided out star.

In the second case, let C' be the subdigraph of D such that $\text{Un}(C') = C$. Note that the spine R of C' contains at most one vertex w with $d_R^+(w) = 2$, since otherwise it would contain a vertex with in-degree 2 as well. Hence as before, we may assume without loss of generality that the spine R of C' contains no such vertex w with $d_R^+(w) = 2$. And as before, every path from the spine to a tooth u of the comb is a directed R - u path. Hence C' is the desired forward out-comb. \square

In contrast to Lemma 6.2.7 whose statement contains an assumption about the reachability of an infinite vertex set from some vertex, the statement of the next lemma has a similar assumption, but about the reachability from some end.

Lemma 6.2.8. *Let D be a weakly connected digraph and $U \subseteq V(D)$ an infinite vertex set. If $U \subseteq N_\infty^+(\omega)$ (resp. $U \subseteq N_\infty^-(\omega)$) holds for some end ω of D , then there exists an infinite subset $U' \subseteq U$ such that one of the following assertions is true:*

- (a) *there exists a subdivided out-star (resp. subdivided in-star) in D whose set of leaves is U' ; or*
- (b) *there exists a forward out-comb (resp. backward in-comb) in D whose set of teeth is U' .*
- (c) *there exists a weak backward (resp. forward) comb in D whose set of teeth is U' .*

Proof. We only give a proof for the case that $U \subseteq N_\infty^+(\omega)$ holds for some end ω of D since the case that $U \subseteq N_\infty^-(\omega)$ follows from the first statement by reversing all edges in D .

Let ω be an end of D as in the statement. Suppose first that there exists a backwards directed ω -ray R in D such that an infinite set $U_1 \subseteq U$ exists with the property that $U_1 \subseteq N_\infty^+(R)$. If there already exists a finite segment I of R such that $U_2 \subseteq N_\infty^+(I)$ for some infinite set $U_2 \subseteq U_1$, we are done by Lemma 6.2.7.

Hence, we may assume that every finite segment of R reaches only finitely many vertices of U_1 . Using this property it is a standard task to recursively define along R a backward out-comb. We omit this definition here. This completes the proof for the first case.

Now let us consider the remaining case where $U \cap N_\infty^+(R)$ is finite for each backwards directed ω -ray R . For every $u \in U$ let R_u denote some backwards directed ω -ray whose start vertex is u . Note that if $R_u \cap R_v \neq \emptyset$ for $u, v \in U$, then $\{u, v\} \subseteq N_\infty^+(R_u) \cap N_\infty^+(R_v)$. Hence, for each $u \in U$ there are only finitely many $v \in U$ such that $R_u \cap R_v \neq \emptyset$. Using this observation we can recursively define an infinite set $U_2 \subseteq U$ such that $R_u \cap R_v = \emptyset$ holds for all $u, v \in U_2$ with $u \neq v$. Since each ray R_u for $u \in U_2$ is an ω -ray, there are, for any two distinct vertices u, v of U_2 , infinitely many disjoint undirected paths between R_u and R_v

in D . Using this property, it is again a standard task to recursively define a weak backward comb whose spine is any previously chosen R_u for $u \in U_2$. \square

6.2.4. Finitely separable multigraphs

In this section we prove certain size related properties of finitely separable multigraphs using Lemma 6.2.5.

For a multigraph G we call a subgraph $X \subseteq G$ a *2-block* of G if X is a maximal connected subgraph without a cutvertex. Hence a 2-block of a connected multigraph either consists of a set of pairwise parallel edges in G or is a maximal 2-connected subgraph of G . In a digraph D we call a subdigraph X a *2-block* of D if $\text{Un}(X)$ is a 2-block of $\text{Un}(D)$.

One of the tools we will use in this chapter is the so-called *2-block-cutvertex-tree* (cf. [9, Lemma 3.1.4]). Let \mathcal{X} denote the set of all 2-blocks in G , and C the set of all cutvertices in G . Then the bipartite graph with vertex set $\mathcal{X} \cup C$ with edge set $\{cX \mid c \in C, X \in \mathcal{X}, c \in X\}$ is a tree, the *2-block-cutvertex-tree*.

We immediately get the following remark.

Remark 6.2.9. Let G be a multigraph or a digraph.

- (i) Every bond of G is contained in a unique 2-block.
- (ii) Bonds of G that are contained in different 2-blocks are nested.

Lemma 6.2.10. (i) *Every 2-block of a finitely separable multigraph or digraph is countable.*

(ii) *Every 2-block of a finitely separable rayless multigraph or digraph is finite.*

Proof. Let G be a finitely separable multigraph and let X be a 2-block of G . Assume for a contradiction that either X is infinite and rayless, or X is uncountable. Let U be a subset of $V(X)$ with $|U| = \min\{|X|, \aleph_1\}$. Applying Lemma 6.2.5 to U in X , we obtain a subdivided star S_1 in X whose set of leaves L_1 satisfies $|L_1| = |U|$. Let c_1 be the centre of S_1 . Using that X is 2-connected, we now apply Lemma 6.2.5 to L_1 in $G - c_1$, which is still connected. Hence, we obtain a subdivided star S_2 in $G - c_1$ whose set of leaves L_2 satisfies $|L_2| = |L_1|$ and $L_2 \subseteq L_1$. Let c_2 denote the centre of S_2 . Now we get a contradiction to G being finitely separable because S_1 and S_2 have infinitely many common leaves in L_2 . So $G[V(S_1) \cup V(S_2)]$ contains

infinitely many internally disjoint c_1 - c_2 paths, witnessing that c_1 and c_2 cannot be separated by a finite cut of G .

To complete the proof we still need to consider for a contradiction a 2-block X of G whose vertex set is countable (in case (i)) or finite (in case (ii)) but whose edge set is uncountable (in case (i)) or infinite (in case (ii)). A contradiction to the fact that X is finitely separable arises by an easy application of the pigeonhole principle to the two-element subsets of $V(X)$. \square

Together with Remark 6.2.9 we obtain the following immediate corollary.

Corollary 6.2.11. *A finitely separable rayless multigraph has no infinite bond.* \square

6.2.5. Quotients

Let G be a digraph or a multigraph.

For a set $N \subseteq E(G)$ let G/N denote the contraction minor of G which is obtained by contracting inside G all edges of N and deleting all loops that might occur. Similarly, we define $G.N := G/(E(G) \setminus N)$. For a vertex $v \in V(G)$ and any contraction minor $G.N$ with $N \subseteq E(G)$ let \dot{v} denote the vertex in $G.N$ which corresponds to the contracted, possibly trivial, (weak) component of $G - N$ containing v .

We state the following basic lemma without proof.

Lemma 6.2.12. *Let $B, N \subseteq E(G)$ with $B \subseteq N$ and let $v, w \in V(G)$. Then B is a cut (or dicut/bond/dibond, respectively) of G that separates v and w if and only if B is a cut (or dicut/bond/dibond, respectively) of $G.N$ that separates \dot{v} and \dot{w} .*

Moreover, two cuts $B_1, B_2 \subseteq N$ are nested as cuts of G if and only if they are nested as cuts of $G.N$. \square

Given a set \mathcal{B} of cuts of G , we define an equivalence relation on $V(G)$ by setting $v \equiv_{\mathcal{B}} w$ if and only if we cannot separate v from w by a cut in \mathcal{B} . It is easy to check that $\equiv_{\mathcal{B}}$ indeed defines an equivalence relation. For $v \in V(G)$ we shall write $[v]_{\equiv_{\mathcal{B}}}$ for the equivalence class with respect to $\equiv_{\mathcal{B}}$ containing v .

Let $G/\equiv_{\mathcal{B}}$ denote the digraph, or multigraph respectively, which is obtained from G by identifying the vertices in the same equivalence class of $\equiv_{\mathcal{B}}$ and deleting loops. Furthermore, let $\hat{X} := \{[x]_{\equiv_{\mathcal{B}}} \mid x \in X\}$ for every set $X \subseteq V(G)$, as well as $\tilde{X} := \{y \in V(G) \mid [y]_{\equiv_{\mathcal{B}}} \in \hat{X}\}$ for every set $X \subseteq V(G)/\equiv_{\mathcal{B}}$.

Proposition 6.2.13. *Let G be a digraph or a multigraph and let \mathcal{B} be a set of cuts of G . Then the following statements hold.*

- (i) $G/\equiv_{\mathcal{B}}$ is (weakly) connected if G is (weakly) connected.
- (ii) Every cut (or dicut/bond/dibond, respectively) $E(X, Y) \in \mathcal{B}$ of G is also a cut (or dicut/bond/dibond, respectively) of $G/\equiv_{\mathcal{B}}$, and $E(X, Y) = E(\hat{X}, \hat{Y})$.
- (iii) Every cut (or dicut, respectively) $E(X, Y)$ of $G/\equiv_{\mathcal{B}}$ is also a cut (or dicut, respectively) of G , and $E(X, Y) = E(\tilde{X}, \tilde{Y})$.
- (iv) Two cuts in \mathcal{B} are nested as cuts of G if and only if they are nested as cuts of $G/\equiv_{\mathcal{B}}$.
- (v) $G/\equiv_{\mathcal{B}}$ is \mathcal{B} -separable.

Proof. For the sake of readability we will phrase the proof just for cuts and bonds. The arguments for dicuts and dibonds are analogous.

Note that if $G[X]$ is (weakly) connected for some $X \subseteq V(G)$, then $G/\equiv_{\mathcal{B}}[\hat{X}]$ is (weakly) connected as well. Hence statement (i) is immediate.

If $E(X, Y) \in \mathcal{B}$, then for every $x \in X$ all vertices in $[x]_{\equiv_{\mathcal{B}}}$ are contained in X by definition of $\equiv_{\mathcal{B}}$. Analogously, all vertices in $[y]_{\equiv_{\mathcal{B}}}$ lie in Y for each $y \in Y$. Hence, $E(\hat{X}, \hat{Y}) = E(X, Y)$ and is a cut of $D/\equiv_{\mathcal{B}}$. If $E(X, Y)$ is a bond of G , then so it is as a bond of $G/\equiv_{\mathcal{B}}$ by the observation on connectivity of the sides from above. This proves statement (ii).

For statement (iii) let $E(X, Y)$ be a cut of $G/\equiv_{\mathcal{B}}$. By definition of $\equiv_{\mathcal{B}}$ we obtain that $E(X, Y)$ is a cut of D as well as $M = \tilde{X}$ and $N = \tilde{Y}$ yielding $E(X, Y) = E(\tilde{X}, \tilde{Y})$.

For any subsets $X, Y \subseteq V(G)$ if $X \subseteq Y$, then $\hat{X} \subseteq \hat{Y}$. Moreover, for any subsets $X, Y \subseteq V(G)/\equiv_{\mathcal{B}}$ if $X \subseteq Y$, then $\tilde{X} \subseteq \tilde{Y}$. With these observations, statement (iv) is immediate.

In order to show statement (v), let $[v]_{\equiv_{\mathcal{B}}}$ and $[w]_{\equiv_{\mathcal{B}}}$ be two different vertices of $V(G/\equiv_{\mathcal{B}})$. Since v and w are not contained in the same equivalence class, there must exist a cut $E(X, Y) \in \mathcal{B}$ separating them. By statement (ii) we get that $E(\tilde{X}, \tilde{Y})$ is a cut of $G/\equiv_{\mathcal{B}}$ and it separates $[v]_{\equiv_{\mathcal{B}}}$ from $[w]_{\equiv_{\mathcal{B}}}$ by definition of $\equiv_{\mathcal{B}}$. □

We will apply this proposition mostly with the set of all finite bonds of a multigraph G , or the set $\mathfrak{B}_{\text{fin}}$ of all finite dibonds of a digraph D , yielding a multigraph which is finitely separable or a digraph which is finitely diseparable.

Let D be any digraph and let $\mathfrak{B}_{\text{fin}}$ be the set of finite dibonds of D . For ease of notation let \sim denote the relation $\equiv_{\mathfrak{B}_{\text{fin}}}$.

Next we characterise the relation $v \sim w$ for any two vertices v, w . An edge set W is a *witness* for $v \sim w$, if it meets every finite cut that separates v and w in both directions, i.e. $W \cap \vec{E}(X, Y) \neq \emptyset \neq W \cap \vec{E}(Y, X)$. Hence the existence of a witness for $v \sim w$ is an obvious obstruction. The whole edge set is similar trivially a witness for $v \sim w$. Note that there exists always an inclusion-minimal witness for $v \sim w$ by Zorn's Lemma.

The following lemmas tell us that given a minimal witness W for $v \sim w$, all vertices incident with an edge of W are also equivalent to v with respect to \sim .

Lemma 6.2.14. *Let $v \sim w$ for two vertices $v, w \in V(D)$. Then a minimal witness W for $v \sim w$ also witnesses $v \sim y$ for any $y \in V(D[W])$.*

Proof. Let W be a minimal witness for $v \sim w$. Now suppose for a contradiction that there is a $y \in V(D[W])$ which is separated from v by a finite dibond $B = \vec{E}(X, Y)$ of D and $W \cap B = \emptyset$. Without loss of generality let $y \in Y$. Since W witnesses $v \sim w$, both vertices v and w have to lie on the same side of B , namely X . We claim that $W' := W \cap E(D[X])$ also witnesses $v \sim w$. This would be a contradiction to the minimality of W as y is incident with an edge of W both of whose endvertices lie in Y since $W \cap B = \emptyset$.

Let $E(M, N)$ be a finite cut of D separating v and w , say with $v \in M$ and $w \in N$. Since $E(X \cap M, Y \cup N)$ is also a finite cut, but $W \cap E(X \cap M, Y) = \emptyset$, we obtain $W' \cap \vec{E}(M, N) \neq \emptyset \neq W' \cap \vec{E}(N, M)$ as desired. \square

Corollary 6.2.15. *Let $v \sim w$ for two vertices $v, w \in V(D)$. Then a minimal witness W for $v \sim w$ induces a strongly connected digraph $D[W]$.*

Proof. Assume for a contradiction that there is a dicut $\vec{E}(X, Y)$ separating some vertices $w_1, w_2 \in W$. By Lemma 6.2.14, W is also a witness for $w_1 \sim w_2$, contradicting that $\vec{E}(Y, X) = \emptyset$. \square

Given a nested set \mathcal{B} of disjoint dicuts of D , there is the notion of a *structure tree* defined by \mathcal{B} . To define this tree, consider the digraph $T(\mathcal{B})$ obtained from

deleting parallel edges from $D/\equiv_{\mathcal{B}}$. It is well known that if \mathcal{B} is a nested set of disjoint dicuts, then this graph is a directed tree. See for example Dicks and Dunwoody [8] for more details.

Let us close this subsection with the following corollary of Proposition 6.2.13 and Lemma 6.2.10(i).

Corollary 6.2.16. *Let \mathcal{B} be a set of finite cuts of G . Each 2-block of $G/\equiv_{\mathcal{B}}$ is countable.* \square

6.2.6. Quotients of rayless digraphs

Let D be any digraph, let $\mathfrak{B}_{\text{fin}}$ be the set of finite dibonds of D , and let $\mathfrak{B}_{\text{fin}}^*$ be the set of finite bonds of D . As in the previous subsection, we denote for the sake of readability the relation $\equiv_{\mathfrak{B}_{\text{fin}}}$ by \sim . Moreover, we denote the relation $\equiv_{\mathfrak{B}_{\text{fin}}^*}$ by \approx .

Note that since $v \approx w$ implies that $v \sim w$ for all $v, w \in V(F)$, we obtain that \sim induces an equivalence relation on $V(D/\approx)$. Since moreover the set of finite dibonds of D/\approx equals the set of finite dibonds of D by Proposition 6.2.13, we obtain the following remark.

Remark 6.2.17. $(D/\approx)/\sim = D/\sim$ \square

The aim of this subsection is to show that if D is rayless, then so is D/\sim . The analogous statement for the relation \approx is proven by an easy construction.

Remark 6.2.18. If D is rayless, then D/\approx is rayless as well.

Proof. Suppose for a contradiction that D is rayless but $R = [v_0]_{\approx}[v_1]_{\approx} \dots$ is a ray in D/\approx . For each $i \in \mathbb{N}$ let $v'_i \in [v_i]_{\approx}$ and $v''_{i+1} \in [v_{i+1}]_{\approx}$ be the endvertices of the edge $[v_i][v_{i+1}] \in E(R)$ seen in D . To arrive at a contradiction, we will construct a ray in D inductively. Let P_0 be the trivial path containing just v'_0 . Assume for $i > 0$ that there is a $j \geq i$ such that that P_i is a $[v_0]_{\approx} - [v_j]_{\approx}$ -path which contains P_{i-1} and is internally disjoint to $[v_k]_{\approx}$ for all $k \geq j$. Let v'''_j be the endvertex of P_i in $[v_j]_{\approx}$. By definition of \approx there is a $v'''_j - v''_j$ path P disjoint from P_i . If P is disjoint from $[v_k]_{\approx}$ for all $k > j$, then let P_{i+1} be concatenation of the paths P_i , P and the edge $v''_j v'_{j+1}$. Otherwise let w be the first vertex of P in $[v_k]_{\approx}$ for some $k > j$ and let P_{i+1} be the concatenation of P_i with $v'''_j P w$. In both cases P_{i+1} satisfies the desired properties and $\bigcup_{i \in \mathbb{N}} P_i$ is the desired ray in D . \square

Before we can prove the analogue for digraphs, we have to prepare some lemmas. The first is about inclusion-minimal edge sets witnessing the equivalence of two vertices with respect to \sim in digraphs whose underlying multigraph is rayless.

Lemma 6.2.19. *Let D be a rayless and finitely separable digraph. Let $v, w \in V(D)$ with $v \sim w$. Then any minimal edge set of D witnessing $v \sim w$ is finite.*

Proof. Let $W \subseteq E(D)$ be an inclusion-minimal witness for $v \sim w$. Let us consider the 2-block-cutvertex tree T of D . Let P denote the finite path in T whose endvertices are the 2-blocks of D containing v and w , respectively. By Remark 6.2.9, each bond of D separating v and w is a bond of the finitely many 2-blocks corresponding to the vertices of P . This implies that all edges in W are contained in the finitely many 2-blocks which correspond to vertices of P . However, each 2-block of D is finite since D is finitely separable and rayless and such multigraphs do not have infinite 2-blocks by Lemma 6.2.10(ii). So W is contained in a finite set and thus finite itself. \square

Proposition 6.2.20. *If D is rayless, then so is D/\sim .*

Proof. By Remarks 6.2.17 and 6.2.18 we may assume without loss of generality that D is finitely separable. Suppose for a contradiction that D is rayless but $R = [v_0]_{\sim}[v_1]_{\sim} \dots$ is a ray in D/\sim . For each $i \in \mathbb{N}$ let $v'_i \in [v_i]_{\sim}$ and $v''_{i+1} \in [v_{i+1}]_{\sim}$ be the endvertices of the edge $[v_i]_{\sim}[v_{i+1}]_{\sim} \in E(R)$ seen in D . Furthermore, let W_i be an inclusion-minimal witness for $v''_i \sim v'_{i+1}$ for every $i \in \mathbb{N}$ with $i \geq 1$. We know by Lemma 6.2.14 that each W_i is completely contained in $[v_i]_{\sim}$. By Corollary 6.2.15 and Lemma 6.2.19 each W_i is strongly connected and finite. Since each W_i is completely contained in $[v_i]_{\sim}$, we get that $W_i \cap W_j = \emptyset$ holds for all $i, j \in \mathbb{N}$ with $i \neq j$. Let P_i be a directed $v''_i - v'_{i+1}$ path that is contained in W_i for every $i \in \mathbb{N}$ with $i \geq 1$. Now the union of these paths together with the edges between v'_i and v''_{i+1} is a ray in D , a contradiction. \square

6.2.7. Cost functions and feasible potentials

Let D be a weakly connected digraph.

