Path spaces

Dissertation

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

vorgelegt
im Fachbereich Mathematik
von
Hendrik Heine

Hamburg
2022

Erstgutachter und Betreuer: Prof. Nathan Bowler
Zweitgutachter: Dr. Max Pitz
Datum der Disputation: 25. Oktober. 2022

Contents

0.1 Introduction

In graphs paths determine connectivity. Yet even in infinite graphs, all paths are finite. This causes significant trouble when trying to generalize theorems from finite to infinite graphs, sometimes explicitly as for Hamilton cycles and Euler tours and sometimes subtly as for the tree packing theorem. To remedy this, [16] defined a topological space on a locally finite graph and its ends, the so called Freudenthal compactification, and worked with topological arcs in this space instead of paths. This method proved quite successful for problems like the ones mentioned above, as both their original article and successive research demonstrate. For instance, while [16] already tackled the Euler tour problem, [37] gives a tree-packing theorem and [8] some results on Hamilton cycles.

Since the topological spaces obtained from graphs have some common properties, this also sparked some interest in general classes including these, like those considered in [49] and later [23]. Topological spaces with some graphic properties have been studied by topologists for much longer, see for example [9], but they can also be of help for combinatorial aims.

An example of this important for this thesis is [5]. There the goal is to represent infinite graphic matroids via graph-like spaces. However, cycles of infinite matroids can have arbitrary cardinality, so not all of them can be represented via images of S^1. Thus they introduce so called pseudo-lines whose edges can be arranged like any linear order, like the long line whose edges are arranged like ω_1. Similar long arcs have been considered by topologists as linearly ordered topological spaces, see for example [30].

Despite some differences in these path-like topological notions, proofs of combinatorial statements in them are often parallel. This raises the question whether it is possible to find some more combinatorial structure in which these proofs can be done for different notions simultaneously. Such a structure, which we call path spaces, is the topic of this dissertation. The main idea is abstracting from topological arcs to their induced linear orders. These sorts of sets of linear orders can then be axiomatized in such a way that many definitions and results from graph theory still make sense here. The role that compactness plays in proofs about topological paths is replaced by an axiom which makes sure that the intersection of two paths is closed in the order topology. One further advantage of this approach is that some definitions, like that of torsos, become much easier when we do not need to worry about specifying a topology.

After some short preliminaries in Section 0.2, we give our definition of path spaces and list some examples in Section 0.3. In Section 0.4 we then develop a basic toolkit of definitions for working with path spaces and in Section 0.4.1 we focus in on minors. Our goal in the remainder of the dissertation is then to prove results about them, in particular to translate many graph theorems to path spaces. This will happen over the course of four chapters, which can be read mostly independently . We will only sketch their contents here, since they each begin with their own more substantial introduction. Each chapter also concludes with an outlook on further research.

Chapter 1 deals with general connectivity theory, in particular we prove two duality theorems. The first is a version of Menger's theorem, the second is about the number of paths starting and ending in the same set.

Chapter 2 is about a version of tree decompositions for path spaces; we give decompositions into blocks as well as along separators of order two and prove a grid theorem.

Chapter 3 is about the structure of separations of path spaces. We show how to distinguish profiles in path spaces and investigate its flowers.

Chapter 4 considers ubiquity. We show that minor-like embeddings of finite, planar graphs are ubiquitous.

0.2 Preliminaries

Let us first note that when talking about graphs we mostly use the notation of [14] except that we write $R_{m,n}$ for the $m \times n$ grid. Now we will delve into some more particular topics. These are not necessary for every chapter and can be consulted when needed.

0.2.1 Separation systems

We will start by introducing separation systems; the relevant definitions are taken from [13]. Recall that a separation of a graph is a pair (A, B) with $A \cup B = V(G)$ such that there is no edge between A and B. If we leave out the condition about the edges and work on an arbitrary set, this defines the more general vertex separations. Either way, these separations have a natural order given by $(A, B) \leq (C, D)$ if and only if $A \subseteq C$ and $D \subseteq B$. Note that if $(A, B) \leq (C, D)$ then also $(D, C) \leq (B, A)$. These characteristics suffice for general framework for identifying cohesive substructures.

A *separation system* is a triple $(S, \leq, ^*)$, where $(S, \leq)$ is a partial order and * is an involution on S such that $s \leq t$ if and only if $t^* \leq s^*$ for all $s, t \in S$. We call the elements of S *(oriented) separations* and sets of the form $\{s, s^*\}$ for $s \in S$ unoriented separations. For $s \in S$ we call s^* the *inverse* of s.

A separation s is called *small* if $s \leq s^*$, *trivial* if there exists a separation t such that $s < t$ and $s < t^*$ and *degenerate* if $s = s^*$. Note that every trivial separation is small. In a separation system of separations of a graph G, the small separations are those of the form $(A, V(G))$ for some $A \subseteq V(G)$, but they need not necessarily be trivial depending on the surrounding separation system. $(V(G), V(G))$ is the only degenerate separation of a graph G.

A pair of unoriented separations is *nested* if they contain comparable elements and a pair of oriented separations (s, t) is *nested* if the unoriented separation containing s is nested with the one containing t. If two (unoriented) separations are not nested, we say that they *cross*. A separation in some separation system is *good* if it is nested with all other separations of that system and neither it nor its inverse is small. Suprema of chains of good separations are good unless they are small or cosmall.