We call a function $c : E(D) \rightarrow \mathbb{R}$ a *cost function* on D . Furthermore, we call a cost function c on D *integer-valued* if c maps into the integers, so $\text{Im}(c) \subseteq \mathbb{Z}$. Given a finite edge set $Z \subseteq E(D)$ we define for convenience $c(Z) := \sum_{z \in Z} c(z)$. Similarly,

we define $c(D') := c(E(D'))$ for a finite digraph $D' \subseteq D$. We shall call an edge set $Z \subseteq E(D)$ *negative with respect to c* if $c(Z) < 0$ holds. Since in our arguments we shall often have a fixed cost function c , we will call an edge set $Z \subseteq E(D)$ just *negative*, but mean negative with respect to c . Similarly, we call a finite digraph $D' \subseteq D$ *negative* if $E(D')$ is negative.

A function $\pi : V(D) \rightarrow \mathbb{R}$ is called a *potential* on D . As for cost functions on D we say that π is *integer-valued* if $\text{Im}(\pi) \subseteq \mathbb{Z}$ holds. Given a cost function c on D we call a potential π on D *feasible with respect to c* if the following inequality is satisfied for every edge $uv \in E(D)$:

$$\pi(v) - \pi(u) \leq c(uv). \quad (*)$$

As before, since our cost function c will be fixed during our argumentation, we shall call a potential π on D just *feasible*, but mean feasible with respect to c .

A well-known theorem due to Gallai about finite digraphs with a cost function establishes a dichotomy between the existence of a negative directed cycle and a feasible potential, which if the cost-function is integer-valued can be chosen to be integer-valued as well.

We extend this theorem to the class of infinite digraphs which are strongly connected. The proof of our theorem uses very similar ideas as a proof for the finite theorem.

Theorem 6.2.21. *Let D be a strongly connected digraph with a cost function c on D . Then either D contains a negative directed cycle or there exists a feasible potential on D , but not both together.*

Moreover, if c is integer-valued and D does not contain a negative directed cycle, then there exists a feasible integer-valued potential as well.

Proof. Let us suppose for a contradiction that both, a negative directed cycle $C \subseteq D$ and a feasible potential π on D exist together. Let $V(C) = \{v_0, v_1, \dots, v_n\}$ for some $n \in \mathbb{N}$ with $n \geq 1$ such that $E(C) = \{v_i v_{i+1} \mid 0 \leq i \leq n-1\} \cup \{v_n v_0\}$. Now we obtain the following contradicting chain of inequalities:

$$0 = \sum_{v \in V(C)} \pi(v) - \sum_{v \in V(C)} \pi(v) \leq c(C) < 0.$$

Note that the first inequality holds since π is feasible and so $(*)$ holds for each edge of C . The latter proper inequality holds by assuming that C is negative.

Now let us assume that D does not contain a negative directed cycle. We shall define a certain feasible potential π on D . For this we fix an arbitrary vertex r of D , on which the definition of π will depend. For every $v \in V(D)$ we now define π as follows:

$$\pi(v) := \inf\{c(P) \mid P \text{ is a directed } r\text{-}v \text{ path}\}.$$

Let us first prove that π is well-defined by showing $\text{Im}(\pi) \subseteq \mathbb{R}$. Suppose for a contradiction there is a vertex $v \in V(D)$ such that $\pi(v) = -\infty$. Now choose a directed v - r path P_{vr} in D . Such a path exists because D is strongly connected. Let P_{vr} have $k \in \mathbb{N}$ many edges and let $c_{\max} := \max\{c(N) \mid N \subseteq E(P_{vr})\}$. Now choose an r - v path P_{rv} such that $c(P_{rv}) < \min\{0, -kc_{\max}\}$. This is possible because $\pi(v) = -\infty$ holds. Since P_{vr} and P_{rv} have the same endvertices, we know that P_{rv} contains at most k many directed subpaths which are disjoint from P_{vr} except from their endvertices. Hence, there is one such subpath P^* of P_{rv} such that $c(P^*) < \min\{0, -c_{\max}\}$. Let p^* and q^* be the endvertices of P^* , which lie on P_{rv} , and let P_{rv}^* denote the directed p^* - q^* subpath of P_{rv} . Then, by definition of c_{\max} , we know that $c(P_{rv}^*) \leq c_{\max}$. So by our choice of P^* we obtain that $P^* \cup P_{rv}^*$ is a negative directed cycle, which contradicts our assumption. Hence π is a well-defined potential on D .

It remains to prove that π is feasible. Suppose for a contradiction that property $(*)$ fails for some edge $uv \in E(D)$. Then we obtain that the inequality

$$\pi(v) - \pi(u) > c(uv) + \varepsilon$$

must hold as well for some sufficiently small $\varepsilon \in \mathbb{R}$ with $\varepsilon > 0$. Let P_u be a directed r - u path, such that the following holds:

$$c(P_u) \leq \pi(u) + \frac{\varepsilon}{2}.$$

If P_u together with the edge uv forms a directed r - v path, call it P'_v , then we would obtain the following contradiction to the definition of π :

$$c(P'_v) = c(P_u) + c(uv) \leq \pi(u) + \frac{\varepsilon}{2} + c(uv) < \pi(v) - \frac{\varepsilon}{2}.$$

Hence, v is contained in the directed r - u path P_u . Note that

$$c(rP_u v) \geq \pi(v)$$

holds as $rP_u v$ is also a directed $r-v$ path. Let C denote the directed cycle consisting of the directed $v-u$ path $vP_u u$ together with the edge uv . Now we are able to point out a contradiction by showing that C is a negative directed cycle:

$$c(C) = c(P_u) - c(rP_u v) + c(uv) < \pi(v) - \frac{\varepsilon}{2} - c(rP_u v) \leq -\frac{\varepsilon}{2} < 0.$$

So property $(*)$ holds for every edge of D , which shows that π is a feasible potential on D .

Note for the remaining statement of the theorem that the feasible potential π defined as above will automatically be integer-valued if c is integer-valued. \square

If a strongly connected digraph D with a cost function c on it has no negative directed cycle, we shall call a potential π on D defined as in the proof of Theorem 6.2.21 a *rooted potential* on D . Furthermore, we call the vertex r on which the definition of π depends the *root* of π .

6.2.8. More on dicuts

Lemma 6.2.22. *Let B be a dicut and $\mathcal{B} = \{B_i \mid i \in I\}$ be a decomposition of B into dibonds. Let C, C' be two weak components of $\text{in}(B)$ such that there is an index $i \in I$ with $C \subseteq \text{in}(B_i)$ and $C' \subseteq \text{out}(B_i)$. Then there is a $j \in I$ with $C' \subseteq \text{in}(B_j)$ and $C \subseteq \text{out}(B_j)$.*

Proof. By Lemma 6.2.12 it is sufficient to consider this problem in the digraph $D.B$. Note that in this digraph every vertex is either a source or a sink, and there are sinks c, c' corresponding to the components C and C' .

Let K be an undirected cycle in $D.B$ and let $v, w \in V(K)$ be distinct. We colour the edges of K red if they point towards v and green if they point towards w . It is easy to check that each dibond A of $D.B$ has the following properties:

- A meets K evenly;
- A meets each of the $v-w$ paths on K in alternating colours with respect to the linear order given by the paths;
- if A meets both $v-w$ paths on K , then the first edges (and last edges, respectively) on these paths have the same colour;
- A has v in its in-shore if and only if the first edges of the paths are red;

- A has w in its in-shore if and only if the last edges of the paths are green;
- A separates v and w if and only if A meets each of the v - w paths on K oddly;

If v is a source and w is a sink, then each of the v - w paths on K have an odd number of edges. Since \mathcal{B} is a decomposition of B and thus covers all edges of K we obtain from the properties above that there is a dibond $A \in \mathcal{B}$ with $v \in \text{out}(A)$ and $w \in \text{in}(A)$. A similar argument yields that if v and w are both sinks (or sources, respectively) and there is a dibond $A \in \mathcal{B}$ with $v \in \text{out}(A)$ and $w \in \text{in}(A)$, then there is also a dicut $A' \in \mathcal{B}$ with $v \in \text{in}(A')$ and $w \in \text{out}(A')$.

Consider in the 2-block-cut-vertex tree of $D.B$ the shortest path $P = b_1 c_1 \cdots c_n b_n$ between a block b_1 containing c and a block b_n containing c' . Since both $c =: c_0$ and $c' =: c_{n+1}$ are both sinks, we consider two cases.

If not all c_i are sinks, then there are $0 \leq j < k \leq n$ such that c_j is a sink, c_{j+1} is a source, c_k is a source and c_{k+1} is a sink. Then b_j contains a bond in \mathcal{B} with c_j (and hence c) in its in-shore and c_{j+1} (and hence c') in its out-shore since either b_j is such a bond or we obtain such a bond by fixing a cycle through c_j and c_{j+1} and obtain the bond as above. Similarly b_k contains a bond in \mathcal{B} with c_{k+1} (and hence c') in its in-shore and c_k (and hence c) in its out-shore, as desired.

If all c_i are sinks, then since by assumption there is a dibond in \mathcal{B} with c in its out-shore and c' in its in-shore, this dibond is contained in a block b_j for some $1 \leq j \leq n$, and hence has c_j in its out-shore and c_{j+1} in its in-shore. This block cannot be a set of parallel edges since c_j is a sink. By again fixing a cycle through c_j and c_{j+1} we obtain as above a bond in B with c_j (and hence c) in its in-shore and c_{j+1} (and hence c'), as desired. \square

6.3. Comparing Conjecture 6.1.3 with Conjecture 6.1.5

In this section we shall compare Conjecture 6.1.3 with Conjecture 6.1.5 more closely by looking at two examples. In both examples we will see an indication why Conjecture 6.1.5 might be properly stronger than Conjecture 6.1.3. To put it straight, both examples show the following:

There exist finitary dijoins that are part of an optimal pair, but of no nested optimal pair.

This is severely different from finite digraphs. There, we could always keep the dijoin F of any optimal pair (F, \mathcal{B}) and just iteratively ‘uncross’ all dicuts of \mathcal{B} , yielding a set \mathcal{B}' of nested disjoint dicuts such that (F, \mathcal{B}') is a nested optimal pair. We illustrate this uncrossing process in the proof of Lemma 6.5.1.

Let us now describe the first example.

Example 6.3.1. Consider the infinite weakly connected digraph D_1 depicted twice in Figure 6.3.1. Before we analyse D in detail, let us define D properly.

Let $A = \{a_i \mid i \in \mathbb{N}\}$ and $B = \{b_i \mid i \in \mathbb{N}\}$ be two disjoint countably infinite sets. Additionally, let r be some set which is neither contained in A nor in B . Now we set

$$V(D) := A \cup B \cup \{r\}.$$

We define $E_1 := \{a_i b_i \mid i \in \mathbb{N}\}$, $E_2 := \{a_i b_{i+1} \mid i \in \mathbb{N}\}$ and $E_3 := \{b_i r \mid i \in \mathbb{N}\}$. We complete the definition of D by setting

$$E(D) := E_1 \cup E_2 \cup E_3.$$

Next consider the set E_2 of grey edges in the left instance of D_1 depicted in Figure 6.3.1, call it F_L . It is easy to check that F_L forms a finitary dijoin of D_1 . Furthermore, we can easily find a nested optimal pair for D_1 in which F_L features. Hence, D_1 is not a counterexample to Conjecture 6.1.5.

In the right instance of D_1 depicted in Figure 6.3.1, the set of grey edges $E_1 \cup \{b_0 r\}$, call it F_R , also forms a finitary dijoin. And again we can easily find an optimal pair for D_1 in which F_R features. However, no matter which finite dijoin we choose which contains the grey edge adjacent to r , it cannot be nested with all the finite dicuts we choose for all the other edges of F_R . Therefore, F_R does not feature in any nested optimal pair for D_1 .

Let us now consider another example, witnessing the same behaviour of finitary dijoins as Example 6.3.1 does. However, the structure of the digraph D_2 in the following example is rather different from D_1 . In particular, D_2 is a locally finite digraph, i.e. every vertex is incident with only finitely many edges.




Figure 6.3.1.: Two instances of the digraph D_1 . All edges are meant to be directed from left to right. The grey edges in the left instance of D_1 form a finitary dijoin featuring in a nested optimal pair for D_1 . The grey edges in the right instance form a finitary dijoin featuring in an optimal pair for D_1 , but not in any nested optimal pair for D_1 .

Example 6.3.2. Consider the infinite weakly connected digraph D_2 depicted in Figure 6.3.2. We first define vertex set of D_2 as

$$V(D_2) := \left\{ (x, y) \in \mathbb{Z} \times \mathbb{Z} \mid \frac{x}{2} - y \leq 1 \right\}.$$

Note that for each $(x, y) \in V(D_2)$ both $(x, y + 1)$ and $(x - 1, y)$ are in $V(D_2)$ as well. Next, we define the sets $E_1 := \{(x, y + 1)(x, y) \mid (x, y) \in V(D_2)\}$ and $E_2 := \{(x - 1, y)(x, y) \mid (x, y) \in V(D_2)\}$. Finally, we define the edge set of D_2 by

$$E(D_2) := E_1 \cup E_2.$$

Now consider the set of dashed grey edges in Figure 6.3.2,

$$F_d := \left\{ (x, y + 1)(x, y) \mid \frac{x}{2} - y = 1 \right\}.$$

It is an easy exercise to check that F_d forms a finitary dijoin of D_2 which also features in a nested optimal pair for D_2 . Therefore, the digraph D_2 is also no counterexample to Conjecture 6.1.5.

In contrast to this, let us now consider the set of sustainedly grey edges in Figure 6.3.2,

$$F_s := \left\{ (x - 1, y)(x, y) \mid \frac{x}{2} - y = \frac{1}{2} \text{ or } \frac{x}{2} - y = 1 \right\}.$$




Figure 6.3.2.: The digraph D_2 . The edges are meant to be directed from left to right and from top to bottom. The dashed grey edges form a finitary disjoint featuring in a nested optimal pair for D_2 . The sustainedly grey edges form a finitary disjoint featuring in an optimal pair for D_2 , but not in any nested optimal pair for D_2 .

Again it is easy to check that F_s forms a finitary disjoint of D_2 . However, F_s is not part of any nested optimal pair for D_2 . This is not difficult to prove using the fact that F_d is a finitary disjoint of D_2 as well. We leave this proof to the reader.

6.4. Reductions for the Infinite

Lucchesi-Younger Conjecture

In this section we prove some reductions for Conjecture 6.1.3 and Conjecture 6.1.5 in the sense that it suffices to solve these conjectures on a smaller class of digraphs. We begin by reducing these conjectures to finitely separable digraphs via the following lemma.

Lemma 6.4.1. *Let D be a weakly connected digraph and \mathfrak{B} be a class of dibonds of D . Then (F, \mathfrak{B}) is a (nested) \mathfrak{B} -optimal pair for D if and only if it is a (nested) \mathfrak{B} -optimal pair for $D / \equiv_{\mathfrak{B}}$.*

Proof. Note first that by Proposition 6.2.13 $D / \equiv_{\mathfrak{B}}$ is weakly connected and that

\mathfrak{B} is also a set of dibonds of $D/\equiv_{\mathfrak{B}}$.

Suppose (F, \mathcal{B}) is a (nested) \mathfrak{B} -optimal pair for D . Then F is still a subset of $E(D/\equiv_{\mathfrak{B}})$ since each edge of F lies on some dibond $B \in \mathcal{B} \subseteq \mathfrak{B}$. Hence, F is still a \mathfrak{B} -dijoin of $D/\equiv_{\mathfrak{B}}$, and (F, \mathcal{B}) is indeed a (nested) \mathfrak{B} -optimal pair for $D/\equiv_{\mathfrak{B}}$, again by Proposition 6.2.13.

Similarly, if (F, \mathcal{B}) is a (nested) \mathfrak{B} -optimal pair for $D/\equiv_{\mathfrak{B}}$, then so it is for D , again by Proposition 6.2.13. \square

The next reduction of Conjecture 6.1.3 and Conjecture 6.1.5 tells us that we can restrict our attention also to digraphs whose underlying multigraph is 2-connected.

Lemma 6.4.2. *Let D be a weakly connected digraph and \mathfrak{B} be a class of dibonds of D . Let \mathcal{X} denote the set of all 2-blocks of D . Then the following statements are true.*

(i) *For each $X \in \mathcal{X}$ the set $\mathfrak{B}_X := \{B \in \mathfrak{B} \mid B \subseteq E(X)\}$ is a class of dibonds of X and $\mathfrak{B} = \dot{\bigcup}_{X \in \mathcal{X}} \mathfrak{B}_X$. Moreover, if \mathfrak{B} is finite-corner-closed, then so is \mathfrak{B}_X .*

(ii) *If (F, \mathcal{B}) is a (nested) \mathfrak{B} -optimal pair for D , then (F_X, \mathcal{B}_X) is a (nested) \mathfrak{B}_X -optimal pair for every $X \in \mathcal{X}$, where $F_X := F \cap E(X)$ and $\mathcal{B}_X := \{B \in \mathcal{B} \mid B \subseteq E(X)\}$.*

(iii) *If (F_X, \mathcal{B}_X) is a (nested) \mathfrak{B}_X -optimal pair for every $X \in \mathcal{X}$, then (F, \mathcal{B}) is a (nested) \mathfrak{B} -optimal pair for D , where $F := \bigcup_{X \in \mathcal{X}} F_X$ and $\mathcal{B} := \bigcup_{X \in \mathcal{X}} \mathcal{B}_X$.*

Proof. Let X be a 2-block of D . By Remark 6.2.9 every dibond $B \in \mathfrak{B}$ is either contained in $E(X)$ and hence a dibond of X , or disjoint to $E(X)$. Vice versa, every dibond of X is a dibond of D as well. Statement (i) is now easy to check.

For statement (ii), let $X \in \mathcal{X}$ and let (F, \mathcal{B}) be a (nested) \mathfrak{B} -optimal pair for D . Then by just translating the definitions we obtain that $(F \cap E(X), \mathcal{B}_X)$ is a (nested) \mathfrak{B}_X -optimal pair for D , as well as for X .

Now we show that statement (iii) is true. So let us assume that (F_X, \mathcal{B}_X) is a (nested) \mathfrak{B}_X -optimal pair for every $X \in \mathcal{X}$. With statement (i) (and Remark 6.2.9(ii)) we immediately get that with (F, \mathcal{B}) is a (nested) \mathfrak{B} -optimal pair for D . \square

We can now close this section by proving Theorem 6.1.6. In order to do this we basically only need to combine Lemma 6.4.1 and Lemma 6.4.2. Let us restate the theorem.

Theorem 6.1.6. *If Conjecture 6.1.3 (or Conjecture 6.1.5, respectively) holds for all countable finitely diseparable digraphs whose underlying multigraph is 2-connected, then Conjecture 6.1.3 (or Conjecture 6.1.5, respectively) holds for all weakly connected digraphs.*

Proof. Let D be any weakly connected digraph and let $\mathfrak{B}_{\text{fin}}$ the set of finite dibonds of D . We know by Proposition 6.2.13 that $D/\equiv_{\mathfrak{B}_{\text{fin}}}$ is a weakly connected and finitely diseparable digraph, and so is every 2-block of it. Furthermore, Corollary 6.2.16 yields that each 2-block of $D/\equiv_{\mathfrak{B}_{\text{fin}}}$ is countable. By our assumption we know that Conjecture 6.1.3 (or Conjecture 6.1.5, respectively) holds for every countable 2-block of $D/\equiv_{\mathfrak{B}_{\text{fin}}}$. So using Lemma 6.4.2 we obtain a (nested) optimal pair for $D/\equiv_{\mathfrak{B}_{\text{fin}}}$. Then we also obtain an optimal pair for D by Lemma 6.4.1. \square

6.5. Special cases

In this section we prove some special cases of Conjecture 6.1.5, or more precisely cases of Question 6.2.4.

6.5.1. Finite parameters

Let D be a weakly connected digraph. Let \mathfrak{B} be a class of finite dibonds of D . Before we come to the first special case, we state a basic observation.

Lemma 6.5.1. *The following statements are equivalent:*

- (i) *There is \mathfrak{B} -dijoin of D of finite size.*
- (ii) *The maximal number of disjoint dibonds in \mathfrak{B} is finite.*

If \mathfrak{B} is finite-corner-closed, then (i) and (ii) are also equivalent with the following statement:

- (iii) *The maximal number of disjoint and pairwise nested dibonds in \mathfrak{B} is finite.*

Proof. We start by proving the implication from (i) to (ii). Let F be a \mathfrak{B} -dijoin of D of finite size. Then, by definition, we can find at most $|F|$ many disjoint dibonds in \mathfrak{B} .

For the implication (ii) to (i) note that for any inclusion-wise maximal set \mathcal{B} of disjoint dibonds in \mathfrak{B} the set $F := \bigcup \mathcal{B}$ is a finite \mathfrak{B} -dijoin of D .

The implication from (ii) to (iii) is immediate, even if \mathfrak{B} is not finite-corner-closed.

Finally, we assume statement (iii) and that \mathfrak{B} is finite-corner-closed, and we prove statement (i).

Suppose that for some finite set $\mathcal{B} \subseteq \mathfrak{B}$ of pairwise disjoint and pairwise nested finite dibonds which is of maximum size there is some dibond $A \in \mathfrak{B}$ which is disjoint to each dibond in \mathcal{B} . Without loss of generality, let \mathcal{B} and A be chosen such that the number of dibonds in \mathcal{B} that cross A is of minimum size among all possible choices.

Let $B \in \mathcal{B}$ be chosen such that A and B cross and either $\text{in}(B)$ (first case) or $\text{out}(B)$ (second case) is inclusion-minimal among all sides of the elements of \mathcal{B} that cross A .

In the first case we consider the dicut $A \wedge B \in \mathfrak{B}^\oplus$. Note that since both A and B are dibonds, the outshore of $A \wedge B$ induces a weakly connected digraph. Hence an easy case analysis shows that any dibond in its decomposition into dibonds in \mathfrak{B} is nested with every dibond in \mathcal{B} as well as with each other. In particular, $A \wedge B$ is a dibond in \mathfrak{B} , since otherwise it would contradict the maximality of \mathcal{B} . Moreover, let A' be any dibond appearing in the decomposition of $A \vee B$ into dibonds in \mathfrak{B} . As before, we can show that A' is nested with $A \wedge B$, as well as with any dibond in \mathcal{B} which is nested with A . And since $\mathcal{B}' := (\mathcal{B} \setminus \{B\}) \cup \{A \wedge B\}$ is a set of pairwise disjoint dibonds in \mathfrak{B} and A' crosses strictly fewer dicuts in \mathcal{B}' than A crosses in \mathcal{B} , the pair \mathcal{B}' and A' contradicts the choice of \mathcal{B} and A . In the second case the same argument works with the roles of $A \wedge B$ and $A \vee B$ reversed.

In any case, this contradicts the existence of such a set \mathcal{B} and such a dibond A . Therefore, for any finite set $\mathcal{B} \subseteq \mathfrak{B}$ of pairwise disjoint and pairwise nested finite dibonds which is of maximum size the set $\bigcup \mathcal{B}$ is a finite \mathfrak{B} -dijoin. \square

Given an edge set $N \subseteq E(D)$, let $\mathfrak{B} \upharpoonright N$ denote the set $\{B \in \mathfrak{B} \mid B \subseteq N\}$. Note that $\mathfrak{B} \upharpoonright N$ is a class of finite dibonds of the contraction minor $D.N$ and if \mathfrak{B}

is finite-corner-closed, then so is $\mathfrak{B}\upharpoonright N$. The following lemma uses a standard compactness argument to show the existence of a (nested) optimal pair for D based on the existence of (nested) optimal pairs of bounded size for its finite contraction minors.

Lemma 6.5.2. *Let $n \in \mathbb{N}$. If for every finite $N \subseteq E(D)$ there is a (nested) $\mathfrak{B}\upharpoonright N$ -optimal pair (F_N, \mathcal{B}_N) for $D.N$ with $|F_N| \leq n$, then there is a (nested) \mathfrak{B} -optimal pair (F, \mathcal{B}) for D .*

Proof. Let \mathcal{B} be a maximal (nested) set of disjoint dicuts in \mathfrak{B} . Note that $|\mathcal{B}| \leq n$, since otherwise a subset $\mathcal{B}' \subseteq \mathcal{B}$ of size $n + 1$ would contradict the assumption for $N = \bigcup \mathcal{B}'$.

Let $N \subseteq E(D)$ be a finite set of edges such that $\bigcup \mathcal{B} \subseteq N$ holds. Since $D.N$ is a finite weakly connected digraph, there exists a (nested) $\mathfrak{B}\upharpoonright N$ -optimal pair (F_N, \mathcal{B}_N) for $D.N$ by assumption. By the choice of N and Lemma 6.2.12 we know that each element of \mathcal{B} is also a finite dicut of $D.N$. Furthermore, each finite dicut in $D.N$ is also one in D and, thus, \mathcal{B}_N is a set of disjoint finite dicuts in D . Hence, $|\mathcal{B}| = |\mathcal{B}_N| = |F_N|$. Using that the elements in \mathcal{B} are pairwise disjoint (and nested) finite dicuts, we get that (F_N, \mathcal{B}) is a (nested) $\mathfrak{B}\upharpoonright N$ -optimal pair for $D.N$ as well. Given a finite edge set $M \supseteq N$ with a (nested) $\mathfrak{B}\upharpoonright M$ -optimal pair (F_M, \mathcal{B}_M) for $D.M$ we obtain that (F_M, \mathcal{B}) is also a nested optimal pair for $D.N$.

Note that for any finite edge set $N \subseteq E(D)$ satisfying $\bigcup \mathcal{B} \subseteq N$ there are only finitely many possible edge sets $F_N \subseteq \bigcup \mathcal{B}$ such that (F_N, \mathcal{B}) is a (nested) $\mathfrak{B}\upharpoonright N$ -optimal pair for $D.N$. Hence, we get via the compactness principle an edge set $F \subseteq \bigcup \mathcal{B}$ with $|F \cap B| = 1$ for every $B \in \mathcal{B}$ such that (F, \mathcal{B}) is a (nested) $\mathfrak{B}\upharpoonright M$ -optimal pair for $D.M$ for every finite edge set $M \subseteq E(D)$ satisfying $\bigcup \mathcal{B} \subseteq M$.

We claim that (F, \mathcal{B}) is a (nested) \mathfrak{B} -optimal pair for D . We already know by definition that \mathcal{B} is a (nested) set of disjoint finite dicuts in \mathfrak{B} and that $F \subseteq \bigcup \mathcal{B}$ with $|F \cap B| = 1$ for every $B \in \mathcal{B}$. It remains to check that F is a \mathfrak{B} -dijoin of D . So let $B' \in \mathfrak{B}$. Then the set $N' := B' \cup \bigcup \mathcal{B}$ is also finite and B' is a finite dicut of $D.N'$. Since (F, \mathcal{B}) is also a nested optimal pair for $D.N'$, we know that $F \cap B' \neq \emptyset$ holds, which proves that F is a \mathfrak{B} -dijoin of D . \square

Lemmas 6.5.1 and 6.5.2 together with Theorem 6.1.4 yield Theorem 6.1.7 (i), (ii) and (iii).

6.5.2. Every edge lies in only finitely many dibonds and reductions to this case

We continue with another special case. Its proof is also based on a compactness argument. However, we need to choose the set up for the argument more carefully.

Lemma 6.5.3. *Conjecture 6.1.5 holds for weakly connected digraphs in which every edge lies in only finitely many finite dibonds.*

Proof. Let D be a weakly connected digraph where every edge lies in only finitely many finite dibonds. For an edge $e \in E(D)$ let \mathcal{B}_e denote the set of finite dibonds of D that contain e . Our assumption on D implies that \mathcal{B}_e is a finite set. For a finite set \mathcal{B} of finite dibonds of D we define $\hat{\mathcal{B}} = \cup\{\mathcal{B}_e \mid e \in \cup\mathcal{B}\}$. Again our assumption on D implies that $\hat{\mathcal{B}}$ is finite. Note that $\mathcal{B} \subseteq \hat{\mathcal{B}}$ holds.

Given a finite set \mathcal{B} of finite dibonds of D , we call $(F_{\mathcal{B}}, \mathcal{B}')$ a *nested pre-optimal pair* for $D.(\cup\mathcal{B})$ if the following hold:

1. $F_{\mathcal{B}}$ intersects every element of \mathcal{B} ,
2. $\mathcal{B}' \subseteq \hat{\mathcal{B}}$,
3. the elements of \mathcal{B}' are pairwise nested,
4. $F_{\mathcal{B}} \subseteq \cup\mathcal{B}'$, and
5. $|F_{\mathcal{B}} \cap B'| = 1$ for every $B' \in \mathcal{B}'$.

We know that for every finite set \mathcal{B} of finite dibonds of D there exists a nested pre-optimal pair for $D.(\cup\mathcal{B})$, since a nested optimal pair for $D.(\cup\hat{\mathcal{B}})$ is one and it exists by Theorem 6.1.4. However, there can only be finitely many nested pre-optimal pairs for $D.(\cup\mathcal{B})$ as $\cup\hat{\mathcal{B}}$ is finite.

Now let \mathcal{B}_1 and \mathcal{B}_2 be two finite sets of finite dibonds of D with $\mathcal{B}_1 \subseteq \mathcal{B}_2$, and let $(F_{\mathcal{B}_2}, \mathcal{B}'_2)$ be a nested pre-optimal pair for $D.(\cup\mathcal{B}_2)$. Then $(F_{\mathcal{B}_2} \cap \cup\mathcal{B}_1, \mathcal{B}'_2 \cap \hat{\mathcal{B}}_\infty)$ is a nested pre-optimal pair for $D.(\cup\mathcal{B}_1)$. Now we get by the compactness principle an edge set $F'_D \subseteq E(D)$ and a set \mathcal{B}_D of finite dibonds of D such that $(F'_D \cap \cup\mathcal{B}, \mathcal{B}_D \cap \hat{\mathcal{B}})$ is a nested pre-optimal pair for $D.(\cup\mathcal{B})$ for every finite set \mathcal{B} of finite dibonds of D . Furthermore, let $F_D \subseteq F'_D$ be such that each element of F_D lies on a finite dibond of D and $(F_D \cap \cup\mathcal{B}, \mathcal{B}_D \cap \hat{\mathcal{B}})$ is still a nested pre-optimal pair for $D.(\cup\mathcal{B})$ for every finite set \mathcal{B} of finite dibonds of D .

We claim that (F_D, \mathcal{B}_D) is a nested optimal pair for D . First we verify that F_D is a finitary dijoin of D . Let B be any finite dibond of D . Then F_D meets B , because $(F_D \cap B, \mathcal{B}_D \cap \widehat{B})$ is a nested pre-optimal pair for $D.B$. So F_D is a finitary dijoin of D .

Next consider any element $e \in F_D$. By definition of F_D we know that $e \in B_e$ holds for some finite dibond B_e of D . Using again that $(F_D \cap B_e, \mathcal{B}_D \cap \widehat{B_e})$ is a nested pre-optimal pair for $D.B_e$, we get that $e \in \bigcup \mathcal{B}_D$. So the inclusion $F_D \subseteq \bigcup \mathcal{B}_D$ is valid.

Given any $B_D \in \mathcal{B}_D$ we know that $(F_D \cap B_D, \mathcal{B}_D \cap \widehat{B_D})$ is a nested pre-optimal pair for $D.B_D$. Hence, $|F_D \cap B| = 1$ holds for every $B \in \mathcal{B}_D \cap \widehat{B_D}$. Especially, $|F_D \cap B_D| = 1$ is true because $B_D \in \mathcal{B}_D \cap \widehat{B_D}$.

Finally, let us consider two distinct elements B_1 and B_2 of \mathcal{B}_D . We know that $(F_D \cap (B_1 \cup B_2), \mathcal{B}_D \cap \widehat{B_1, B_2})$ is a nested pre-optimal pair for $D.(B_1 \cup B_2)$. Therefore, B_1 and B_2 are disjoint and nested. This shows that (F_D, \mathcal{B}_D) is a nested optimal pair for D and completes the proof of this lemma. \square

The next lemma can be used together with Lemma 6.5.3 to deduce that Conjecture 6.1.5 holds for weakly connected digraphs without infinite dibonds.

Lemma 6.5.4. *In a weakly connected digraph D where some edge e lies in infinitely many finite dibonds of D there is an infinite dibond containing e .*

Proof. We construct with a compactness argument a dibond containing $e =: vw$ that is distinct from every finite dibond.

Let $W \subseteq V(D)$ be finite with $v, w \in W$. Consider the set \mathcal{B}_W consisting of those bipartitions (A, B) of W with $v \in A$ and $w \in B$ such that $\vec{E}(B, A)$ is empty, but $\vec{E}(A, B)$ contains no finite dibond of D . Obviously, \mathcal{B}_W is finite. For any dibond $\vec{E}(X, Y)$ containing e that is not contained in $E(D[W])$ the bipartition $(X \cap W, Y \cap W)$ is in \mathcal{B}_W . And since e lies in infinitely many dibonds, such a dibond always exists. Moreover, for $W \subseteq W'$ and $(A, B) \in \mathcal{B}_{W'}$ we have $(A \cap W, B \cap W) \in \mathcal{B}_W$. Hence by compactness there is a bipartition (A, B) of $V(D)$ such that $(A \cap W, B \cap W) \in \mathcal{B}_W$ for every finite $W \subseteq V(D)$ with $v, w \in W$. Now $\vec{E}(A, B)$ is a dicut of D which does not contain any finite dibond of D , since these properties would already be witnessed for some finite $W \subseteq V(D)$. Therefore, $\vec{E}(A, B)$ is an infinite dicut of D containing only infinite dibonds of D . \square

As noted before, we obtain the following corollary.

Corollary 6.5.5. *Conjecture 6.1.5 holds for weakly connected digraphs without infinite dibonds.* \square

We close this section with a last special case where we can show that Conjecture 6.1.5 holds.

Corollary 6.5.6. *Conjecture 6.1.5 holds for rayless weakly connected digraphs.*

Proof. Let D be a rayless weakly connected digraph. We know by Proposition 6.2.20 that (D/\sim) is rayless as well, and by Proposition 6.2.13 that D/\sim is weakly connected and finitely separable. So we obtain from Corollary 6.2.11 that D/\sim has no infinite dibond. Now Corollary 6.5.5 implies that Conjecture 6.1.5 is true in the digraph D/\sim . Using again that D/\sim is finitely separable, any nested optimal pair for D/\sim directly translates to one for D by Lemma 6.4.1. Hence, Conjecture 6.1.5 is true for D as well. \square

6.6. A matching problem about infinite hypergraphs

In this section we discuss how Conjecture 6.1.3 is related to more general questions about infinite hypergraphs, where the initial one was posted by Aharoni. We shall give an example, which then negatively answers Aharoni's original question. However, we leave a modified version as a conjecture which then is still open. Then we shall strengthen the latter conjecture to obtain a new one, which is closely related to Conjecture 6.1.3 and the infinite version of Menger's Theorem. Before we can do this we have to give some definitions and set notation.

Let us fix a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$. We call \mathcal{H} *simple*, if no hyperedge is contained in another one. Given a set $F \subseteq \mathcal{E}$, we shall write $\mathcal{H}[F]$ for the hypergraph $(\cup F, F)$ and call it a *subhypergraph* of \mathcal{H} . Moreover, a subhypergraph \mathcal{K} of \mathcal{H} is called *finite*, if there exists some finite $F \subseteq \mathcal{E}$ such that $\mathcal{K} = \mathcal{H}[F]$. Note that \mathcal{K} might have infinitely many vertices since a hyperedge can contain infinitely many vertices. We call \mathcal{H} *locally finite* if each vertex of \mathcal{H} lies in only finitely hyperedges. Furthermore, we say that \mathcal{H} has *finite character* if no hyperedge of \mathcal{H} contains infinitely many vertices.

A set of hyperedges $\mathcal{M} \subseteq \mathcal{E}$ is called a *matching of \mathcal{H}* if any two hyperedges in \mathcal{M} are pairwise disjoint. A set of vertices $A \subseteq \mathcal{V}$ is called a *cover of \mathcal{H}* if every hyperedge of \mathcal{H} contains a vertex from A . Now the hypergraph \mathcal{H} is said to have the *Kőnig property* if a pair (\mathcal{M}, A) exists such that the following statements hold:

1. \mathcal{M} is a matching of \mathcal{H} .
2. A is a cover of \mathcal{H} .
3. $A \subseteq \bigcup \mathcal{M}$.
4. $|M \cap A| = 1$ for every $M \in \mathcal{M}$.

We call such a pair (\mathcal{M}, A) an *optimal pair for \mathcal{H}* .

Now we are able to state the original problem on infinite hypergraphs posted by Aharoni.

Problem 6.6.1. [1, Prob. 6.7] *Let \mathcal{H} be a hypergraph and suppose that every finite subhypergraph of \mathcal{H} has the Kőnig property. Does then \mathcal{H} have the Kőnig property?*

We shall now point out that, in full generality, this problem has a negative answer by stating a certain infinite hypergraph \mathcal{H} . For this, consider the digraph in Figure 6.1.1, call it D . Let \mathcal{B} denote the set of all dicuts of D . We now define the hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ by setting $\mathcal{V} = E(D)$ and $\mathcal{E} = \mathcal{B}$. As discussed in the introduction, \mathcal{H} does not have the Kőnig property, since we cannot even find two disjoint hyperedges, but we need infinitely many vertices of \mathcal{V} to cover all hyperedges. However, for every non-empty finite subset F of \mathcal{E} we can find one vertex of \mathcal{V} covering all hyperedges of (\mathcal{V}, F) .

As noticed in the introduction, \mathcal{H} does not have any finite hyperedges. This motivates us to modify Problem 6.6.1 to include only hypergraphs of finite character.

Conjecture 6.6.2. *Let \mathcal{H} be a hypergraph of finite character and suppose that every finite subhypergraph of \mathcal{H} has the Kőnig property. Then \mathcal{H} has the Kőnig property.*

Variations of Problem 6.6.1, particularly Conjecture 6.6.2, are very general problems about infinite hypergraphs and probably difficult to answer. Not much is known about them, not even partial answers. However, relaxed questions

involving fractional matchings and covers have more successfully been studied, see [1, Section 6] for a brief survey on such results.

Now we modify Conjecture 6.6.2 even further yielding the following stronger conjecture.