Lemma 0.2.1. *Let C be a chain of good separations in a separation system S. Then the supremum of C is nested with every separation of S.*

Proof. Let s be the supremum of C and let t be a separation which is not small or cosmall. Without loss of generality t is comparable with every separation in C. If $t \leq c$ for some $c \in C$, then clearly also $t \leq s$ and we are done. So we may assume that t is an upper bound for C and thus $s \leq t$ since s is the supremum. $\square$

A *partial orientation* of a separation system S is a subset of S meeting each unoriented separation at most once. It is an *orientation* of S if it meets each unoriented separation exactly once. A partial orientation is *consistent* if it contains no separations $s \neq t$ such that $s^* < t$. Given a separation system S we write $O(S)$ for the consistent orientations of S. In graphs consistent orientations are induced by many cohesive substructures, for example clique minors or more generally highly connected sets.

We call a set of separations $T \subseteq S$ *regular* if it contains no small separations and *nested* if all pairs of separations contained in it are nested. A *tree set* is a nested separation system containing no trivial or degenerate elements.

The following special case of [13, Lemma 4.1] will be useful.

Lemma 0.2.2. *Let S be a regular tree set and P a consistent partial orientation of S. Then P extends to a consistent orientation and for any maximal element of P there is a unique consistent orientation extending P in which that element is maximal.*

Tree sets are the main structure we are looking for in abstract separation systems, since they correspond to tree decompositions in graphs, as we will detail later. To actually construct them, however, we need an additional tool. If there are two crossing separations which we would like in our tree set, to find replacements for one or the other a natural choice would be a supremum or infimum, if it exists. In graphs (or more generally for vertex separations) it is easy to define a supremum of two separations (A, B) and (C, D) as $(A \cup C, B \cap D)$ and an infimum analogously. To ensure that such a separation is actually in our system, in graphs we usually work with the set of all separations whose order, defined as the size of separator, is at most some natural number k. This motivates the following definitions.

A separation system S together with binary functions $\vee$ and $\wedge$ on S is called a *universe* if $s \vee t$ is the supremum of $s, t \in S$ and $s \wedge t$ is their infimum. A *submodular universe* consist of a universe U together with a function from U into $\mathbb{N} \cup \{\infty\}$. (denoted as $|u|$ for $u \in U$) such that $|r| = |r^*|$ and $|r \vee s| + |r \wedge s| \leq |r| + |s|$ for all $r, s \in U$.[1] We call $|s|$ the *order* of s and if s has order k, we call it a *k-separation*. In a submodular universe the separation system of separations of order at most k contains the supremum or the infimum of any pair of its separations, solving the problem discussed before. Given two separations v and w in a universe, we call the (up to) eight separations obtained by applying $\vee$ and $\wedge$ to some orientation of v and some orientation of w their *corners*.

0.2.2 Graph-like spaces

Next we proceed to the graph-like spaces of [5]. They define a *graph-like space* as a topological space G together with a vertex set V, an edge set E and a continuous map $t_e^G : [0, 1] \mapsto G$ for every $e \in E$ satisfying the following conditions:

1. V and $(0, 1) \times E$ are disjoint.

2. The underlying set of G is $V \cup (0, 1) \times E$.

3. $\{t_e^G(0), t_e^G(1)\} \subseteq V$ for every $e \in E$.

[1] The source allows real values here, but we do not

4. $\{t_e^G(x); x \in (0,1)\} \subseteq (0,1) \times \{e\}$ for every $e \in E$

5. $t_e(G)$ restricted to $(0,1)$ is an open map for every $e \in E$.

6. For $v, w \in V$ there exist disjoint open subsets U, U' of G such that $v \in U$, $w \in U'$ and $V(G) \subseteq U \cup U'$.

We write $V(G)$ for the vertex set of G and $E(G)$ for the edge set. A *graph-like subspace* H of a graph-like space G is a graph-like space such that its topological space is a subspace of that of G, its vertex and edge sets are subsets of those of G and the map for any edge of H is the same as the map for the same edge in G. For an edge e of G we call $(0,1) \times e$ the set of *inner points* of e and write $G - e$ for the subspace obtained by deleting these points. We call a compact, topologically connected graph-like subspace A of G a *pseudo-line (with endvertices $x, y \in V(A)$)* if for every edge e the vertices x and y are contained in different topological components of $A - e$ and for every $v, w \in V(A)$ there exists some edge f such that v and w are in different topological components of $G - f$. Any nontrivial pseudo-line is the closure of the set of inner points of its edges ([5, Corollary 4.7]).

A *tree-like space*, as defined in [28], is a compact graph-like space, such that between any two vertices there is a unique pseudo-line. We will write $L_T(s,t)$ for the unique pseudo-line between s and t in a tree-like space T. Call $t \in V(T)$ a *limit point* of T if there is a nontrivial pseudo-line containing t in which t is not an endpoint of an edge.

0.3 Definition and examples

We need to start with some preliminary notions to set up the definition of path spaces. Call a subset Y of a linearly ordered set X *complete*, if for any nonempty $Z \subseteq Y$ there exists a supremum and infimum in X and these are contained in Y. We define a *path* as a nonempty linearly ordered set complete in itself. For a path P with $x \leq y$ in P we call the closed interval from x to y with the induced order its *segment* from x to y. We will call a path *nontrivial* if it has more than one element. A set of paths $\mathcal{P}$ is called *compatible* if for any $P, Q \in \mathcal{P}$ the set $P \cap Q$ is complete in P and *strongly compatible* if any two segments of paths in $\mathcal{P}$ are compatible. Given two paths P and Q we say that P *connects to* Q, if $P \cap Q = \{x\}$ where x is the maximum of P and the minimum of Q. In this case, we call the union of P and Q with the induced order their *concatenation*. The *inverse* of a path is obtained by reversing its order.

Now we have enough to define path spaces (and their directed counterparts). A set of paths is called a *dipath space* if is compatible and closed under segments and concatenations. A dipath space is a *path space* if it is also closed under inverses.

For defining some of the upcoming examples, we require one additional tool. Let $\mathcal{P}$ be a set of paths. We can close it under segments by just adding all segments of its paths and close it under concatenations by inductively adding in countably many steps all concatenations of paths constructed so far. Now we define $\widehat{\mathcal{P}}$ by closing under segments and then closing under concatenations.

Lemma 0.3.1. *Let $\mathcal{P}$ be a strongly compatible set of paths. Then $\widehat{\mathcal{P}}$ is a dipath space.*

Proof. Let us first show that $\widehat{\mathcal{P}}$ is closed under segments. Let Q be a segment of some $P \in \widehat{\mathcal{P}}$. Now P is the concatenation of finitely many $P_1, \dots, P_n$ which are segments of paths in $\mathcal{P}$. We may assume w.l.o.g. that Q meets all of these in a nontrivial segment, otherwise we may move to a shorter concatenation. Let Q_1 be the segment of P_1 which Q meets and define Q_n similary. Then Q is the concatenation of $Q_1, P_2, \dots, P_{n-1}, Q_n$, which are all segments of paths in $\mathcal{P}$, and thus Q is contained in $\hat{\mathcal{P}}$.

Now it remains to show that $\widehat{\mathcal{P}}$ is compatible. Since $\mathcal{P}$ is strongly compatible, closing it under segments keeps it compatible. Thus it suffices to prove that closing a compatible set under concatenations leaves it compatible. For this let P_1, P_2, Q be three paths from a compatible set and let P be the concatenation of P_1 and P_2. Let $Z \subseteq P \cap Q$ be nonempty. Clearly there is a supremum of Z in Q as required, namely the maximum of the suprema of $Z \cap P_1$ and $Z \cap P_2$ in Q. Furthermore if Z does not meet P_2, then the supremum of Z in P_1 is the supremum of Z in P and otherwise the supremum of Z in P_2 is the supremum of Z in P. For infima the proof proceeds analogously. □

This justifies calling $\widehat{\mathcal{P}}$ the *directed completion* of $\mathcal{P}$. If we close $\mathcal{P}$ under inverses before taking the directed completion then we obtain a path space $\overline{\mathcal{P}}$, which we call the *completion* of $\mathcal{P}$.

The rest of this section will provide various examples of (di)path spaces. Since the proofs required are all straightforward and very similar, we omit most of them.

The most basic examples of path spaces are of course given by finite and infinite graphs, as mentioned in the introduction.

Example 0.3.2. Given a graph G we obtain a path space as the completion of the set of single point orders given by each vertex of G and two point orders on the vertices of each edge of G. Equivalently, we can also define this path space as containing for each path P of G the linear order of $V(P)$ which is given by the sequence of the vertices. Of course, each path of this path space is finite.

Conversely given any path space $\mathcal{P}$ with only finite paths we can define a graph inducing it by taking the vertex set as $V(\mathcal{P})$ and adding an edge vw whenever $\mathcal{P}$ contains a two element path from v to w. In particular all finite path spaces come from finite graphs.

Directed graphs induce dipath spaces in an analogous manner using the directed completion and all dipath spaces with all paths finite arise from directed graphs.

A more general class of examples is given by topological spaces.

Example 0.3.3. Let X be a Hausdorff space and $\Phi : [0, 1] \to X$ an arc in X. Then Φ induces a path P_Φ on its image given by the standard ordering on $[0, 1]$. Let $A(X)$ be the set of all these P_Φ for all arcs Φ in X (together with the trivial path for each point). Let us now prove that $A(X)$ is a path space. Clearly $A(X)$ is closed under inverses.

Let $P_\Phi \in A(X)$ be arbitrary and let a, b be two points in the image of Φ, w.l.o.g. different. Then by scaling the interval between their preimages, we can obtain a new arc Ψ between a and b such that P_Ψ is a segment of P_Φ.

Let P_Φ and P_Ψ in $A(X)$ be such that the maximal element of P_Φ is the minimal element of P_Ψ, but otherwise disjoint. By scaling and combining Φ and Ψ we then obtain a new arc Λ, such that P_Λ is the concatenation of P_Φ and P_Ψ.

Now let P_Φ and P_Ψ in $A(X)$ be arbitrary. Since the images of Φ and Ψ are compact in X and X is Hausdorff, they are also closed. Then $P_\Phi \cap P_\Psi$ is again closed and also inherits compactness. Thus this set is complete, since these paths inherit their order from $[0, 1]$.