Conjecture 6.6.3. *Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be a hypergraph of finite character and suppose that for every finite $F \subseteq \mathcal{E}$ there exists some finite set $F' \subseteq \mathcal{E}$ such that $F \subseteq F'$ and (\mathcal{V}, F') has the König property. Then \mathcal{H} has the König property.*

Although even stronger than Conjecture 6.6.2, this conjecture is very important, because it is closely related to the infinite version of Menger's Theorem and Conjecture 6.1.3. In case Conjecture 6.6.3 is verified, this would not only give another proof of the infinite version of Menger's Theorem proved by Aharoni and Berger [2] but also imply Conjecture 6.1.3. The deductions are very similar in both of these cases; namely by defining a suitable auxiliary hypergraph.

For Menger's Theorem where an infinite graph $G = (V, E)$ is given as well as vertex sets $A, B \subseteq V$, we define an auxiliary hypergraph $\mathcal{H}_{A,B} = (\mathcal{V}, \mathcal{E})$ as follows. The vertex set \mathcal{V} of \mathcal{H} consists precisely of those vertices of G that lie on any A - B path in G . Now a subset $F \subseteq \mathcal{V}$ forms a hyperedge of \mathcal{H} if and only if F is the vertex set of an A - B path in G . For every finite set F of A - B paths consider the finite subgraph G_F induced by the vertex set spanned by the paths in F . Let F' be the set of all A - B paths in G_F . Note that F' is a finite superset of F , for which by Menger's Theorem (\mathcal{V}, F') has the König property. Hence, verifying Conjecture 6.6.3 would imply the infinite version of Menger's Theorem.

With respect to the Infinite Lucchesi-Younger Conjecture, consider an infinite weakly connected digraph $D = (V, E)$. We define a auxiliary hypergraph $\mathcal{H}_D = (\mathcal{V}, \mathcal{E})$ as follows. We set $\mathcal{V} = E$. Furthermore, a set $B \subseteq E$ forms a hyperedge of \mathcal{H}_D if and only if B defines a dibond of D . Given a finite set F of dibonds of D we set F' to be the minimal finite-corner-closed set of dibonds containing F . Note that F' is still a finite set. Now by Theorem 6.1.1, respectively Theorem 6.1.2, (\mathcal{V}, F') has the König property. So a positive answer to Conjecture 6.6.3 would imply Conjecture 6.1.3.

Now we conclude this section by translating some results based on compactness arguments of the previous section to yield also verified affirmative answers for special cases of Conjecture 6.6.3. Note first that an analogue version of Lemma 6.5.1

is true for hypergraphs as well:

Lemma 6.6.4. *Let \mathcal{H} be a hypergraph of finite character. Then the following statements are equivalent:*

- (i) \mathcal{H} has a finite cover.
- (ii) The maximal size a matching of \mathcal{H} can have is finite.

Using Lemma 6.6.4 we can verify the following special case via the same compactness argument as used for Lemma 6.5.2.

Lemma 6.6.5. *Let \mathcal{H} be a hypergraph of finite character satisfying the premise of Conjecture 6.6.3. Furthermore, let \mathcal{H} satisfy one of the following conditions:*

- (i) \mathcal{H} has a finite cover.
- (ii) There is a finite maximal size a matching of \mathcal{H} can have.

Then \mathcal{H} has the the Kőnig property.

The other result from Section 6.5 we can lift to hypergraphs is Lemma 6.5.3. Again the proof depends on a compactness argument which can immediately be translated into the setting for hypergraphs.

Lemma 6.6.6. *Let \mathcal{H} be a locally finite hypergraph of finite character satisfying the premise of Conjecture 6.6.3. Then \mathcal{H} has the the Kőnig property.*

6.7. Extending an algorithmic proof of Frank

In this section we extend several parts of the proof of Frank [20, Section 9.7.2] for Theorem 6.1.1 to infinite digraphs. This proof is based on the ideas of the negative circuit method developed for more general submodular frameworks by Fujishige [21] and Zimmermann [52]. Instead of just starting with a dijoin of minimum size, the idea of Frank’s proof is to start with any dijoin and algorithmically “improve” it with the help of cycles of negative cost in an auxiliary digraph whose definition depends on the dijoin. Once the dijoin can no longer be “improved” some structural properties of the auxiliary graph help in fining the desired set of dibonds which together with the dijoin form a nested optimal pair.

This improvement process in finite digraphs strictly reduces its size. We conjecture that even in infinite digraphs we can still obtain a dijoin which can no longer be “improved” in this manner. For such a dijoin the auxiliary graph will still exhibit these structural properties. In certain cases this will allow us to similarly obtain the set of dibonds, although this is not universally true.

6.7.1. Tightness and an auxiliary graph

Let D be a weakly connected digraph. Let \mathfrak{B} be a finite-corner-closed class of finite dibonds of D and let F be a \mathfrak{B} -dijoin.

In this subsection we will introduce some terminology for the dicuts that can appear as part of a \mathfrak{B} -optimal pair for D and prove some properties about them. Furthermore, we will introduce an auxiliary graph depending on D and F using this terminology.

A dibond $B \in \mathfrak{B}$ is called *F-tight* if $|F \cap B| = 1$ holds. A dicut $B \in \mathfrak{B}^\oplus$ is called *F-tight* if $B = \bigoplus \mathcal{B}$ for some set $\mathcal{B} \subseteq \mathfrak{B}$ of *F-tight* dibonds.

Our first useful observations about *F-tight* dicuts is that they themselves are closed under taking finite corners. In particular, if two *F-tight* dibonds cross, then their corners will again be *F-tight* dibonds.

Lemma 6.7.1. *Let $B_1, B_2 \in \mathfrak{B}$ be *F-tight* dibonds such that $\text{in}(B_1) \cap \text{in}(B_2) \neq \emptyset$ and $\text{out}(B_1) \cap \text{out}(B_2) \neq \emptyset$. Then both $B_1 \wedge B_2$ and $B_1 \vee B_2$ are *F-tight* dibonds in \mathfrak{B} .*

Proof. Since F is a \mathfrak{B} -dijoin and \mathfrak{B} is finite-corner-closed, we get with Remark 6.2.2 that $|F \cap (B_1 \wedge B_2)| = 1 = |F \cap (B_1 \vee B_2)|$. Since $B_1 \wedge B_2$ and $B_1 \vee B_2$ are both in \mathfrak{B}^\oplus , they must be dibonds in \mathfrak{B} , since otherwise there would be a dibond in \mathfrak{B} in one of their decompositions which is not met by F . \square

A generalisation shows that corners of *F-tight* dicuts are again *F-tight*.

Lemma 6.7.2. *Let $B_1, B_2 \in \mathfrak{B}^\oplus$ be *F-tight*.*

1. *If $B_1 \wedge B_2 \neq \emptyset$, then $B_1 \wedge B_2 \in \mathfrak{B}^\oplus$ is *F-tight*.*
2. *If $B_1 \vee B_2 \neq \emptyset$, then $B_1 \vee B_2 \in \mathfrak{B}^\oplus$ is *F-tight*.*

Proof. First assume that B_1 and B_2 are dibonds. If $\text{in}(B_1) \cap \text{in}(B_2) = \emptyset$, then $B_1 \wedge B_2 = \emptyset$ and $B_1 \vee B_2 = B_1 \oplus B_2$ is F -tight. If $\text{out}(B_1) \cap \text{out}(B_2) = \emptyset$, then $B_1 \vee B_2 = \emptyset$ and $B_1 \wedge B_2 = B_1 \oplus B_2$ is F -tight. Hence, with Lemma 6.7.1 we are done in all cases.

Otherwise, with Remark 6.2.3 and the statement for dibonds, we get that $B_1 \wedge B_2$ and $B_1 \vee B_2$ are F -tight as well. \square

Remark 6.7.3. Let $u, v \in V(D)$. There is an F -tight dicut $B' \in \mathfrak{B}^\oplus$ with $u \in \text{out}(B')$ and $v \in \text{in}(B')$ if and only if there is an F -tight dibond $B \in \mathfrak{B}$ with $u \in \text{out}(B)$ and $v \in \text{in}(B)$. \square

We define an auxiliary multi-digraph D_F as follows:

First we define the vertex set of D_F as

$$V(D_F) := V(D).$$

In order to define the edge set of D_F we have to define two other edge sets first. We define a set F^* that contains for each $f \in F$ with head v and tail w an edge f^* with head w and tail v , such that the map $f \mapsto f^*$ is a bijection between F and F^* .

Furthermore, we define

$$J_F := \{uv \mid \text{no } F\text{-tight dibond in } \mathfrak{B} \text{ has } v \text{ in its in-shore and } u \text{ in its out-shore}\}.$$

We call the edges in J_F *jumping edges*.

At last, we define the edge set of D_F as the disjoint union of all these edge sets:

$$E(D_F) := E(D) \dot{\cup} F^* \dot{\cup} J_F.$$

We define an integer-valued cost function $c : E(D_F) \rightarrow \mathbb{Z}$ to be constant 1 on $E(D)$, constant -1 on F^* and constant 0 on J_F .

Note that if between two vertices $v, w \in V(D)$ there is a directed v - w -path in D , then $wv \in J_F$. We call a jumping edge for which there is such a path *basic*.

Let us note a few properties of the auxiliary digraph.

Lemma 6.7.4. J_F is transitive, i.e. if $uv, vw \in J_F$ then $uw \in J_F$ as well.

Proof. Suppose for a contradiction that $uv, vw \in J_F$ and $uw \notin J_F$. Then there is an F -tight dicut $B \in \mathfrak{B}$ with $u \in \text{out}(B)$ and $w \in \text{in}(B)$. Now either $v \in \text{in}(B)$, contradicting that $uv \in J_F$, or $v \in \text{out}(B)$, contradicting that $vw \in J_F$. \square

Lemma 6.7.5. *The auxiliary graph D_F is strongly connected.*

Proof. Assume for a contradiction that there is a dicut B of D_F . Since D is weakly connected, B contains an edge e of D with tail v and head w . But then D_F contains the basic jumping edge $wv \in J_F$, contradicting that B is a dicut of D_F . \square

One of the main tools we will use to obtain F -tight dicuts is the following lemma.

Lemma 6.7.6 (Finite Separation Lemma). *Let $U, W \subseteq V(D)$ be two finite vertex sets such that $uw \notin J_F$ for all $u \in U$ and $w \in W$. Then there is a finite F -tight dicut $B \in \mathfrak{B}^\oplus$ with $U \subseteq \text{out}(B)$ and $W \subseteq \text{in}(B)$.*

Proof. Since for every pair $(u, w) \in U \times W$ there is no jumping edge $uw \in J_F$, there is an F -tight dibond $B_{u,w} \in \mathfrak{B}$ with $u \in \text{out}(B_{u,w})$ and $w \in \text{in}(B_{u,w})$.

For a fixed $u \in U$, the dicut

$$B_u := \bigvee_{w \in W} B_{u,w}$$

satisfies $u \in \text{out}(B_u)$, $W \subseteq \text{in}(B_u)$, is in \mathfrak{B}^\oplus , since \mathfrak{B} is finite-corner-closed, and by Lemma 6.7.2 it is F -tight. Then the dicut

$$B := \bigwedge_{u \in U} B_u = \bigwedge_{u \in U} \left(\bigvee_{w \in W} B_{u,w} \right)$$

is the desired dicut again by Lemma 6.7.2. \square

Let us call a (finite or infinite) dicut B F -rigid if there is no jumping edge $vw \in J_F$ with $v \in \text{out}(B)$ and $w \in \text{in}(B)$,

For finite dicuts, the notions of F -tightness and F -rigidity coincide.

Lemma 6.7.7. *A finite dicut $B \in \mathfrak{B}^\oplus$ is F -tight if and only if it is F -rigid.*

Proof. If there is $vw \in J_F$ with $v \in \text{out}(B)$ and $w \in \text{in}(B)$, then B is obviously not F -tight.

So assume that there is no such jumping edge. For each $v \in V(D)$ let C_v denote the weak component of $D - B$ containing v . Moreover, let C'_v denote a finite weakly connected subdigraph of C_v containing $\partial(C_v)$. Then consider

$$U := \bigcup \{V(C'_u) \mid u \text{ is a tail of an edge in } B\} \text{ and}$$

$$W := \bigcup \{V(C'_w) \mid w \text{ is a head of an edge in } B\}.$$

Applying Lemma 6.7.6 to U and W yields that there is an F -tight dicut $B' \in \mathfrak{B}^\oplus$ containing B . Moreover, B' does not separate the vertices in the boundary of any weak component C_v of $D - B$.

Let $\oplus \mathcal{B}'$ be a decomposition of B' into F -tight dibonds in \mathfrak{B} . Assume for a contradiction that some $A \in \mathcal{B}'$ contains an edge $uw \in B$ as well as an edge $xy \in B' \setminus B$.

Then x and y both belong to the same weak component C_v of $D - B$ for some head or tail v of an edge in B . Either $\text{in}(A) \cap C_v$ or $\text{out}(A) \cap C_v$ are disjoint from $\partial(C_v)$, since one of the sides contains C'_v . But the side disjoint from $\partial(C_v)$ is actually a subset of $C_v \setminus C'_v$ and hence contains neither u nor w , since one of them is in $\partial(C_v)$ and the other is not in C_v . This contradicts that $uw \in A$.

Thus every dibond $A \in \mathcal{B}'$ is either contained in B or disjoint from B . With $\mathcal{B} := \{A \in \mathcal{B}' \mid A \subseteq B\}$ we get the desired decomposition $B = \oplus \mathcal{B}$ into F -tight dibonds. \square

While an infinite dibond (which is not in \mathfrak{B}^\oplus) may be F -rigid but not F -tight, we pose the following question about the equivalence of these notions for all dicuts in \mathfrak{B}^\oplus .

Question 6.7.8. *Is every F -rigid dicut $B \in \mathfrak{B}^\oplus$ also F -tight?*

As a corollary of the Finite Separation Lemma we can also obtain information on how the edges of F interact with an F -rigid dicut.

Corollary 6.7.9. *Let B be an F -rigid dicut. Let $C \subseteq \text{out}(B)$, $C' \subseteq \text{in}(B)$ be weak components of $D - B$. Then $|E(C, C') \cap F| \leq 1$.*

Proof. Assume for a contradiction that $u_1w_1, u_2w_2 \in E(C, C') \cap F$. Let P be an undirected u_1 - u_2 -path in C and let Q be an undirected w_1 - w_2 -path in C' . Applying the Finite Separation Lemma 6.7.6 to $U := V(P)$ and $W := V(Q)$ yields an F -tight dicut $B' \in \mathfrak{B}$ with $V(P) \subseteq \text{out}(B)$ and $V(Q) \subseteq \text{in}(B)$. Since no dibond in the decomposition of B' into F -tight dibonds in \mathfrak{B} contains an edge of P or Q , one of these dibonds contains both u_1w_1 and u_2w_2 , contradicting that it is F -tight. \square

With Conjecture 6.1.5 in mind, we are not just interested in any decomposition of a dicut into F -tight dibonds, but one where the set of dibonds in the decomposition

is nested. We call a dicut $B \in \mathfrak{B}^\oplus$ *nestedly F -tight* if $B = \bigoplus \mathcal{B}$ for some nested set $\mathcal{B} \subseteq \mathfrak{B}$ of F -tight dibonds.

Lemma 6.7.10. *Every finite F -tight $B \in \mathfrak{B}^\oplus$ is nestedly F -tight.*

Proof. Let \mathcal{B} be a decomposition of B into finitely many F -tight dibonds such that the number of pairs $(A, A') \in \mathcal{B} \times \mathcal{B}$ such that A and A' cross is minimal. We show that \mathcal{B} is nested. Suppose for a contradiction that $A, A' \in \mathcal{B}$ are crossing. Then

$$\mathcal{B}' := ((\mathcal{B} \setminus \{A, A'\}) \cup \{A \wedge A', A \vee A'\})$$

is a decomposition of B into F -tight dibonds in \mathfrak{B} by Lemma 6.7.1. An easy calculation shows that every dibond which is nested with both A and A' is also nested with both $A \wedge A'$ and $A \vee A'$. Similarly, a dibond which crosses both $A \wedge A'$ and $A \vee A'$ also crosses both A and A' . Hence \mathcal{B}' contradicts the minimality choice of \mathcal{B} . \square

As above, the question whether this is in reality a stronger condition remains open.

Question 6.7.11. *Is every F -tight dicut $B \in \mathfrak{B}^\oplus$ also nestedly F -tight?*

6.7.2. Feasible dijoints

Let D be a weakly connected digraph. Let \mathfrak{B} be a finite-corner-closed class of finite dibonds of D .

In this subsection we will describe a process which will allow us to change a given \mathfrak{B} -dijoin to a different \mathfrak{B} -dijoin. This process is intended to “improve” a dijoin when applying it to negative directed cycles in the auxiliary graph. While in the finite case this improvement process strictly decreases the size of the dijoin, we will arrive at a dijoin whose auxiliary graph does not contain any negative cycles. For the infinite case the existence of such a dijoin whose auxiliary graph does not contain any negative cycle is not so straight forward. For certain classes of graphs this can be proven by a compactness argument, cf. Section 6.8.

Given an edge set $Z \subseteq E(D_F)$ and a vertex set $X \subseteq V(D)$ we denote by $J(Z, X)$ the set $Z \cap \delta_{D_F}^-(X) \cap J_F$. We call a set $Z \subseteq E(D_F)$ *exchangeable* if

(E1) for every finite dicut $\delta_D^-(X) \in \mathfrak{B}^\oplus$ the equation

$$|Z \cap \delta_{D_F}^-(X)| = |Z \cap \delta_{D_F}^+(X)| < \infty$$

holds;

(E2) for every finite dicut $\delta_D^-(X) \in \mathfrak{B}^\oplus$ with $|J(Z, X)| \geq 2$ there is an F -tight dibond $\delta^-(Y) \in \mathfrak{B}$ such that $J(Z, X \cap Y)$ and $J(Z, X \cup Y)$ partition $J(Z, X)$.

Remark 6.7.12. Every directed cycle C in D_F for which every jumping edge it contains is basic is exchangeable. \square

We continue with a lemma that allows us to modify a given \mathfrak{B} -dijoin to another \mathfrak{B} -dijoin given some exchangeable set.

Lemma 6.7.13 (Exchange Lemma). *If $Z \subseteq E(D_F)$ is exchangeable, then*

$$F \circ Z := (F \setminus \{f \in F \mid f^* \in (Z \cap F^*)\}) \cup (Z \cap E(D))$$

is a \mathfrak{B} -dijoin of D .

Proof. By property (E1), it suffices to show that $|\delta_D^-(X) \cap F| \geq |J(Z, X)| + 1$ for each dicut $\delta_D^-(X) \in \mathfrak{B}^\oplus$, since then either $Z \cap \delta_D^-(X)$ is non-empty or $\{f \in F \cap \delta_D^-(X) \mid f^* \in Z \cap F^*\}$ is a proper subset of $F \cap \delta_D^-(X)$. We show this claim by induction on $|J(Z, X)|$.

For $|J(Z, X)| = 0$ this is trivial. For $|J(Z, X)| = 1$, note that $|\delta_D^-(X) \cap F| > 1$ since $\delta_D^-(X)$ is not F -tight.