We could also consider the set of injective continous maps from any compact, connected ordered topological space (see [12]) to X and obtain a path space the same way. This still holds true if we choose an infinite cardinal κ and consider only maps from ordered spaces of size less than κ.

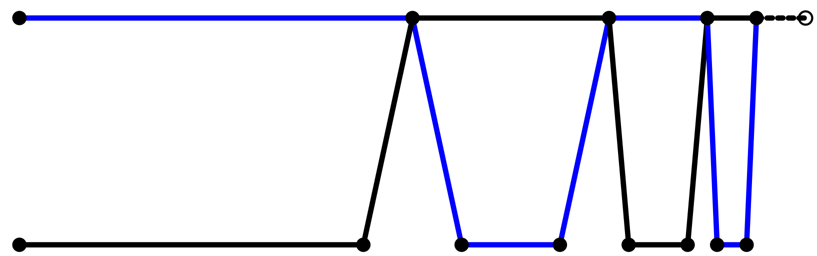

Figure 1: An illustration of Example 0.3.6. The hollow circle signifies that this point only forms a path with the top and bottom ray, not for example with the one indicated in blue.

Since we just saw that long arcs in topological spaces form a path space, it is not all that surprising that pseudo-lines in graph-like spaces also do.

Example 0.3.4. Let G be a graph-like space and L a pseudo-line in G. Then L together with an endvertex v of L induces a path Q_L^v on the vertices of L given by the order of L with minimal element v. Let $P(G)$ be the set of all these Q_L^v. This is a path space.

Furthermore given some set o of orientations of edges of G, we can define $P^o(G)$ as the set of all Q_L^v such that the orientation towards v of every edge of L is contained in o. This is a dipath space. As in the prior example we can again choose an infinite cardinal κ and consider only those pseudo-lines of size at most κ.

If a property is defined piecewise, as in the following example, we can impose it on our paths and retain a path space.

Example 0.3.5. Let $k \in \{1, 2, \dots, \infty\}$. If M is a differentiable manifold and f an injective, piecewise k-times (continously) differentiable curve on M, let Q_f be the image of f with the order induced by $[0, 1]$. Then the set of all these Q_f forms a path space.

All the classes of path spaces considered so far are in some sense topological. We will now look at a path space which does not belong to any of them and is slightly pathological.

Example 0.3.6. Let $X = [0, 1] \times \{0, 1\}$. Let $\sim$ be the relation on X with $(x, a) \sim (y, b)$ if we have $x = y$ and at least one of $a = b$, $x = 1$ or $x = 1 - \frac{1}{2^n}$

for some $n \in \mathbb{N}$. Note that $\sim$ is an equivalence relation. Let X' be its set of equivalence classes and $f : X \to X'$ the natural surjection. Let Y consist of the images under f of the paths $[0, 1] \times \{0\}$ and $[0, 1] \times \{1\}$. Then Y is a strongly compatible set of paths and so $\mathcal{P} = \overline{Y}$ is a path space. Each individual path of Y could be represented by a topological arc, but we will now show that $\mathcal{P}$ is not induced by a Hausdorff space.

More specifically, we will show that given any Hausdorff space on X in which the two paths in Y are arcs, then so is another path not contained in $\mathcal{P}$. Let $g : [0, 1] \to X$ be given by $g(x) = f((x, 0))$ for $x \in [1 - \frac{1}{2^n}, 1 - \frac{1}{2^{n+1}})$ with n even, $g(x) = f((x, 1))$ for $x \in [1 - \frac{1}{2^n}, 1 - \frac{1}{2^{n+1}})$ with n odd and $g(1) = f((1, 0))$. Since g is clearly injective, to show that it is an arc it is enough to show that it is continuous. Thus we need to show that for any $x \in [0, 1]$ and open neighborhood U of $g(x)$ there is some open interval containing x whose image is contained in U. If x is not $1 - \frac{1}{2^n}$ for some natural number n or 1, this is clear. Otherwise the continuity of the arcs corresponding to the paths in Y shows that there are open sets I_0 and I_1 in $[0, 1]$ containing x with $f[I_i \times \{i\}] \subseteq U$ for $i \in 2$. Then $g(I_0 \cap I_1) \subseteq U$.

We have now shown that g is an arc, but the path corresponding to g is not a concatenation of segments of the paths in Y and so it is not contained in $\mathcal{P}$.

0.4 Basic notions

Here we will establish a basic set of definitions for working with (di)path spaces, many of which generalize those for graphs. In particular, whenever we give something a name known for graphs, it does indeed reduce to the corresponding graph definition for finite path spaces, unless mentioned otherwise.

The *ground set* $V(\mathcal{P})$ of a dipath space $\mathcal{P}$ is the union of the ground sets of its paths. Given a (di)path space $\mathcal{P}$ a *subspace* is a subset of $\mathcal{P}$ that is a (di)path space. One way to obtain such a subspace is to take the *restriction* $\mathcal{P}[X]$ of $\mathcal{P}$ to a set $X \subseteq V(\mathcal{P})$, which is defined as the set of all its paths whose ground set is completely contained in X. We write $\mathcal{P} - X$ for $\mathcal{P}[V(\mathcal{P}) \setminus X]$.

While just omitting a set of paths from a (di)path space will usually not yield a subspace, given a dipath space $\mathcal{P}$ and a set of paths $\mathcal{Q}$ the set of all $P \in \mathcal{P}$ which do not share a nontrivial segment with an element of $\mathcal{Q}$ is a subspace of $\mathcal{P}$. Thus when we speak of *deleting* a set of paths from a dipath space, this is what we mean. For path spaces we also again obtain a path space if $\mathcal{Q}$ is closed under inverses.

Now we want to translate some useful concepts from graphs. Let us start with connectivity. Given a path space $\mathcal{P}$ we write $x \sim y$ for x, y in its ground set if there exists some $P \in \mathcal{P}$ with minimum x and maximum y.

Lemma 0.4.1. *$\sim$ is an equivalence relation.*

Proof. It is clearly reflexive and symmetric, so let $x, y, z \in V(\mathcal{P})$ with $x \sim y$ and $y \sim z$ be given and let P and Q be paths witnessing this respectively. Since $P \cap Q$ is complete in Q, it has a maximum m. Let P' be the segment of P up to m and Q' the segment of Q starting from m. Then P' connects to Q' and their concatenation witnesses $x \sim z$. $\square$

We call the equivalence classes defined by $\sim$ *components* of $\mathcal{P}$ and call $\mathcal{P}$ *connected* if it has just one component. In general dipath spaces we define components and connectedness via the completion.

Next we want to find an equivalent to degrees. Degrees will not be helpful in every space, but they will be in the controlled spaces we will apply them in. Given a path space $\mathcal{P}$ and some $v \in V(\mathcal{P})$ we can define a relation on the nontrivial paths of $\mathcal{P}$ starting at v where two paths are equivalent if there exists a nontrivial path starting at v in $\mathcal{P}$ which is an initial segment of both. Clearly, this an equivalence relation. We call its classes the *outdirections* at v. The number of these directions is the *outdegree* of v. Similarly we can define *indirections* and the *indegree* of v. If $\mathcal{P}$ is a path space, the indegree and outdegree of v are the same and we call this the *degree* of v.

A circuit in our setting just consists of two independent paths. More precisely, if P is a path from some v to some different w and Q is a path from w to v not meeting P otherwise, then their (directed) completion is called a *(directed) circuit.* Any path space containing two different nontrivial paths with the same endpoints contains a circuit.

Lemma 0.4.2. *Let $\mathcal{P}$ be a path space and let $P, Q \in \mathcal{P}$ be different nontrivial paths with the same endpoints. Then $\mathcal{P}$ contains a circuit.*

Proof. Let x be the supremum in P of all y such that the segment of P up to y is also a segment of Q. By compatibility the segment of P up to x is then also a segment of Q. Let R be the segment of P from x to the first point z of P on Q after x and let S be the segment of Q between x and z. Then the completion of R and S is a circuit. $\square$

We will also need some niceness properties for spaces and subsets. A dipath space is called *finitary* if it is the directed completion of a finite set of paths. Given some dipath space $\mathcal{P}$, a set $X \subseteq V(\mathcal{P})$ is *simple* if it is a finite union of paths in $\mathcal{P}$. We also say that a dipath space is *simple* if its ground set is simple with respect to itself. Clearly every finitary dipath space is simple. We call a subset X of $V(\mathcal{P})$ *closed* if $X \cap P$ is complete for every $P \in \mathcal{P}$. It is easy to see that the set of the closed subsets of $V(\mathcal{P})$ is the set of closed sets of a topology. Since paths are closed by definition, so are finite unions of paths and thus every simple set is closed.

Note that given two closed sets A and B, every path from A to B contains a segment that is an *A-B-path*, that is a path which meets A exactly in its first point and B exactly in its last point.

0.4.1 Minors

We also give a definitions of minors and topological minors. Given a path space $\mathcal{Q}$ an $I\mathcal{Q}$ is a path space $\mathcal{P}$ with a map $\phi : V(\mathcal{P}) \to V(\mathcal{Q})$ such that for every $q \in V(\mathcal{Q})$ the set $\phi^{-1}(q)$ is connected in $\mathcal{Q}$ and such that for every path $Q \in \mathcal{Q}$ there exists some $P \in \mathcal{P}$ such that ϕ restricts to an order-preserving surjection from P to Q. If $\mathcal{P}$ has an $I\mathcal{Q}$ as a subspace, then $\mathcal{Q}$ is a *minor* of $\mathcal{P}$. For topological minors we need a bit more setup. Given an injective map $\psi : V(\mathcal{Q}) \to V(\mathcal{P})$ between path spaces $\mathcal{P}$ and $\mathcal{Q}$ we write ψ^* for the function mapping each path P of $\mathcal{P}$ with endpoints in the image of ψ to $\psi^{-1}(P)$ with the induced order. We say that a path space $\mathcal{P}$ is a $T\mathcal{Q}$ if there is an injective $\psi : V(\mathcal{Q}) \to V(\mathcal{P})$ such that ψ^* satisfies the following conditions:

1. ψ^* is injective with image $\mathcal{Q}$.
2. Whenever two paths in $\mathcal{Q}$ are internally disjoint, so are their inverse images under ψ^*.
3. Every path of $\mathcal{P}$ is a concatenation of segments of paths which ψ^* maps to a path of $\mathcal{Q}$.

If $\mathcal{P}$ has a $T\mathcal{Q}$ as a subspace, then $\mathcal{Q}$ is a *topological minor* of $\mathcal{Q}$. First, let us verify that topological minors are minors.