For $|J(Z, X)| \geq 2$ let $\delta^-(Y)$ be as in property (E2). Since \mathfrak{B} is finite-corner-closed, the corners $\delta_D^-(X \cap Y)$ and $\delta_D^-(X \cup Y)$ are in \mathfrak{B}^\oplus . Now $|J(Z, X \cap Y)|$ and $|J(Z, X \cup Y)|$ are both less than $|J(Z, X)|$ by property (E2). Hence we get $|\delta_D^-(X \cap Y) \cap F| \geq |J(Z, X \cap Y)| + 1$ and $|\delta_D^-(X \cup Y) \cap F| \geq |J(Z, X \cup Y)| + 1$ by induction. Moreover, since $\delta^-(Y)$ is an F -tight dibond we get that $|J(Z, Y)| = 0$ and $|\delta_D^-(Y) \cap F| = 1$. Finally these inequalities together with Remark 6.2.2 and property (E2) yield

$$\begin{aligned} |\delta_D^-(X) \cap F| &= |\delta_D^-(X \cap Y) \cap F| + |\delta_D^-(X \cup Y) \cap F| - |\delta_D^-(Y) \cap F| \\ &\geq |J(Z, X \cap Y)| + 1 + |J(Z, X \cup Y)| + 1 - 1 \\ &= |J(Z, X \cap Y)| + |J(Z, X \cup Y)| + 1 \\ &= |J(Z, X)| + 1, \end{aligned}$$

finishing the proof. \square

Before we continue, we fix the following notation. Let C be a directed cycle in any digraph D . If $e \in E(D)$ is an edge both of whose endvertices lie on C , then there exists a unique directed cycle in $C + e$ that contains e . We shall call that cycle $C(e)$.

In order to prove a lemma telling us that we can obtain an exchangeable negative directed cycle from any negative directed cycle in D_F , we need the following technical lemma.

Lemma 6.7.14. *Let D be a digraph, let $Z \subseteq D$ a directed cycle and let $M = \{s_i t_i \in E(Z) \mid i \in [m]\}$ a set of $m \geq 2$ edges such that $s_i t_{i+1} \in E(D)$ for every $i \in [m]$ (where $t_{m+1} := t_1$). Then there exists an integer $q > 0$ such that each edge of $E(Z) \setminus M$ lies in precisely q many of the directed cycles of $\{C(s_i t_{i+1}) \mid i \in [m]\}$.*

Proof. Let $\mathcal{C} := \{C(s_i t_{i+1}) \mid i \in [m]\}$. We prove the statement by induction on m . For $m = 2$ we immediately see that $q = 1$ holds. Now suppose the statement holds for some integer $m \geq 2$ with any digraph and any cycle that satisfy the conditions. Let a digraph D , a cycle Z and a set M as in the statement be given such that $|M| = m + 1$. We define the auxiliary digraph $D' := (D - \{s_1 t_2, s_2 t_3\} + s_1 t_3) / s_2 t_2$. Note that in D' , the set $E(Z) \setminus \{s_2 t_2\}$ forms a directed cycle Z' containing the m edges of the set $M' := M \setminus \{s_2 t_2\}$. Applying the induction hypothesis to M' together with D' and Z' yields some integer $q' > 0$ as in the statement of the lemma.

Next we again consider D , Z and M , and distinguish two cases: If the directed s_1 - t_3 path on Z contains $s_2 t_2$, then every edge of $E(Z) \setminus M$ lies in precisely $q' + 1$ many directed cycles of \mathcal{C} . If the directed s_1 - t_3 path on Z does not contain $s_2 t_2$, then each edge of $E(Z) \setminus M$ lies in precisely q' many directed cycles of \mathcal{C} . Since we get in both cases an integer as required, this completes the proof of the lemma. \square

The following lemma uses the same ideas as [20, Lemma 9.7.13], including Lemma 6.7.14, whose proof has been omitted in [20, Lemma 9.7.13]. We include a full proof for the sake of completeness.

Lemma 6.7.15. *For every negative directed cycle C in D_F there is an exchangeable negative directed cycle Z of D_F with $Z \setminus J_F \subseteq C \setminus J_F$.*

Proof. Let C be a negative directed cycle in D_F . If there is a chord e of C in J_F such that $C(e)$ is still negative and strictly smaller than C , then we consider

this cycle instead. Iterating this process yields a negative directed cycle Z with $Z \setminus J_F \subseteq C \setminus J_F$ such that no further such chord exists. We claim that Z is exchangeable.

Since Z is a finite directed cycle it trivially satisfies property (E1). For property (E2) we consider the following claim.

Claim. *There is an enumeration x_1y_1, \dots, x_ny_n of $Z \cap J_F$ such that $x_jy_i \notin J_F$ for all $i, j \in [n]$ with $1 \leq i < j \leq n$.*

First we show that if this claim is true, then property (E2) holds. Consider a dicut $\delta_D^-(X) \in \mathfrak{B}$ with $x_ky_k, x_\ell y_\ell \in J(Z, X)$ for some $k, \ell \in [n]$ with $k < \ell$ such that $x_iy_i \notin J(Z, X)$ for all i with $k < i < \ell$. Then applying Lemma 6.7.6 to $U := \{x_j \mid \ell \leq j \leq n\}$ and $W := \{y_i \mid 1 \leq i \leq k\}$ gives an F -tight dicut $\delta^-(Y)$ with $U \subseteq V(D) \setminus Y$ and $W \subseteq Y$. By Lemma 6.7.7 we have that $y_j \in V(D) \setminus Y$ for all j with $\ell \leq j \leq n$ and $x_i \in Y$ for all i with $1 \leq i \leq k$. Hence with this Y we get the desired bipartition of $J(Z, X)$ into $J(Z, X \cap Y)$ and $J(Z, X \cup Y)$.

In proving the claim we may assume that $|Z \cap J_F| \geq 2$ as otherwise the claim holds trivially. Now we consider an auxiliary digraph H on the vertex set $Z \cap J_F$ where there is an edge with tail xy and head $x'y'$ if and only if $xy' \in J_F$. Then the desired enumeration is an enumeration of $V(H)$ with no backwards edges and hence its existence is equivalent to the non-existence of a directed cycle in H . So assume for a contradiction there is a directed cycle $(x_1y_1) \cdots (x_{m+1}y_{m+1})$ in H (with $x_1y_1 = x_{m+1}y_{m+1}$). Hence this cycle corresponds to chords $x_iy_{i+1} \in J_F$ of Z . By the minimality of Z the cycles $C(x_iy_{i+1})$ are non-negative for all $i \in [m]$. By applying Lemma 6.7.14 with D_F , Z and the edges x_1y_1, \dots, x_my_m , there exists an integer $q > 0$ such that each edge in $Z - \{x_iy_i \mid i \in [m]\}$ lies in precisely q many cycles $C(x_iy_{i+1})$. Therefore,

$$\sum_{i=1}^n c_F(C(x_iy_{i+1})) = q \cdot c_F(Z),$$

contradicting that Z is negative. □

Let us call a \mathfrak{B} -dijoin F *feasible* if the auxiliary graph D_F together with the cost-function c_F does not contain any negative cycles. Theorem 6.2.21 and Lemma 6.7.5 then imply that for D_F and c_F there always is a feasible integer-valued potential π if F is feasible.

As discussed before, Lemmas 6.7.13 and 6.7.15 immediately yield as a corollary that finite weakly connected digraphs do always contain a feasible \mathfrak{B} -dijoin, since the exchange process strictly decreases the size of the dijoin.

Corollary 6.7.16. *For every finite weakly connected digraph D and every finite-corner-closed class \mathfrak{B} of finite dibonds of D there is a feasible \mathfrak{B} -dijoin. \square*

The question about the existence of a feasible potential in infinite digraphs is a cornerstone of this proof method. While we conjecture that they always exist, we are not able to prove this conjecture in its entirety. We will prove some special cases of this conjecture in Section 6.8.

Conjecture 6.7.17. *For every weakly connected digraph D and every finite-corner-closed class \mathfrak{B} of finite dibonds of D there is a feasible \mathfrak{B} -dijoin.*

This conjecture is weaker than Conjecture 6.1.3, since as the following lemma illustrates each \mathcal{B} -dijoin that features in a \mathfrak{B} -optimal pair is feasible.

Lemma 6.7.18. *Let D be a weakly connected digraph, let \mathfrak{B} be a finite-corner-closed class of finite dibonds of D , and let (F, \mathcal{B}) be a \mathfrak{B} -optimal pair for D . Then F is feasible.*

Proof. Suppose for a contradiction that D_F contains a negative directed cycle C . Let $F^* \cap E(C) = \{f_1^*, f_2^*, \dots, f_n^*\}$ for some $n \in \mathbb{N}$. Also let $f_i \in E(D)$ denote the edge that is mapped to f_i^* by $*$ for every $i \in [n]$. Since (F, \mathcal{B}) is a \mathfrak{B} -optimal pair for D , there exists a unique finite dicut $B_i \in \mathcal{B}$ containing f_i for every $i \in [n]$. Furthermore, B_i contains no further edge of F by definition. Hence, every B_i is an F -tight dicut of D . Let $X_i, Y_i \subseteq V(D)$ be the sides of B_i such that $E_D(X_i, Y_i) = B_i$ holds for every $i \in [n]$. Since B_i is F -tight, we know by Lemma 6.7.7 that B_i is also F -rigid for every $i \in [n]$.

Next let us consider the intersection of C with any cut $E_{D_F}(X_i, Y_i)$ in D_F . Since C is a directed cycle, we know that

$$|E(C) \cap \vec{E}_{D_F}(X_i, Y_i)| = |E(C) \cap \vec{E}_{D_F}(Y_i, X_i)|$$

holds. However, since B_i is F -rigid and contains only the edge f_i from F , there exists an edge $e_i \in E(C) \cap E_{D_F}(X_i, Y_i)$ such that $c(e_i) = 1$ holds for every $i \in [n]$. Furthermore, we know that $e_i \neq e_j$ holds if $i, j \in [n]$ and $i \neq j$, because all elements

of \mathcal{B} are disjoint by definition. Now we have a contradiction to the negativity of C by the following inequality:

$$c_F(C) \geq \sum_{i=1}^n c_F(f_i^*) + \sum_{i=1}^n c_F(e_i) \geq 0. \quad \square$$

The converse of the previous lemma, that every feasible \mathfrak{B} -dijoin features in a \mathfrak{B} -optimal pair is not true. We illustrate this fact with the following example.

Example. In this example we shall consider the weakly connected infinite digraph D depicted twice in Figure 6.7.1. Before we analyse D in detail, let us define D properly. Let $A = \{a_i \mid i \in \mathbb{N}\}$ and $B = \{b_i \mid i \in \mathbb{N}\}$ be two disjoint countably infinite sets. Additionally, let r be some set which is neither contained in A nor in B . Now we set

$$V(D) := A \cup B \cup \{r\}.$$

Next, we define $E_1 := \{a_i b_i \mid i \in \mathbb{N}\}$, $E_2 := \{a_i b_{i+1} \mid i \in \mathbb{N}\}$, $E_3 := \{b_i r \mid i \in \mathbb{N}\}$ and $E_4 := \{b_i b_{i+1} \mid i \in \mathbb{N}\}$. Finally, we complete the definition of D by setting

$$E(D) = E_1 \cup E_2 \cup E_3 \cup E_4.$$

Now let us first consider the left instance of D in Figure 6.7.1 more closely. The set of grey edges in that instance is $F_L := E_2$ and it is easy to check that F_L is a finitary dijoin of D . In particular, F_L is feasible, which follows by checking that the following map $\pi_{F_L} : V(D) \rightarrow \mathbb{Z}$ is a feasible potential:

$$\pi_{F_L}(v) = \begin{cases} -1 & \text{if } v \in A, \\ 0 & \text{otherwise.} \end{cases}$$

Actually, π_{F_L} is even a feasible rooted potential with r as its root. Here we see that the potential threshold τ_{-1} decomposes into a nested set of disjoint dicuts forming a nested optimal pair for D together with F_L for D . Hence, D is not a counterexample to Conjecture 6.1.5.

Next, let us consider the right instance of D in Figure 6.7.1. There the set of grey edges is $F_R := E_1 \cup \{b_0 r\}$. Again it is easy to verify that F_R is a finitary dijoin of D . Furthermore, F_R is feasible, which is witnessed by the following map $\pi_{F_R} : V(D) \rightarrow \mathbb{Z}$ being a feasible potential:

$$\pi_{F_R}(v) = \begin{cases} -2 & \text{if } v \in A, \\ -1 & \text{if } v \in B, \\ 0 & \text{if } v = r. \end{cases}$$




Figure 6.7.1.: Two instances of the digraph D . All edges that are not already directed are meant to be directed from left to right. Each instance of D contains a feasible finitary dijoin F_L and F_R , resp., marked by the grey edges. These dijoints each give rise to a feasible potential π_{F_L} and π_{F_R} , resp.. However, F_R is not part of any optimal pair for D , while F_L is part of some optimal pair for D .

As for π_{F_L} , the map π_{F_R} is even a feasible rooted potential with root r .

Now we shall see that F_R does not feature in any optimal pair for D . In order to see this, consider an edge $a_i b_i$ from F_R for some arbitrary $i \in \mathbb{N}$. The only dibond $a_i b_i$ is contained in is $\delta^+(\{a_i\})$. Hence, in order for F_R to feature in an optimal pair (F_R, \mathcal{B}) for D , we know that $\delta^+(\{a_i\}) \in \mathcal{B}$ must hold for every $i \in \mathbb{N}$. However, every finite dicut containing the edge $b_0 r \in F_R$ must also contain an edge $a_i b_{i+1}$ for some $i \in \mathbb{N}$. This shows that we cannot find an optimal pair for D in which F_R features as a finitary dijoin of D . Note that this is in contrast to Example 6.3.1, where F_R could feature as a finitary dijoin in an optimal pair, just not in a nested one.

We conclude this subsection with a lemma that will help us prove Conjecture 6.7.17 for special classes of digraphs in Section 6.8.

Lemma 6.7.19. *Let D be a weakly connected digraph and let \mathfrak{B} be a finite-corner-closed class of dibonds of D . Then there is a \mathfrak{B} -dijoin F such that for all finite $X \subseteq V(D)$ there is a \mathfrak{B} -dijoin F_X of D such that $F \cap E(D[X]) = F_X \cap E(D[X])$ and such that $D_{F_X}[X]$ does not contain any negative cycles.*

Proof. Without loss of generality we may assume by Lemma 6.4.1 that D is \mathfrak{B} -separable. We prove this lemma by compactness.

For a finite $X \subseteq V(D)$, let \mathcal{J}_X be the set of all $J \subseteq E(D[X])$ such that there is a \mathfrak{B} -dijoin $F_X \subseteq E(D)$ of D with $F_X \cap E(D[X]) = J$ with the property that $D_{F_X}[X]$ does not contain any negative directed cycles. This set is non-empty since by applying Lemmas 6.7.13 and 6.7.15 successively we can eliminate all negative cycles in $D_{F_X}[X]$, as well as finite, since $E(D[X])$ is finite due to X being finite and D being \mathfrak{B} -separable. Moreover, for finite sets $X \subseteq Y \subseteq V(D)$ we obtain that $J \cap E(D[X]) \in \mathcal{J}_X$ for each $J \in \mathcal{J}_Y$ as witnessed by the same \mathfrak{B} -dijoin F_Y of D . By the compactness principle there is a set $F \subseteq E(D)$ such that $F \cap E(D[X]) \in \mathcal{J}_X$ for all finite $X \subseteq V(D)$. We claim that this F is a \mathfrak{B} -dijoin as desired.

First note that for any dibond $B \in \mathfrak{B}$ there is a finite set $X \subseteq V(D)$ with $B \subseteq E(D[X])$. But since the \mathfrak{B} -dijoin F_X witnessing that $F \cap E(D[X]) \in \mathcal{J}_X$ meets B , so does F . Thus F is indeed a \mathfrak{B} -dijoin, finishing the proof. \square

6.7.3. The auxiliary graph for a feasible dijoin

Let D be a weakly connected digraph. Let \mathfrak{B} be a finite-corner-closed class of finite dibonds of D . Let F be a feasible \mathfrak{B} -dijoin of D and let π be a feasible integer-valued potential of the auxiliary graph D_F .

In this subsection we make a couple of remarks regarding the structure of the auxiliary graph D_F .

Lemma 6.7.20. (1) *If there is a jumping edge $uv \in J_F$, then $\pi(v) \leq \pi(u)$.*

(2) *If there is an edge of D from u to v , then $\pi(u) \in \{\pi(v), \pi(v) - 1\}$.*

(3) *If there is an edge in F from u to v , then $\pi(u) = \pi(v) - 1$.*

Proof. (1) is a direct consequence of the feasibility of π and the fact that $c(uv) = 0$ for all jumping edges $uv \in J_F$.

For (2), if there is such an edge of D from u to v we have $\pi(v) - \pi(u) \leq 1$. With (1) it follows that $vu \in J_F$ and hence $\pi(u) \leq \pi(v)$. These two inequalities can only be satisfied for $\pi(u) \in \{\pi(v), \pi(v) - 1\}$.

(3) follows similarly since if there is such an edge in F from u to v we have $\pi(v) - \pi(u) \leq 1$, and with $vu \in F^*$ we get $\pi(u) - \pi(v) \leq -1$. Hence together we get the desired equality. \square

Given any integer-valued potential π on a digraph D we want to introduce the following notation:

$$\Pi_i = \{v \in V(D) \mid \pi(v) = i\} \quad \text{and} \quad \Pi_{>i} = \{v \in V(D) \mid \pi(v) > i\}.$$

The sets $\Pi_{\geq i}$, $\Pi_{<i}$ and $\Pi_{\leq i}$ are analogously defined.

Define $\tau_i := \vec{E}_D(\Pi_{\leq i}, \Pi_{>i})$, which we call the i -th potential threshold of D (with respect to π).

Remark 6.7.21. τ_i is an F -rigid dicut of D for each $i \in \mathbb{Z}$.

Moreover, τ_i and τ_j are disjoint for all $i, j \in \mathbb{Z}$ with $i \neq j$.

6.7.4. The standard decomposition of F -tight dicuts

Let D be a weakly connected digraph. Let \mathfrak{B} be a finite-corner-closed class of finite dibonds of D . Let F be a feasible \mathfrak{B} -dijoin and let π be a feasible integer-valued potential of the auxiliary graph D_F .

By Lemmas 6.7.7 and 6.7.10 and Remark 6.7.21 each finite potential threshold is nestedly F -tight. If Questions 6.7.8 and 6.7.11 are affirmed, this is true for any potential threshold. But this may not be sufficient for the union of their decompositions to be nested as well.

In order to obtain nested dibonds across different potential thresholds in some cases, we consider a specific decomposition of them.

Lemma 6.2.22 yields the following corollary.

Corollary 6.7.22. *Let $B \in \mathfrak{B}^\oplus$ be an F -tight dicut and let \mathcal{B} be a decomposition of B into F -tight dibonds. Then no dicut $A \in \mathcal{B}$ separates a component of $D_F[\text{in}(B)]$ or a component of $D_F[\text{out}(B)]$.*

Proof. Assume for a contradiction that a dicut $A \in \mathcal{B}$ separates a component K of $D_F[\text{in}(B)]$. Then it would contain some components of $D[\text{in}(B)]$ that are contained in K in its out-shore and some in its in-shore. Let K_1 and K_2 be such components with $K_1 \subseteq \text{in}(A)$ and $K_2 \subseteq \text{out}(A)$ and a jumping edge between these components. Such a choice is possible since K is a component of $D_F[\text{in}(B)]$. Since A is F -tight, this jumping edge has its tail in K_1 and its head in K_2 .

By Lemma 6.2.22 there is a dicut $A' \in \mathcal{B}$ with $K_1 \subseteq \text{out}(A')$ and $K_2 \subseteq \text{in}(A')$. But this dicut is not F -rigid, contradicting that it is F -tight by Lemma 6.7.7.