Lemma 0.4.3. *Every $T\mathcal{Q}$ is an $I\mathcal{Q}$.*

Proof. Let $\mathcal{P}$ be a $T\mathcal{Q}$ with witnessing function ψ. Let X be the image of ψ. For $v \in V(\mathcal{P})$ we write $R(v)$ for the set of those $x \in X$ from which v can be reached by a path of $\mathcal{P}$ avoiding $X - x$.

First we will show that $R(v)$ is nonempty for all $v \in V(\mathcal{P})$. If $v \in X$ we have $R(v) = \{v\}$, so assume otherwise. By the third condition there is a path $P \in \mathcal{P}$ with both endpoints in X and $v \in P$. Let Z consist of all those $z \in \psi^*(P)$ such that $\psi(z)$ is less than v in P. Furthermore, let s be the supremum of Z in $\psi^*(P)$ and let $x = \psi(s)$. Then the segment between v and x in P witnesses $x \in R(v)$. Indeed, if x is smaller than v in P than any point of X between them would show that s was not an upper bound for Z, and if x is larger than v in P any point of X between them would show that s was not the smallest upper bound.

We now fix a well-order on X and define a function ϕ mapping each point of $\mathcal{P}$ to the least element of $R(v)$ in that well-order. We claim that ϕ witnesses that $\mathcal{P}$ is an $I\mathcal{Q}$. First let us show that for every $q \in V(\mathcal{Q})$ the set $\phi^{-1}(q)$ of $\mathcal{Q}$ is connected. It is enough to show that for every $v \in \phi^{-1}(q)$ there is a path from $\psi(q)$ to x contained in $\mathcal{P}[\phi^{-1}(q)]$. Since $\psi(q) \in R(v)$ is witnessed by some path $P \in \mathcal{P}$ and then also $\psi(q) \in R(p)$ for every $p \in P$, it is enough to show that there is no smaller $x \in X$ with $x \in R(p)$ for some $p' \in P$. Indeed, if there was this would be witnessed by some path Q and then a concatenation of suitable segments of P and Q would show $v \in R(x)$, a contradiction.

Finally, we need to show that for every $Q \in \mathcal{Q}$ there is some $P \in \mathcal{P}$ such that ϕ restricts to an order-preserving surjection from P to Q. We claim that $P = \psi^{*-1}(Q)$ is as desired. It is enough to show that any path with both endpoints in X which meets $\phi^{-1}(q)$ for some $q \in \mathcal{Q}$ also contains $\psi(q)$, since any point of P where ϕ is not as desired will then give rise to a point of X which contradicts the choice of P.

To show this, let $R \in \mathcal{P}$ with endpoints $\psi(a)$ and $\psi(b)$ which contains some point of $\phi^{-1}(q)$ for $q \in V(\mathcal{Q})$. In particular, there is a path $S \in \mathcal{P}$ which starts in $\psi(q)$ and ends in its first point of R otherwise avoiding X. Concatenating S with segments of R gives paths R_1 and R_2 both ending in $\psi(q)$ and starting in $\psi(a)$ and $\psi(b)$ respectively. Then the preimages of R_1 and R_2 under ψ^* are internally disjoint and thus by the second condition so are they. Thus S consists only of $\psi(q)$ and $\psi(q) \in R$, as was desired.

□

These relations are indeed transitive.

Lemma 0.4.4. *Let $\mathcal{P}, \mathcal{Q}, \mathcal{R}$ be path spaces. If $\mathcal{R}$ is a (topological) minor of $\mathcal{Q}$ and $\mathcal{Q}$ is a (topological) minor of $\mathcal{P}$, then $\mathcal{R}$ is a (topological) minor of $\mathcal{P}$.*

Proof. Let us first consider minors. Let ϕ_1 be a function witnessing an $I\mathcal{Q}$ in $\mathcal{P}$ and ϕ_2 a function witnessing an $I\mathcal{R}$ in $\mathcal{Q}$ and let ϕ be the concatenation of ϕ_1 and ϕ_2 on the set A where this gives a value. Now ϕ witnesses that the induced space on A is an $I\mathcal{R}$ in $\mathcal{P}$.

Similarly, let ψ_1 be a function witnessing a $T\mathcal{Q}$ in $\mathcal{P}$ and ψ_2 a function witnessing a $T\mathcal{R}$ in $\mathcal{Q}$ and let ψ be the concatenation of ψ_1 and ψ_2. Let $\mathcal{S}$ be the path space induced by the image of $\psi_1^{*-1} \circ \psi_2^{*-1}$. Then ψ witnesses that $\mathcal{S}$ is a $T\mathcal{R}$ in $\mathcal{P}$. □

If we are only checking whether a graph G is a (topological) minor of a path space, we can work with simpler equivalent definitions. We define an IG as a path space $\mathcal{P}$ consisting of disjoint connected spaces $(A_v)_{v \in V(G)}$ and paths $(P_{vw})_{vw \in E(G)}$[2], where P_{vw} is a path from A_v to A_w not meeting any other A_x such that no two of these paths meet outside of the union of the A_v. We call the A_v its *branch sets*. An TG is an IG with all branch sets singletons. We then call the elements of its branch sets *branch vertices*.

Lemma 0.4.5. *Let G be a graph and $\mathcal{Q}$ its corresponding path space. Then $\mathcal{P}$ is an IG if and only if it is an $I\mathcal{Q}$ and $\mathcal{P}$ is a TG if and only if it is an $T\mathcal{Q}$.*

Proof. If $\mathcal{P}$ is an IG with branch sets $(A_v)_{v \in V(G)}$ and paths $(P_{vw})_{vw \in E(G)}$, then we can obtain a function ϕ showing that $\mathcal{P}$ is an $I\mathcal{Q}$ by mapping $x \in V(\mathcal{P})$ to v if $x \in A_v$ and to y if $x \in P_{yz} \setminus A_z$. If $\mathcal{P}$ is a TG with branch vertices a_v for $v \in \mathcal{Q}$, then the function ψ mapping every $v \in V(G)$ to a_v witnesses that $\mathcal{P}$ is a $T\mathcal{Q}$.

Conversely, if $\mathcal{P}$ is an $I\mathcal{Q}$ with function ϕ, we can set A_v for $v \in V(G)$ as $\phi^{-1}(v)$ and choose the P_{vw} such that ϕ restricts to an order-preserving surjection from P_{vw} to the path corresponding to the edge vw. If $\mathcal{P}$ is a $T\mathcal{Q}$ with function ψ, we can set A_v for $v \in V(G)$ as $\{\psi(v)\}$ and P_{vw} for $vw \in E(G)$ as the unique preimage under ψ^* of the path defined by that edge. □

0.4.2 Separations

Finally, we will define separations of path spaces. Recall that a separation of a graph is a pair (A, B) with $A \cup B = V(G)$ such that there is no edge between A

[2]Note that $P_{vw} = P_{wv}$, since these are the same edge.

and B. This is equivalent to saying that there is no path from A to B avoiding $A \cap B$ and we use this equivalent condition to formulate a definition for path spaces.

Thus given a path space $\mathcal{P}$ a *separation* is a pair (A, B) of subsets of $V(\mathcal{P})$ with $A \cup B = V(\mathcal{P})$ such that every path from A to B meets $A \cap B$. Just like for a graphs we define its *order* $|(A, B)|$ as $|A \cap B|$. We will call a separation of order k a *k-separation*. For a separation (A, B) we call $A \cap B$ its *separator* and A and B its *sides*.

A separation is small or cosmall if and only if one side is the whole ground set of $\mathcal{P}$. We call such separations *improper* and the other separations *proper*. We define $\mathcal{P}$ to be k-connected if it has no proper l-separation for $l < k$.[3]

As usual for separations of this form, we define the *inverse* $(A, B)^*$ of a separation (A, B) as (B, A) and a relation $\leq$ between separations as follows: $(A, B) \leq (C, D)$ if and only if $A \subseteq C$ and $D \subseteq B$. Clearly, if (A, B) is a k-separation, then so is (B, A). Thus * is an involution on the set of k-separations. By definition of $\leq$, * is also order-reversing, i.e. $(A, B) \leq (C, D)$ if and only if $(C, D)^* \leq (A, B)^*$. Define $S_k(\mathcal{P})$ for each finite k as the set of all k-separations of $\mathcal{P}$. Then $(S_k, \leq, ^*)$ is a separation system.

Lemma 0.4.6. *Let (A, B) be a finite order separation of a path space $\mathcal{P}$. Then A is closed in $\mathcal{P}$.*

Proof. Let $P \in \mathcal{P}$ be arbitrary. Since P meets $A \cap B$ only finitely often, there is a set of finitely many segments of P covering it which only meet $A \cap B$ in endpoints. Each of these segments is either completely contained in A or in B because (A, B) is a separation. Then $P \cap A$ is the finite union of those segments contained in A and the finite set $A \cap B \cap P$, so it is complete in P. $\square$

An essential tool for working with separations is the existence of suprema and infima. Recall that for graphs the supremum of separations (A, B) and (C, D) is $(A \cup C, B \cap D)$ and at least within the realm of finite order separations this is also true for path spaces.

Lemma 0.4.7. *Let (A, B) and (C, D) be finite order separations. Then $(A \cup C, B \cap D)$ is a finite order separation.*

Proof. Since any point that is not contained in A or C must be contained in both B and D we have $(A \cup C) \cup (B \cap D) = V(\mathcal{P})$. Thus, to show that $(A \cup C, B \cap D)$

[3]Usually one also requires a minimum size, but here this definition is more convenient.

is a separation, it suffices to prove that every path P from some $y \in A \cup C$ to some $z \in B \cap D$ meets $(A \cup C) \cap (B \cap D)$.

Since $(A \cap B) \cup (C \cap D)$ is finite and P meets this set, P has a last point x in this set. If $x \in A \cap B \cap C \cap D$, we are done. Otherwise let P' be the segment of P from x to z. Then P' avoids one of $A \cap B$ and $C \cap D$ and must therefore lie completely in $D \setminus C$ or $B \setminus A$. Therefore, $x \in (A \cap B \cap D) \cup (C \cap B \cap D) = (A \cup C) \cap (B \cap D)$.