The argument for a component of $D_F[\text{out}(B)]$ is similar. □

Given an F -tight dicut $B \in \mathfrak{B}$, let \mathcal{K}_B be the set of component of $D_F[\text{in}(B)]$. Then for each $K \in \mathcal{K}_B$ the dicut $\delta^-(K)$ is F -rigid. If $\delta^-(K)$ is F -tight as well, then by Corollary 6.7.22 it has a unique decomposition \mathcal{B}_K into F -tight dibonds. Hence if for all $K \in \mathcal{K}_B$ the dicut $\delta^-(K)$ is F -tight, we call the decomposition $\mathcal{B}_B := \cup\{\mathcal{B}_K \mid K \in \mathcal{K}_B\}$ the *standard decomposition of B* . It is easy to see that the standard decompositions of different potential thresholds are nested if they are well-defined.

Lemma 6.7.23. *Given $i, j \in \mathbb{Z}$, the union of the standard decompositions of τ_i and τ_j are nested, if they are well-defined.*

Proof. Without loss of generality let $i < j$. Let $B_i \in \mathcal{B}_i$ and $B_j \in \mathcal{B}_j$. Let $K_i \in \mathcal{K}_{\tau_i}$ be the component such that $B_i \subseteq \delta^-(K_i)$ and $K_j \in \mathcal{K}_{\tau_j}$ be such that $B_j \subseteq \delta^-(K_j)$. Then $B_i = \delta^+(X_i)$ for some component X_i of $D - K_i$ and $B_j = \delta^+(X_j)$ for some component X_j of $D - K_j$. We distinguish two cases.

If $K_j \subseteq K_i$, then either $X_j \supseteq X_i$ or $X_j \cap X_i = \emptyset$, therefore B_i and B_j are nested.

Thus, we consider the case that K_j and K_i are disjoint. Note that since K_i is connected in $D_F[\text{out}(\delta^-(K_j))]$ we get by Corollary 6.7.22 that there is a unique component X'_j of $D - K_j$ that contains K_i . Similarly there is a unique component X'_i of $D - K_i$ that contains K_j . If $X_i = X'_i$ and $X_j = X'_j$, then X_j contains every component of $D - K_i$ except X_i and hence $\text{in}(B_i) \subseteq X_j$. If $X_i = X'_i$ and $X_j \neq X'_j$, then $X_i \subseteq X_j$. If $X_i \neq X'_i$ and $X_j = X'_j$, then $X_j \subseteq X_i$. Finally, if $X_i \neq X'_i$ and $X_j \neq X'_j$, then $X_i \cap X_j = \emptyset$. In any case, B_i and B_j are nested, as desired. \square

For finite F -tight dicuts the standard decomposition is well-defined.

Corollary 6.7.24. *The standard decomposition of a finite F -tight dicut $B \in \mathfrak{B}^\oplus$ is well-defined.*

Proof. If B is finite, then so is $\delta^-(K)$ for any $K \in \mathcal{K}_B$. Since $\delta^-(K)$ is F -rigid, it is also F -tight by Lemma 6.7.7 and the standard decomposition of B is well-defined. \square

Question 6.7.25. *Is the standard decomposition of each F -tight dicut $B \in \mathfrak{B}^\oplus$ well-defined?*

Affirming this question would then also affirm Question 6.7.11.

6.7.5. The auxiliary graph for a healthy feasible dijoin

Let D be a weakly connected digraph. Let \mathfrak{B} be a finite-corner-closed class of finite dibonds of D .

Let us call a feasible integer-valued potential π of the auxiliary graph D_F *healthy* if there is a nested set $\mathcal{B} \subseteq \mathfrak{B}$ of F -tight dibonds which is a disjoint union $\bigcup_{i \in \mathbb{Z}} \mathcal{B}_i$ such that $\tau_i = \bigoplus \mathcal{B}_i$ for all $i \in \mathbb{Z}$. In this case we say π *displays* such a set \mathcal{B} .

Lemma 6.7.26. *If (F, \mathcal{B}) is a nested \mathfrak{B} -optimal pair for D , then there is a healthy feasible integer-valued potential π on D_F displaying \mathcal{B} .*

Proof. Since \mathcal{B} is nested, consider the structure tree $T(\mathcal{B})$ and let $r \in V(T(\mathcal{B}))$ be arbitrary. For each node $t \in V(T(\mathcal{B}))$ consider the unique path P_t from r to t . Now this path contains n_t edges oriented from r to t and m_t edges oriented from t to r . Let $p(t)$ denote $n_t - m_t$. Lastly, for each $v \in t$ we define $\pi(v) = p(t)$. Since each dibond in \mathcal{B} is F -tight, it is easy to check that π is indeed a feasible potential and that for each $B \in \mathcal{B}$ there is a unique $i \in \mathbb{Z}$ such that $B \subseteq \tau_i$. For each $i \in \mathbb{Z}$ the set $\mathcal{B}_i := \{B \in \mathcal{B} \mid B \subseteq \tau_i\}$ defines a decomposition of τ_i into F -tight dibonds in \mathfrak{B} , since each $B \in \mathcal{B}$ is F -tight. Hence, π is healthy and displays \mathcal{B} . \square

By the definition of healthy we have proven the following theorem.

Theorem 6.7.27. *Let D be a weakly connected digraph, let \mathfrak{B} be a finite-corner-closed class of finite dibonds of D and let F be a \mathfrak{B} -dijoin. Then there is a set $\mathcal{B} \subseteq \mathfrak{B}$ such that (F, \mathcal{B}) is a nested \mathfrak{B} -optimal pair for D if and only if there is a healthy feasible integer-valued potential π_F on the auxiliary graph D_F .* \square

From Lemma 6.7.23 with Corollary 6.7.24, and Remark 6.7.21 with Lemma 6.7.7, we immediately obtain the following corollary for finite potential thresholds.

Corollary 6.7.28. *If π is a feasible integer-valued potential such that each potential threshold is finite, then π is healthy.* \square

6.7.6. F -tightly corner-closed classes of finite dibonds

Let D be a weakly connected digraph. Let \mathfrak{B} be a finite-corner-closed class of finite dibonds of D . Let F be a \mathfrak{B} -dijoin of D .

In this subsection we will establish a sufficient condition in terms of F and \mathfrak{B} that ensures that any feasible integer-valued potential will be healthy.

We call a set \mathfrak{C} of dicuts of D *F-tightly corner-closed* if for all $\mathcal{C} \subseteq \mathfrak{C}$ such that each $c \in \mathcal{C}$ is *F-tight* the following statements hold:

1. $\bigwedge \mathcal{C}$ is either empty or nestedly *F-tight*.
2. $\bigvee \mathcal{C}$ is either empty or nestedly *F-tight*.

The following remark is a direct consequence of Remark 6.2.3.

Remark 6.7.29. If \mathfrak{B} is *F-tightly corner-closed*, then so is \mathfrak{B}^\oplus .

This strong condition allows us to affirm Questions 6.7.8 and 6.7.11.

Corollary 6.7.30. *If \mathfrak{B} is F-tightly corner closed, then every F-rigid dicut B is in \mathfrak{B}^\oplus and nestedly F-tight.*

Proof. For each $v \in \text{out}(B)$ and $w \in \text{in}(B)$ let $B_{v,w} \in \mathfrak{B}$ be an *F-tight* dibond with $v \in \text{out}(B_{v,w})$ and $w \in \text{in}(B_{v,w})$. Then $B = \bigwedge_{v \in \Pi_{\leq i}} \bigvee_{w \in \Pi_{> i}} B_{v,w} \in \mathfrak{B}^\oplus$ is nestedly *F-tight* by definition and Remark 6.7.29. \square

Hence, we obtain that the condition that \mathfrak{B} is *F-tightly corner-closed* is sufficient for any feasible integer-valued potential to be healthy, since the dicuts in the standard decomposition of the potential thresholds are *F-rigid*.

Proposition 6.7.31. *Let D be a weakly connected digraph, \mathfrak{B} be a finite-corner-closed class of dibonds of D and F be a \mathfrak{B} -dijoin. Suppose \mathfrak{B} is F-tightly corner-closed. Then any feasible integer-valued potential π of the auxiliary graph D_F is healthy.* \square

To show that some finite-corner-closed class \mathfrak{B} of dibonds of D is *F-tightly corner-closed* we actually only need to show the statement for some sets of *F-tight* dicuts, as the following lemmas will show.

Given a ordinal-indexed sequence $\mathcal{S} = (s_\alpha \mid \alpha < \lambda)$ of length λ and some $\mu < \lambda$ let $\mathcal{S} \upharpoonright \mu$ denote the initial segment $(s_\alpha \mid \alpha < \mu)$ of \mathcal{S} . Given an ordinal λ , we call a sequence $\mathcal{S} = (B_\alpha \mid \alpha < \lambda)$ of *F-tight* dibonds

- \wedge -good if $\bigwedge \mathcal{S} \neq \emptyset$ and $\text{out}(B_\alpha) \cap \text{out}(\bigwedge \mathcal{S} \upharpoonright \alpha)$ is non-empty for each $\alpha \in \lambda$;
- \vee -good if $\bigvee \mathcal{S} \neq \emptyset$ and $\text{in}(B_\alpha) \cap \text{in}(\bigvee \mathcal{S} \upharpoonright \alpha)$ is non-empty for each $\alpha \in \lambda$.

Lemma 6.7.32. *Let $\mathcal{S} \subseteq \mathfrak{B}$ be a sequence of length λ of F-tight dibonds.*

(1) If \mathcal{S} is \wedge -good and if for all $\alpha \leq \lambda$ we have $\wedge \mathcal{S} \upharpoonright \alpha \in \mathfrak{B}$, then $\wedge \mathcal{S}$ is F -tight.

(2) If \mathcal{S} is \vee -good and if for all $\alpha \leq \lambda$ we have $\vee \mathcal{S} \upharpoonright \alpha \in \mathfrak{B}$, then $\vee \mathcal{S}$ is F -tight.

Proof. We show just show (1) since the proof of (2) is analogous.

Let $\mathcal{S} = (B_\alpha \mid \alpha < \lambda)$ be a \wedge -good sequence of length λ . We show by induction on the length λ of the sequence that $\wedge \mathcal{S}$ is F -tight.

Suppose that $\lambda = \beta + 1$. Then $\wedge \mathcal{S} = \wedge \{B_\alpha \mid \alpha < \beta\} \wedge B_\beta$ is an F -tight dibond in \mathfrak{B} by Lemma 6.7.1.

Suppose λ is a limit ordinal. We assume for a contradiction that $\wedge \mathcal{S}$ contains two distinct edges of F . But then there is a $\alpha < \lambda$ such that the dibond $\wedge \mathcal{S} \upharpoonright \alpha \in \mathfrak{B}$ already contains these edges, contradicting that this dibond is F -tight. Hence $\wedge \mathcal{S} \in \mathfrak{B}$ is indeed F -tight. \square

Lemma 6.7.33. *If for each \wedge -good sequence \mathcal{S} we have $\wedge \mathcal{S} \in \mathfrak{B}$ and for each \vee -good sequence \mathcal{S} we have $\vee \mathcal{S} \in \mathfrak{B}$, then \mathfrak{B} is F -tightly corner-closed.*

Proof. Let $\mathcal{B} \subseteq \mathfrak{B}$ be a set of F -tight dicuts in \mathfrak{B} . We only show that $\wedge \mathcal{B}$ is either empty or an F -tight dicut in \mathfrak{B}^\oplus since the other case is analogous. So we assume $\wedge \mathcal{B}$ is non-empty. By Lemma 6.7.32 we may assume without loss of generality that \mathcal{B} contains $\wedge \mathcal{S}$ for any \wedge -good sequence $\mathcal{S} \subseteq \mathcal{B}$.

We claim that $\wedge \mathcal{B} = \bigoplus \{\wedge \mathcal{S} \mid \mathcal{S} \text{ is a maximal } \wedge\text{-good sequence in } \mathcal{B}\}$.

Suppose for a contradiction that for two maximal \wedge -good sequences $\mathcal{S}, \mathcal{S}'$ the dibonds $\wedge \mathcal{S}$ and $\wedge \mathcal{S}'$ are distinct but not disjoint. Let $vw \in \wedge \mathcal{S} \cap \wedge \mathcal{S}'$ and $v'w' \in \wedge \mathcal{S}' \setminus \wedge \mathcal{S}$. Then we can append $\wedge \mathcal{S}$ to \mathcal{S}' to obtain another \wedge -good sequence, contradicting the maximality of $\wedge \mathcal{S}'$.

We need to show that every edge vw in $\wedge \mathcal{B}$ is covered by $\wedge \mathcal{S}$ for some maximal \wedge -good sequence \mathcal{S} . There is a dicut $B_0 \in \mathcal{B}$ with $v \in \text{out}(B_0)$. Starting with the sequence $\mathcal{S}_0 = (B_0)$, as long as the sequence \mathcal{S}_α we constructed so far is not maximal, we can greedily extend it to a sequence $\mathcal{S}_{\alpha+1}$. Note that the sequences for the limit steps we obtain this way are automatically \wedge -good. This way we construct some maximal \wedge -good sequence \mathcal{S} such that $\wedge \mathcal{S}$ will contain vw , since every dibond in \mathcal{B} has w in its in-shore.

Lastly note that $\{\wedge \mathcal{S} \mid \mathcal{S} \text{ is a maximal } \wedge\text{-good sequence in } \mathcal{B}\}$ is a nested set of F -tight dibonds: if for any two maximal \wedge -good sequences $\mathcal{S}, \mathcal{S}'$ the dibonds $\wedge \mathcal{S}$ and $\wedge \mathcal{S}'$ are crossing, then we can append $\wedge \mathcal{S}$ to \mathcal{S}' , or vice versa, contradicting the maximality again. \square

6.8. Applications of Frank's proof

6.8.1. Finite parameters in finite-corner-closed classes

Theorem 6.7.27 together with Corollaries 6.7.16 and 6.7.28 immediately yield the following theorem about finite digraphs.

Theorem 6.8.1. *Let D be a finite weakly connected digraph and \mathfrak{B} be a finite-corner-closed class of dibonds of D . Then there is a nested \mathfrak{B} -optimal pair for D .*

□

Note that it has been observed before that such a version about finite-corner-closed classes in finite digraphs also holds by slight modifications of the usual proofs of Theorem 6.1.5, cf. [20, Theorem 9.7.5].

Lemma 6.5.2 with Theorem 6.8.1 yields the following theorem.

Theorem 6.8.2. *Let D be a weakly connected digraph and \mathfrak{B} be a finite-corner-closed class of finite dibonds of D . Suppose D satisfies one of the following properties:*

- (i) *There exists a \mathfrak{B} -dijoin of D of finite size.*
- (ii) *The maximal number of disjoint dibonds in \mathfrak{B} is finite.*
- (iii) *The maximal number of disjoint and pairwise nested dibonds in \mathfrak{B} .*

Then there is a nested \mathfrak{B} -optimal pair for D .

□

6.8.2. Every edge lies on only finitely many dibonds.

Let D be a weakly connected digraph and let \mathfrak{B} be a finite-corner-closed class of dibonds of D . Suppose every edge of D lies on only finitely many dibonds in \mathfrak{B} .

Lemma 6.8.3. *\mathfrak{B} is F -tightly corner-closed for any \mathfrak{B} -dijoin F .*

Proof. By Lemma 6.7.33 it suffices to show that $\bigwedge \mathcal{S}$ (or $\bigvee \mathcal{S}$, respectively) is an element of \mathfrak{B} for any \wedge -good sequence (or \vee -good sequence, respectively) $\mathcal{S} = (B_\alpha \mid \alpha < \lambda)$ of F -tight dibonds. We only show the first statement, as the other is analogous.

We show this by induction on the length λ of \mathcal{S} . If $\lambda = \beta + 1$, then we are done since $\bigwedge \mathcal{S} \upharpoonright \beta \wedge B_\beta$ is an F -tight dibond in \mathfrak{B} by Lemma 6.7.1. If λ is a limit ordinal, then since each edge $e \in \bigwedge \mathcal{S}$ is eventually contained in all $\bigwedge \mathcal{S} \upharpoonright \alpha$, there is a $\beta < \lambda$ such that for all α with $\beta \leq \alpha < \lambda$ we have $\bigwedge \mathcal{S}_\alpha = \bigwedge \mathcal{S}_\beta$. But then $\bigwedge \mathcal{S} = \bigwedge \mathcal{S} \upharpoonright \beta \in \mathfrak{B}$. \square

Before we prove that D indeed contains some feasible \mathfrak{B} -dijoin, we make the following remark.

Remark 6.8.4. For $v, w \in V(D)$ there are at most finitely many dibonds in \mathfrak{B} separating v and w .

Proof. Let P be an undirected path between v and w . Since every dibond that separates v and w contains at least one of the finitely many edges of P , there are at most finitely many such dibonds. \square

We now can conclude the existence of a feasible \mathfrak{B} -dijoin from Lemma 6.7.19.

Lemma 6.8.5. D contains a feasible \mathfrak{B} -dijoin F .

Proof. Let F be the \mathfrak{B} -dijoin as in Lemma 6.7.19. We claim that F is feasible. Suppose for a contradiction that there is a negative cycle C in D_F . Let $X \subseteq V(D)$ be a finite vertex set containing $V(C)$ as well as every vertex incident to some edge of any dibond $B \in \mathfrak{B}$ that separates any two vertices of C . Then since $F \cap D[X] = F_X$, if $vw \in E(C)$ is a jumping edge in D_F , it is also a jumping edge in $D_{F_X}[X]$ since every dibond of D in \mathfrak{B} that separates v and w is a subset of $E(D[X])$. Hence C is a subdigraph of $D_{F_X}[X]$ as well, a contradiction. \square

With Theorem 6.7.27 and 6.7.31 this yields the proofs of items (iv) of Theorem 6.1.7. More precisely, we get the following theorem.

Theorem 6.8.6. Let D be a weakly connected digraph and \mathfrak{B} be a finite-corner-closed class of finite dibonds of D . Suppose that edge of D lies in only finitely many finite dibonds in \mathfrak{B} . Then there is a nested \mathfrak{B} -optimal pair for D . \square

6.8.3. No infinite dibond in arbitrary corners

Let D be a weakly connected digraph and let \mathfrak{B} be a finite-corner-closed class of dibonds of D . Suppose that for every set $\mathcal{B} \subseteq \mathfrak{B}^\oplus$ the dicuts $\bigwedge \mathcal{B}$ and $\bigvee \mathcal{B}$ are in \mathfrak{B}^\oplus and contain no infinite dibond.

Lemma 6.8.7. \mathfrak{B} is F -tightly corner-closed for any \mathfrak{B} -dijoin F .

Proof. By Lemma 6.7.33 it suffices to show that $\bigwedge \mathcal{S}$ (or $\bigvee \mathcal{S}$, respectively) is an element of \mathfrak{B} for any \wedge -good sequence (or \vee -good sequence, respectively) $\mathcal{S} = (B_\alpha \mid \alpha < \lambda)$ of F -tight dibonds. We only show the first statement, as the other is analogous.

We show this by induction on the length λ of \mathcal{S} . If $\lambda = \beta + 1$, then we are done since $\bigwedge \mathcal{S} \upharpoonright \beta \wedge B_\beta$ is an F -tight dibond in \mathfrak{B} by Lemma 6.7.1.

Suppose λ is a limit ordinal. First note that $D[\text{out}(\bigwedge \mathcal{S})]$ is weakly connected. Take a component C of $D[\text{in}(\bigwedge \mathcal{S})]$. Since $\delta^-(C)$ is finite by assumption there is a $\beta < \lambda$ such that for all α with $\beta \leq \alpha < \lambda$ we have that $\bigwedge \mathcal{S} \upharpoonright \alpha$ is equal to $\delta^-(C) \in \mathfrak{B}$. \square

Again, we still need to deduce the existence of a feasible \mathfrak{B} -dijoin.

Lemma 6.8.8. D contains a feasible \mathfrak{B} -dijoin F .

Proof. Let F be the \mathfrak{B} -dijoin as in Lemma 6.7.19. We claim that F is feasible. Suppose for a contradiction that there is a negative cycle C in D_F . Then since for every finite $X \subseteq V(D)$ the digraph $D_{F_X}[X]$ does not contain C there are vertices $v, w \in V(D)$ such that vw is a jumping edge in $C \subseteq D_F$ that is no jumping edge in D_{F_X} for every $X \subseteq V(D)$. Hence there are infinitely many dibonds in \mathfrak{B} with v in its out-shore and w in its in-shore. Analogously to the construction given in Lemma 6.5.4 we can construct an infinite dibond B with v in its out-shore and w in its in-shore. Since for every pair x, y of vertices with $x \in \text{out}(B)$ and $y \in \text{in}(B)$ there is a dibond $B_{x,y}$ of D with x in its out-shore and y in its in-shore, we obtain by setting $\mathcal{B}_y := \{B_{x,y} \mid x \in \text{out}(B)\}$ for every $y \in \text{in}(B)$ that

$$B = \bigvee \left\{ \bigwedge \mathcal{B}_y \mid y \in \text{in}(B) \right\},$$

a contradiction. \square

Again with Theorem 6.7.27 and 6.7.31 this yields the proof of item (v) of Theorem 6.1.7, or more precisely the following theorem.

Theorem 6.8.9. Let D be a weakly connected digraph and \mathfrak{B} be a finite-corner-closed class of finite dibonds of D . Suppose that for every set $\mathcal{B} \subseteq \mathfrak{B}^\oplus$ the dicuts $\bigwedge \mathcal{B}$ and $\bigvee \mathcal{B}$ are in \mathfrak{B}^\oplus and contain no infinite dibond. Then there is a nested \mathfrak{B} -optimal pair for D . \square

6.8.4. Dibonds of bounded size

Let D be a weakly connected digraph and let \mathfrak{B} be a finite-corner-closed class of dibonds of D . Suppose there is an $m \in \mathbb{N}$ such that $|B| \leq m$ for each $B \in \mathfrak{B}$.

Note that we can reduce this case to the case that every edge lies on only finitely many dibonds of \mathfrak{B} with the help of the following lemma, which is due to Thomassen and Woess. This lemma is a helpful tool in infinite graph theory.

Lemma 6.8.10. [50, Prop. 4.1] *Let G be a connected graph, $e \in E(G)$ and $k \in \mathbb{N}$. Then there are only finitely many bonds of G of size k that contain e .*

Note that this lemma trivially translates to dibonds and hence allows us to immediately obtain the following theorem as a corollary from Theorem 6.8.6.

Theorem 6.8.11. *Let D be a weakly connected digraph and \mathfrak{B} be a finite-corner-closed class of finite dibonds of D . Suppose there is an $m \in \mathbb{N}$ such that $|B| \leq m$ for each $B \in \mathfrak{B}$. Then there is a nested \mathfrak{B} -optimal pair for D . \square*

6.8.5. Dibonds of minimum size

Let D be a weakly connected digraph. Let \mathfrak{B}_{\min} be the class of dibonds of D of minimum size. This class is indeed finite-corner-closed (see the proof of the following theorem), and it has been observed that in finite digraphs there exists a \mathfrak{B}_{\min} -optimal pair for D , cf. [20, Corollary 9.7.6]. Our tools now prove the analogue for infinite digraphs.

Theorem 6.8.12. *For each weakly connected digraph D there is a nested \mathfrak{B}_{\min} -optimal pair for D .*

Proof. By Theorem 6.8.11 it suffices to show that the class \mathfrak{B}_{\min} is finite-corner-closed. Let $B_1, B_2 \in \mathfrak{B}_{\min}$.

If $\text{in}(B_1) \cap \text{in}(B_2) = \emptyset$, then $B_1 \wedge B_2 = \emptyset$ and $B_1 \vee B_2 = B_1 \oplus B_2 \in \mathfrak{B}_{\min}^\oplus$.

If $\text{in}(B_1) \cup \text{in}(B_2) = V(D)$, then $B_1 \wedge B_2 = B_1 \oplus B_2 \in \mathfrak{B}_{\min}^\oplus$ and $B_1 \vee B_2 = \emptyset$.

Otherwise, Remark 6.2.2 yields that both $B_1 \wedge B_2$ and $B_1 \vee B_2$ are of size m , since neither of them can be smaller. Hence they are dibonds of minimum size and hence in \mathfrak{B}_{\min} , proving that \mathfrak{B}_{\min} is finite-corner-closed. \square

6.8.6. Another class of digraphs

Let D be a weakly connected and finitely diseparable digraph. Assume that D does not contain any backwards directed ray, and only contains finitely many sources, finitely many sinks and finitely many ends. Consider the class $\mathfrak{B}_{\text{fin}}$ of finite dicuts.

In this subsection, we will reduce Conjecture 6.1.5 to Conjecture 6.7.17 for this class of digraphs.

Since D is finitely diseparable and hence does not contain any directed cycles, we get the following remark.

- Remark 6.8.13.** (1) For each $v \in V(D)$ there is either a source q of D and a directed q - v -path in D , or there is a backwards directed ray with endvertex v .
- (2) For each $v \in V(D)$ there is either a sink s of D and a directed v - s -path in D , or there is a forwards directed ray with start vertex v . \square

As an immediate consequence we get the following remark.

Remark 6.8.14. Let B be a dicut of D .

- (1) Each component of $D[\text{out}(B)]$ contains a source. Hence there are only finitely many such components.
- (2) Each component of $D[\text{in}(B)]$ contains a sink or a forwards directed ray. \square

Suppose there exists a feasible $\mathfrak{B}_{\text{fin}}$ -dijoin F , and let π denote a feasible integer-valued potential of D_F rooted in some vertex r .

We will show that π is healthy by showing that every potential threshold is finite.

Lemma 6.8.15. *For each $i \in \mathbb{Z}$ the potential threshold τ_i contains at most finitely many edges of F .*

Proof. Suppose for a contradiction that $\tau_i \cap F$ is infinite. Since by Corollary 6.7.9 and Remark 6.7.21 no two weak components have more than one edge of F between them, there are infinitely many weak components of $D[\text{in}(\tau_i)]$ with an edge of F entering them. Since there are only finitely many weak components of $D[\text{out}(\tau_i)]$ by Remark 6.8.14(1), and since there are only finitely many sinks

and ends of D , by applying the pidgeonhole principle twice we get with the help of Remark 6.8.14(2) a weak component C of $D[\text{out}(\tau_i)]$ and two weak components C_0, C_1 of $D[\text{in}(\tau_i)]$ such that there are edges $u_0w_0, u_1w_1 \in F$ with $u_0, u_1 \in C, w_i \in C_i$ as well as forwards directed rays $R_i \subseteq D[C_i]$ for $i \in \{0, 1\}$ which both belong to the same end of D . Let P be an undirected u_0 - u_1 -path in $D[C]$ and let P_i be an undirected path from w_i to the start vertex of R_i in $D[C_i]$ for $i \in \{0, 1\}$. Then applying Lemma 6.7.6 to $U := V(P)$ and $W := V(P_0) \cup V(P_1)$ yields a finite F -tight dicut containing u_0w_1 and u_1w_1 with both R_0 and R_1 in its in-shore. But by Corollary 6.7.9 R_0 and R_1 lie in different weak components of the in-shore, contradicting that both belong to the same end of D . \square

Lemma 6.8.16. *Suppose τ_{i-1} is finite for some $i \in \mathbb{Z}$.*

1. *Any forwards directed rays R_0 in $\Pi_{\leq i}$ and R_1 in $\Pi_{> i}$ belong to different ends of D .*
2. *No forwards directed ray $R_0 \in \Pi_{\leq i}$ is dominated by a vertex d_1 in $\Pi_{> i}$.*
3. *No forwards directed ray $R_1 \in \Pi_{> i}$ is dominated by a vertex d_0 in $\Pi_{\leq i}$.*

Proof. We define the sets $X := \partial(\Pi_{> i-1})$, $Y := \{\text{tail}(f) \mid f \in \tau_i \cap F\}$ and $Z := \{\text{head}(f) \mid f \in \tau_i \cap F\}$. In case (1) or (3), let d_1 denote the start vertex of R_1 , and in case (2), let d_1 denote any vertex in $\Pi_{> i}$.

Let $U := X \cup Y$ if $\pi(r) \neq i$, and let $U := X \cup Y \cup \{r\}$ if $\pi(r) = i$. Furthermore, let $W := Z \cup \{d_1\}$. Applying Lemma 6.7.6 to U and W yields a finite F -tight dicut B with $V(R_1) \subseteq \text{in}(B)$ in case (1) or (3) and $d_1 \in \text{in}(B)$ in case (2). We claim that $\pi^{-1}(i) \subseteq \text{out}(B)$. For each $v \in V(D)$ with $\pi(v) = i$, let P_v be an r - v -path in D_F of cost i . If P_v only contains jumping edges, then $\pi(r) = i$ since B is F -tight, $v \in \text{out}(B)$ as well. Otherwise let e be the last edge of P_v which is not a jumping edge. But then the head of e is in $X \cup Y \subseteq \text{out}(B)$, and since B is F -tight, $v \in \text{out}(B)$ as well. Hence any R_0 - R_1 path (in case (1)), any d_0 - R_1 -path (in case (3)) or R_0 - d_1 -path (in case (2)) has to contain an edge from B . Hence R_0 and R_1 do not belong to the same end, d_0 does not dominate R_1 or d_1 does not dominate R_0 . \square

Lemma 6.8.17. *π is healthy.*

Proof. We show by induction on i that τ_i is finite, which suffices by Lemmas 6.7.7 and 6.7.10.

Since D contains no backwards directed ray and each non-empty out-shore of a potential threshold contains at least one of the finitely many sources, there is a smallest $s \in \mathbb{Z}$ such that τ_s is non-empty.

So inductively assume that τ_{i-1} is finite for some $i \geq s$. Assume for a contradiction that τ_i is infinite. By Lemma 6.8.15 there is a weak component C_1 of $D[\text{in}(\tau_i)]$ for which $\delta^-(C_1)$ is infinite.

If $\partial(C_1)$ is finite, then let d_1 denote a vertex in $\partial(C_1)$ with infinitely many neighbours in $\Pi_{\leq i}$ and let $W := \{d_1\}$.

Claim. *If $\partial(C_1)$ is infinite, then there is either a comb whose spine R_1 is a forwards directed ray in $D[C_1]$ and whose teeth W are an infinite subset of $\partial(C_1)$, or a subdivided star in $D[C_1]$ with centre d_1 whose leaves W are an infinite subset of $\partial(C_1)$.*

Proof of Claim. If $D[C_1]$ contains two forwards directed rays which do not belong to the same end of $D[C_1]$ but to the same end of D , then the infinitely many disjoint undirected paths between them meet $\partial(C_1)$, yielding the desired comb. Otherwise, since each forwards directed ray in $D[C_1]$ belongs to one of the finitely many ends of D , and since $D[C_1]$ contains only finitely many sinks, we can apply the pigeonhole principle to find a sink or an end of D which is reachable from an infinite subset of $\partial(C_1)$. We get the desired comb or star by Lemma 6.2.7 or Lemma 6.2.8. \square

In either case, we look at the neighbourhood of W in $\Pi_{\leq i}$. By Remark 6.8.14(1) there is a weak component C_0 of $D[\Pi_{\leq i}]$ such that $\vec{E}(C_0, W)$ is infinite. Let $U := N(W) \cap C_0$.

If U is finite, then there is a vertex $d_0 \in U$ with $\delta^+(d_0) \cap \vec{E}(C_0, W)$ infinite. But then either d_0 and d_1 cannot be separated from each other by a finite dicut or d_0 dominates R_1 , contradicting Lemma 6.8.16.

If U is infinite, then there is a source q of D in C_0 such that an infinite subset U' of U is reachable from q in $D[C_0]$. By Lemma 6.2.7 there is a forward out-comb in $D[C_0]$ whose spine is a forwards directed ray R_0 and whose infinitely many teeth are in U' or a subdivided out-star with centre d_0 and infinitely leaves in U' . But this contradicts again either that D is finitely separable or Lemma 6.8.16.

Hence by induction τ_i is finite for each $i \in \mathbb{Z}$ and therefore π is healthy by Corollary 6.7.28. \square

Together with Theorem 6.7.27, Lemma 6.8.17 proves the following theorem.

Theorem 6.8.18. *Let D be a weakly connected finitely diseparable digraphs with no backwards directed ray and only finitely many sources, sinks and ends. If D contains some feasible finitary dijoin, then there is a nested optimal pair (F, \mathcal{B}) . \square*

Appendix

English summary

This dissertation deals with different aspects of connectivity and tree structure in infinite graphs, which make it part of the research area of structural infinite graph theory. It consists of two parts: simple graphs are considered in Part I, and directed graphs, or digraphs, are considered in Part II. Part I consists of Chapters 2, 3 and 4, while Part II consists of Chapters 5 and 6.

A typical theme of this research area are duality theorems that offer a dichotomy between the existence of some ‘highly connected part’ in a graph with the non-existence of some tree-decomposition of that graph which, whenever it exists, clearly precludes the existence of such a ‘highly connected part’. In infinite graphs such tree-decompositions may not exist for other reasons. Hence, the notion of a nested separation system, has turned out to be useful for generalisations of the aforementioned type of duality theorem to infinite graphs.

In Chapter 2 we work with *tree sets*, a more abstract version of these nested separation systems, and characterise the tree sets that can be represented as infinite graph-theoretic trees. Moreover, we introduce the notion of *tree-like spaces*, a special form of a graph-like space which exhibits many of the analogous properties that graph-theoretical trees exhibit. We then give a construction how regular tree sets can be represented by tree-like spaces and vice versa.

In Chapter 3 we deal with a different aspect of connectivity by looking at ends of infinite graphs. An interesting property of the rays in a normal spanning tree of a graph is that every ray in an end meets the normal ray corresponding to that end in the normal spanning tree. A ray with this property is called *end-devouring*. Georgakopoulos introduced this concept for finite families of disjoint rays and proved their existence with some feasible set of prescribed start vertices.

We answer a question of Georgakopoulos about the existence of a countable end-devouring family of disjoint rays in ends of countable degree where we can prescribe some feasible set of start vertices in that chapter.

Chapter 4 is dedicated to the investigation of one aspect of higher connectivity, the k -connected sets for some positive integer k . In the spirit of the duality theorems mentioned earlier we prove the equivalence of the existence of a k -connected set of a certain cardinality κ with the non-existence of a tree set in that graph of width less than κ and adhesion less than k , which, whenever it exists, preclude the existence of such a k -connected set.

Moreover, we characterise these k -connected sets via the existence certain minors (or topological minors). For fixed k and cardinal κ there are only finitely many such (topological) minors. Both these results extend similar theorems of Geelen and Joeris in finite graphs.

In Chapter 5 we investigate an extension of Edmonds' Branching Theorem to locally finite digraphs. Similar to how in undirected locally finite graphs the topological framework developed by Diestel and his research group proved to be useful for an extension of tree-packing results (which fail in their straightforward generalisation as shown by Aharoni and Thomassen), we prove some facts about the topological setting in locally finite digraphs and extend a corresponding packing result for pseudo-arborescences, which we introduce here. Furthermore, we study the structure of these pseudo-arborescences as well.

The last chapter, Chapter 6, introduces a conjecture of Heuer which, if affirmed, generalises a min-max theorem of Lucchesi and Younger that the maximum number of disjoint dicuts in a digraph equals the minimum size of a dijoin, i.e. an edge set meeting every dicut. The conjecture states that a structural version of this theorem that restricts itself to finite dicuts in an infinite graph still holds.

We prove a reduction of this conjecture to countable digraphs whose underlying multigraph is 2-connected. Lastly, we affirm the conjecture in a variety of special cases.

Deutsche Zusammenfassung

Diese Dissertation behandelt verschiedene Aspekte von Zusammenhang und Baumstruktur in unendlichen Graphen. Als solche ist sie im Forschungsbereich der strukturellen unendlichen Graphentheorie einzuordnen. Sie besteht aus zwei Teilen, wobei Teil **I** sich mit einfachen Graphen und Teil **II** sich mit gerichteten Graphen beschäftigt. Teil **I** besteht aus den Kapiteln **2**, **3** und **4**, während Teil **II** aus den Kapiteln **5** und **6** besteht.

Eine typische Ausprägung dieses Forschungsgebietes sind Dualitätssätze die eine Dichotomie zwischen der Existenz von bestimmten ‘hoch zusammenhängenden Teilen’ in einem Graphen und der nicht-Existenz einer Baumzerlegung herstellt, die, wenn sie existiert, offenbar die Existenz solcher ‘hoch zusammenhängender Teile’ ausschließt. In unendlichen Graphen können solche Baumzerlegungen allerdings auch aus anderen Gründen nicht existieren. In Folge dessen hat sich der Begriff des geschachtelten Teilungssystems als hilfreich herausgestellt, um solche Dualitätssätze vom Endlichen ins Unendliche zu generalisieren.

In Kapitel **2** arbeiten wir mit *Baummengen*, einer abstrakten Version solcher geschachtelten Teilungssysteme, und Charakterisieren die Baummengen, die als unendliche graphentheoretischen Bäume repräsentiert werden können. Außerdem führen wir den Begriff des *baumartigen Raumes* ein, eine bestimmte Form von graphenartigen Räumen, die viele analoge Aspekte von graphentheoretischen Bäumen ausweisen. Dann geben wir eine Konstruktion an, wie reguläre Baummengen als solche baumartigen Räume dargestellt werden können, und umgekehrt.

In Kapitel **3** betrachten wir ein anderen Aspekt von Zusammenhang, indem wir uns Enden von unendlichen Graphen anschauen. Eine interessante Eigenschaft von Strahlen in normalen Spannbäumen ist, dass jeder Stahl in einem Ende den normalen Stahl im normalen Spannbaum, der zu dem Ende gehört, trifft. Wir sagen, dass so ein Stahl das Ende *verschlingt*. Georgakopoulos hat dieses Konzept auf endliche Familien von disjunkten Strahlen erweitert und die Existenz solcher

Familien mit einer realisierbaren vorgeschriebenen Menge von Startecken gezeigt.

In diesem Kapitel beantworteten wir eine Frage von Georgakopoulos über die Existenz von abzählbaren endenverschlingenden Familien disjunkter Stahlen in Enden mit abzählbarem Grad, wobei wir auch hier eine realisierbare Menge von Startecken vorschreiben können.

Kapitel 4 ist einem Aspekt von höherem Zusammenhang gewidmet, nämlich *k-zusammenhängenden Mengen* für eine positive natürliche Zahl k . Im Sinne der zuvor beschriebenen Dualitätssätze zeigen wir die Äquivalenz der Existenz einer k -zusammenhängenden Menge von Kardinalität κ mit der nicht-Existenz einer Baummenge mit Weite kleiner als κ und Adhäsion kleiner als k , welche, wenn sie existiert, die Existenz einer solchen k -zusammenhängenden Menge ausschließt.

Außerdem charakterisieren wir k -zusammenhängende Mengen via bestimmter Minoren, beziehungsweise topologischer Minoren. Für festes k und Kardinalität κ haben wir nur endlich viele solcher Minoren, die auftreten können. Beide Resultate erweitern Theoreme von Geelen und Joeris über endliche Graphen.

In Kapitel 5 untersuchen wir eine Erweiterung von Edmonds' Arboreszenzpackungssatz in lokal endliche Digraphen. Im ungerichteten Fall hat sich die topologische Herangehensweise, die von Diestel und seiner Arbeitsgruppe entwickelt wurde, als hilfreich herausgestellt um Baumpackungssätze ins lokal endliche zu erweitern (dessen direkte Verallgemeinerung fehlschlägt, wie Aharoni und Thomassen gezeigt haben). Wir zeigen einige Fakten über diese topologische Herangehensweise in lokal endlichen Digraphen und beweisen ein Packungssatz über Pseudoarboreszenzen, die wir hier einführen. Ferner untersuchen wir die Struktur dieser Pseudoarboreszenzen.

Im letzten Kapitel, Kapitel 6 stellen wir eine Vermutung auf, die, falls bestätigt ein Min-Max Theorem von Lucchesi und Younger, welches aussagt, dass die maximale Anzahl gerichteter Schnitte in einem Digraphen gleich der kleinsten Größe einer Menge ist, die alle diese gerichteten Schnitte überdeckt. Wir vermuten, dass eine strukturelle Version dieses Theorems, dass sich auf endliche gerichtete Schnitte beschränkt, auch in unendlichen Digraphen wahr ist.

Wir zeigen, dass wir uns für den Beweis dieser Vermutung auf abzählbare Digraphen, dessen unterliegender Multigraph 2-zusammenhängend ist, beschränken können. Letztlich beweisen wir die Vermutung in diversen Spezialfällen.

Publications related to this dissertation

The following articles are related to this dissertation:

Part I:

1. Chapter 2 is based on [28].
2. Chapter 3 is based on [26].
3. Chapter 4 is based on [25].

Part II:

4. Chapter 5 is based on [24].
5. Chapter 6 is based on [27].

Declaration on my contributions

The research conducted in this thesis is based on collaborative research efforts with two different co-authors. All chapters are based on research articles with either Karl Heuer or Jakob Kneip as co-authors, as mentioned in the previous section of the appendix ‘Publications related to this dissertation’. Both on the research involved in the articles, as well as the creation of the articles, my respective co-authors and I share an equal amount of work. In many cases, the initial drafts of the articles were also developed collaboratively. I will point out, when this differs.

The research of Chapter 2 is a continuation of the research questions I worked on for my bachelor’s thesis. Independently of each other, Jakob Kneip and I both proved the results of Section 2.3. After learning about graph-like spaces, I conjectured the results of Section 2.4 and came up with the rough ideas for the constructions involved in the proof. I then introduced this problem to Jakob Kneip, and together we worked out the details of the proofs. Jakob Kneip did provide a first draft of the proof in Section 2.5.

The result of Chapter 3 answers a question of Georgakopoulos [23], which Karl Heuer introduced to me after doing some preliminary work on the question for locally finite graphs. Together, we came up with a different construction that allowed us to obtain the full result.

The research conducted in Chapter 4 was inspired by a research seminar talk by Reinhard Diestel at the University of Hamburg, in which both Karl Heuer and I participated. The project grew immensely when we were confronted with the difficulties that arose in the case of singular cardinals, whereafter we worked out the proof strategy collaboratively, I drafted a first version of the proof for that case.

Reinhard Diestel encouraged both Karl Heuer and me to extend Edmonds’ Branching Theorem to locally finite digraphs by making use of a similar topological

setting that allowed for the generalisation of tree-packing results in undirected graphs. After completing the common research on the results of Chapter 5, Karl Heuer provided a first draft of the proofs in the article.

The extension of Lucchesi-Younger's Theorem to infinite digraphs is based on a conjecture of Karl Heuer. He found the counterexample of the non-finitary version of this conjecture and found compactness proofs for a few special cases of the finitary version before introducing the problem to me. Collaboratively we developed all the other results of Chapter 6. While Karl Heuer drafted a first version of the proofs in Sections 6.4 and 6.5, I drafted first version of the proofs in Sections 6.7 and 6.8.

Acknowledgement

I would like to thank Reinhard Diestel, my supervisor. You became interested in my mathematical development very early on and guided me through my studies while shaping my research interests. You encouraged and supported me every step along the way, for which I am very grateful to you.

Next, I would like to thank Karl Heuer. We spend most of our time during our master's and doctoral studies working closely together. You are not just my closest collaborator, but a very dear friend as well. I am very grateful for all the fun times we shared together doing mathematics and beyond.

My thanks also for all my colleagues of the discrete mathematics research group in Hamburg. All of you not only provided a great working environment but influenced much of my development as a researcher. A special thanks goes out to my collaborators and to those of you, who became my friends even outside of work.

Thank you to all my friends! Not just for all the fun and games, but all your support as well.

And finally, I would like to thank my family for all their love and support, especially my mother, Ute Gollin. *Danke Mama, für alles!*

Bibliography

- [1] R. Aharoni, *Infinite matching theory*, Discrete Math. **95** (1991), no. 1-3, 5–22, DOI [10.1016/0012-365X\(91\)90327-X](https://doi.org/10.1016/0012-365X(91)90327-X). Directions in infinite graph theory and combinatorics (Cambridge, 1989). MR1141929 ↑[6.1](#), [6.6.1](#), [6.6](#)
- [2] R. Aharoni and E. Berger, *Menger's theorem for infinite graphs*, Invent. Math. **176** (2009), no. 1, 1–62, DOI [10.1007/s00222-008-0157-3](https://doi.org/10.1007/s00222-008-0157-3). MR2485879 ↑[1.1](#), [6.1](#), [6.6](#)
- [3] R. Aharoni and C. Thomassen, *Infinite, highly connected digraphs with no two arc-disjoint spanning trees*, J. Graph Theory **13** (1989), no. 1, 71–74, DOI [10.1002/jgt.3190130110](https://doi.org/10.1002/jgt.3190130110). MR982868 ↑[1.2.4](#), [5.1](#)
- [4] J. Bang-Jensen and G. Gutin, *Digraphs*, 2nd ed., Springer Monographs in Mathematics, Springer-Verlag London, Ltd., London, 2009. Theory, algorithms and applications. MR2472389 ↑[1.3.1](#), [5.1.2](#)
- [5] P. Bellenbaum and R. Diestel, *Two short proofs concerning tree-decompositions*, Combin. Probab. Comput. **11** (2002), no. 6, 541–547, DOI [10.1017/S0963548302005369](https://doi.org/10.1017/S0963548302005369). MR1940119 ↑[4.2](#)
- [6] N. Bowler, J. Carmesin, and R. Christian, *Infinite graphic matroids*, Combinatorica **38** (2018), no. 2, 305–339, DOI [10.1007/s00493-016-3178-3](https://doi.org/10.1007/s00493-016-3178-3). MR3800843 ↑[1.2.1](#), [2.1](#), [2.4.1](#), [2.4.1](#), [2.4.2](#), [2.4.2](#), [2.4.8](#), ||, **
- [7] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark, *k-blocks: a connectivity invariant for graphs*, SIAM J. Discrete Math. **28** (2014), no. 4, 1876–1891, DOI [10.1137/130923646](https://doi.org/10.1137/130923646). MR3270977 ↑[4.2](#), [4.2.7](#)
- [8] W. Dicks and M. J. Dunwoody, *Groups acting on graphs*, Cambridge Studies in Advanced Mathematics, vol. 17, Cambridge University Press, Cambridge, 1989. MR1001965 ↑[6.2.5](#)
- [9] R. Diestel, *Graph theory*, 5th ed., Graduate Texts in Mathematics, vol. 173, Springer, Berlin, 2018. MR3822066 ↑[1.3.1](#), [1.3.2](#), [3.1](#), [4.1](#), †, [4.2.1](#), [4.2](#), [4.2](#), [4.2](#), [5.1.1](#), [5.1](#), [5.1.4](#), [5.2.1](#), [5.2.2](#), [5.2.3](#), [5.2.4](#), [5.2.5](#), [6.2.4](#)
- [10] ———, *Abstract separation systems*, Order **35** (2018), no. 1, 157–170, DOI [10.1007/s11083-017-9424-5](https://doi.org/10.1007/s11083-017-9424-5). MR3774512 ↑[2.1](#), [2.2.2](#)
- [11] ———, *Locally finite graphs with ends: a topological approach* (2012), available at [arXiv:0912.4213v3](https://arxiv.org/abs/0912.4213v3). Post-publication manuscript. ↑[3.1](#), †, [4.2](#), [5.1](#)

- [12] ———, *Tree sets*, Order **35** (2018), no. 1, 171–192, DOI [10.1007/s11083-017-9425-4](https://doi.org/10.1007/s11083-017-9425-4). MR3774513 ↑[2.1](#), [2.3.1](#), [4.1](#), [4.2](#), [4.9.1](#)
- [13] R. Diestel, T. R. Jensen, K. Yu. Gorbunov, and C. Thomassen, *Highly connected sets and the excluded grid theorem*, J. Combin. Theory Ser. B **75** (1999), no. 1, 61–73, DOI [10.1006/jctb.1998.1862](https://doi.org/10.1006/jctb.1998.1862). MR1666948 ↑[1.2.3](#), [4.1](#)
- [14] R. Diestel and J. Kneip, *Profinite separation systems* (2018), available at [arXiv:1804.01921](https://arxiv.org/abs/1804.01921). ↑[4.9.1](#)
- [15] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, J. Combin. Theory Ser. B **87** (2003), no. 1, 197–206, DOI [10.1016/S0095-8956\(02\)00034-5](https://doi.org/10.1016/S0095-8956(02)00034-5). Dedicated to Crispin St. J. A. Nash-Williams. MR1967888 ↑[5.2.1](#)
- [16] R. Diestel and S.-i. Oum, *Tangle-tree duality in abstract separation systems* (2017), available at [arXiv:1701.02509](https://arxiv.org/abs/1701.02509). Preprint. ↑[2.1](#)
- [17] ———, *Tangle-tree duality: in graphs, matroids and beyond* (2017), available at [arXiv:1701.02651](https://arxiv.org/abs/1701.02651). Preprint. ↑[2.1](#)
- [18] R. Diestel and G. Whittle, *Tangles and the Mona Lisa* (2016), available at [arXiv:1603.06652](https://arxiv.org/abs/1603.06652). Preprint. ↑[2.1](#)
- [19] J. Edmonds, *Edge-disjoint branchings*, Combinatorial algorithms (Courant Comput. Sci. Sympos. 9, New York Univ., New York, 1972), Algorithmics Press, New York, 1973, pp. 91–96. MR0351889 ↑[1.2.4](#), [5.1](#), [5.1.2](#)
- [20] A. Frank, *Connections in combinatorial optimization*, Oxford Lecture Series in Mathematics and its Applications, vol. 38, Oxford University Press, Oxford, 2011. MR2848535 ↑[6.1](#), [6.1](#), [6.1](#), [6.7](#), [6.7.2](#), [6.8.1](#), [6.8.5](#)
- [21] S. Fujishige, *Algorithms for solving the independent-flow problems*, J. Oper. Res. Soc. Japan **21** (1978), no. 2, 189–204, DOI [10.15807/jorsj.21.189](https://doi.org/10.15807/jorsj.21.189) (English, with Japanese summary). MR0484341 ↑[6.7](#)
- [22] J. Geelen and B. Joeris, *A generalization of the Grid Theorem* (2016), available at [arXiv:1609.09098](https://arxiv.org/abs/1609.09098). ↑[1.2.3](#), [4.1](#), [4.1](#)
- [23] A. Georgakopoulos, *Infinite Hamilton cycles in squares of locally finite graphs*, Adv. Math. **220** (2009), no. 3, 670–705, DOI [10.1016/j.aim.2008.09.014](https://doi.org/10.1016/j.aim.2008.09.014). MR2483226 ↑[1.2.2](#), [3.1](#), [6.8.6](#)
- [24] J. P. Gollin and K. Heuer, *An analogue of Edmonds’ Branching Theorem for infinite digraphs* (2018), available at [arXiv:1805.02933](https://arxiv.org/abs/1805.02933). Preprint. ↑[4](#)
- [25] ———, *Characterising k -connected sets in infinite graphs* (2018), available at [arXiv:1811.06411](https://arxiv.org/abs/1811.06411). Preprint. ↑[3](#)
- [26] ———, *Infinite end-devouring sets of rays with prescribed start vertices*, Discrete Math. **341** (2018), no. 7, 2117–2120, DOI [10.1016/j.disc.2018.04.012](https://doi.org/10.1016/j.disc.2018.04.012). MR3802167 ↑[2](#)

- [27] ———, *On the Infinite Lucchesi-Younger Conjecture* (2019). available at <https://www.math.uni-hamburg.de/home/gollin/preprints/ly-infinite.pdf>. Preprint. ↑5
- [28] J. P. Gollin and J. Kneip, *Representations of infinite tree sets* (2019). available at <https://www.math.uni-hamburg.de/home/gollin/preprints/tree-sets-infinite.pdf>. Preprint. ↑4.1, 1
- [29] R. Halin, *Simplicial decompositions of infinite graphs*, Ann. Discrete Math. **3** (1978), 93–109. Advances in graph theory (Cambridge Combinatorial Conf., Trinity Coll., Cambridge, 1977). MR499113 ↑1.2.3, 4.1, 4.6.1
- [30] ———, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Math. Nachr. **30** (1965), 63–85, DOI 10.1002/mana.19650300106 (German). MR0190031 ↑1.3.1, 3.2
- [31] B. Joeris, *Connectivity, tree-decompositions and unavoidable-minors*, Ph.D. Thesis, University of Waterloo, 2015. available at <https://uwspace.uwaterloo.ca/handle/10012/9315>. ↑1.2.3, 4.1, 4.1
- [32] A. Joó, *Edmonds' branching theorem in digraphs without forward-infinite paths*, J. Graph Theory **83** (2016), no. 3, 303–311, DOI 10.1002/jgt.22001. MR3549511 ↑5.1
- [33] K. Kunen, *Set theory*, Studies in Logic and the Foundations of Mathematics, vol. 102, North-Holland Publishing Co., Amsterdam-New York, 1980. An introduction to independence proofs. MR597342 ↑1.3.1, 4.2.6
- [34] L. Lovász, *On two minimax theorems in graph*, J. Combinatorial Theory Ser. B **21** (1976), no. 2, 96–103, DOI 10.1016/0095-8956(76)90049-6. MR0427138 ↑6.1, 6.1
- [35] C. L. Lucchesi and D. H. Younger, *A minimax theorem for directed graphs*, J. London Math. Soc. (2) **17** (1978), no. 3, 369–374, DOI 10.1112/jlms/s2-17.3.369. MR500618 ↑1.2.5, 6.1, 6.1.1, 6.1, 6.1, 6.1.4
- [36] C. St. J. A. Nash-Williams, *Edge-disjoint spanning trees of finite graphs*, J. London Math. Soc. **36** (1961), 445–450, DOI 10.1112/jlms/s1-36.1.445. MR0133253 ↑1.2.4, 5.1, 5.1.1
- [37] B. Oporowski, J. Oxley, and R. Thomas, *Typical subgraphs of 3- and 4-connected graphs*, J. Combin. Theory Ser. B **57** (1993), no. 2, 239–257, DOI 10.1006/jctb.1993.1019. MR1207490 ↑1.2.3, 4.1, 4.6.2
- [38] J. G. Oxley, *On a packing problem for infinite graphs and independence spaces*, J. Combin. Theory Ser. B **26** (1979), no. 2, 123–130, DOI 10.1016/0095-8956(79)90050-9. MR532581 ↑5.1
- [39] R. Rado, *Axiomatic treatment of rank in infinite sets*, Canadian J. Math. **1** (1949), 337–343. MR0031993 ↑1.3.2
- [40] ———, *A selection lemma*, J. Combinatorial Theory Ser. A **10** (1971), 176–177. MR0270920 ↑1.3.2

- [41] N. Robertson and P. D. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Combin. Theory Ser. B **52** (1991), no. 2, 153–190, DOI [10.1016/0095-8956\(91\)90061-N](https://doi.org/10.1016/0095-8956(91)90061-N). MR1110468 ↑[1.1](#), [2.1](#)
- [42] N. Robertson, P. Seymour, and R. Thomas, *Excluding infinite minors*, Discrete Math. **95** (1991), no. 1-3, 303–319, DOI [10.1016/0012-365X\(91\)90343-Z](https://doi.org/10.1016/0012-365X(91)90343-Z). Directions in infinite graph theory and combinatorics (Cambridge, 1989). MR1141945 ↑[2.1](#), [4.1](#)
- [43] ———, *Quickly excluding a planar graph*, J. Combin. Theory Ser. B **62** (1994), no. 2, 323–348, DOI [10.1006/jctb.1994.1073](https://doi.org/10.1006/jctb.1994.1073). MR1305057 ↑[1.2.3](#), [4.1](#)
- [44] P. D. Seymour and R. Thomas, *Graph searching and a min-max theorem for tree-width*, J. Combin. Theory Ser. B **58** (1993), no. 1, 22–33, DOI [10.1006/jctb.1993.1027](https://doi.org/10.1006/jctb.1993.1027). MR1214888 ↑[2.1](#)
- [45] M. Stein, *Extremal infinite graph theory*, Discrete Math. **311** (2011), no. 15, 1472–1496, DOI [10.1016/j.disc.2010.12.018](https://doi.org/10.1016/j.disc.2010.12.018). MR2800973 ↑[4.5.1](#)
- [46] R. Thomas, *A Menger-like property of tree-width: the finite case*, J. Combin. Theory Ser. B **48** (1990), no. 1, 67–76, DOI [10.1016/0095-8956\(90\)90130-R](https://doi.org/10.1016/0095-8956(90)90130-R). MR1047553 ↑[4.2](#)
- [47] C. Thomassen, *Hamiltonian paths in squares of infinite locally finite blocks*, Ann. Discrete Math. **3** (1978), 269–277. Advances in graph theory (Cambridge Combinatorial Conf., Trinity College, Cambridge, 1977). MR499125 ↑[3.1](#)
- [48] ———, *Edmond’s branching theorem in digraphs without backward-infinite paths*. unpublished manuscript, personal communication. ↑[5.1](#)
- [49] C. Thomassen and A. Vella, *Graph-like continua, augmenting arcs, and Menger’s theorem*, Combinatorica **28** (2008), no. 5, 595–623, DOI [10.1007/s00493-008-2342-9](https://doi.org/10.1007/s00493-008-2342-9). MR2501250 ↑[1.2.1](#), [2.1](#), [2.4.1](#)
- [50] C. Thomassen and W. Woess, *Vertex-transitive graphs and accessibility*, J. Combin. Theory Ser. B **58** (1993), no. 2, 248–268, DOI [10.1006/jctb.1993.1042](https://doi.org/10.1006/jctb.1993.1042). MR1223698 ↑[6.8.10](#)
- [51] W. T. Tutte, *On the problem of decomposing a graph into n connected factors*, J. London Math. Soc. **36** (1961), 221–230, DOI [10.1112/jlms/s1-36.1.221](https://doi.org/10.1112/jlms/s1-36.1.221). MR0140438 ↑[1.2.4](#), [5.1](#), [5.1.1](#), [5.1.3](#)
- [52] U. Zimmermann, *Minimization of some nonlinear functions over polymatroidal network flows*, Bonn Workshop on Combinatorial Optimization (Bonn, 1980), Ann. Discrete Math., vol. 16, North-Holland, Amsterdam-New York, 1982, pp. 287–309. MR686314 ↑[6.7](#)

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.