Furthermore, since $(A \cup C) \cap (B \cap D) \subseteq (A \cap B) \cup (C \cap D)$, the separation $(A \cup C, B \cap D)$ has finite order. $\square$

Let $\mathcal{P}$ be a path space and X a closed subset of $V(\mathcal{P})$. Then separations of $\mathcal{P}$ induce separations of $\mathcal{P}[X]$, but not necessarily the other way around. There is, however, a way to get such a correspondence.

Lemma 0.4.8. *Let X be a closed set in a path space $\mathcal{P}$. The set of all nonempty restrictions of paths $P \in \mathcal{P}$ to $P \cap X$ is a strongly compatible set of paths.*

Proof. Since X is closed, this is a set of paths and it is clearly closed under segments, so it is enough to show that it is compatible. Let $P, Q \in \mathcal{P}$ meet X and let P' be the restriction of P to $P \cap X$. Then we need to prove that $P' \cap (Q \cap X) = P' \cap Q$ is complete in P'. Let $R \subset P' \cap Q$ be arbitrary. But R is a subset of $P \cap Q$ as well, which is complete in P. Thus R has a supremum x in $P \cap Q$ which agrees with the one in P. Since $R \subseteq X$ and X is closed, $x \in X$ and thus $x \in P' \cap Q$. It follows that x is the desired supremum. The proof for infima is analogous. $\square$

We call the completion of this set the *torso of X (in $\mathcal{P}$)*. Note that the torso of X in $\mathcal{P}$ has $\mathcal{P}[X]$ as a subspace. For graphs this definition of torso corresponds to adding an edge vw to the induced graph on X whenever there is an X-path between v and w.[4] The relative ease of giving such a definition is an advantage of working with path spaces compared to a more topological notion, where we would need to specify the whole topology.

Lemma 0.4.9. *If (A, B) is a separation of the torso of some $X \subseteq V(\mathcal{P})$, then there is some separation (E, F) with $A \subseteq E$, $B \subseteq F$ and $A \cap B = E \cap F$.*

Proof. Assume that P is a path from some $a \in A$ to some $b \in B$ avoiding $A \cap B$. Then this defines a path on $P \cap X$ from A to B in the torso of X, a contradiction. Thus we may arrange the components of $\mathcal{P} - (A \cap B)$ to obtain a separation as desired. $\square$

[4]This is not quite equivalent to the usual definition of torso, which is defined only for a part of a tree decomposition, but they coincide for well-behaved decompositions.

We will need one additional observation about the torso.

Lemma 0.4.10. *Let X be a closed set in a path space $\mathcal{P}$ and let $a, b \in X$. If P is a path from a to b in the torso of X, then there is some $P' \in \mathcal{P}$ with $V(P') \cap X \subseteq V(P)$.*

Proof. Since P is a path in the torso, it is a concatenation of paths $P_1, \ldots, P_k$, where each P_i is restriction of some $P_i' \in \mathcal{P}$ to $P_i' \cap X$. Then $\bigcup_{1 \leq i \leq k} P_i$ is a connected set which contains a and b and meets X only in points of P, completing the proof. $\square$

Chapter 1

Connectivity theory

1.1 Introduction

Since paths are our basic unit, the area of connectivity seems like a natural place to start. One of the basic tools of infinite graph theory is the star-comb lemma, stating that for every infinite vertex set U there is either an infinite comb with teeth in U or an infinite subdivided star with leaves in U. There is also a finite version of this lemma which considers large sets instead. In Section 1.2 we show that both of these generalize to path spaces.

Next we consider Menger's theorem, which is perhaps the fundamental theorem of connectivity and has therefore been extended and generalized in many different ways. Perhaps the most important of these is the Aharoni-Berger theorem proven in [1], which extends Menger's theorem to infinite sets of paths.

Of particular note for us are the generalizations of Menger's theorem proven for some of the topological path-like objects mentioned in the overall introduction. For instance [7] generalizes the Aharoni-Berger theorem to the Freudenthal compactification of graphs, [26] proves a version of Menger's theorem for topological arcs and [49] shows multiple variations of Menger's theorem for different classes of topological spaces.

In Section 1.3 we prove a version of Menger's theorem for path spaces.

Theorem 1.1.1. *Let $\mathcal{G}$ be a dipath space, $A, B \subseteq V(\mathcal{G})$ and k a natural number. Then either there is a set of size less than k meeting every A-B-path or a set of k disjoint A-B paths.*

This theorem is weaker than some of those mentioned above and in fact this must be the case given the counterexamples we provide later. However, it can be applied in a wide variety of settings, including those for which no version of Menger's theorem existed so far, like for the pseudo-lines from the preliminaries.

In Theorem 1.4.3 we consider a similar duality theorem for the number of disjoint paths starting and ending in the same set. The graph version of this theorem was proved by Gallai in [24] and it corresponds to the induction start of Mader's $\mathcal{S}$-path theorem. Since its dual object is more complicated, we do not state it in this introduction.

The proofs of both theorems use different versions of alternating walks. For Menger's theorem the proof proceeds mostly as for graphs, but for Gallai's theorem the usual proof, seen for instance in [48], translates the graph into a matching, which we cannot do. We instead define a new type of alternating paths and use reachability via them to construct the required witness.

Finally, in Section 1.5 we provide some counterexamples, both those against variations of Theorem 1.1.1 promised above as well as an example showing that path spaces do not admit spanning trees as we know them.

1.2 Star-comb lemma

To start with we should define stars and combs. Let κ be a cardinal and U a set. A *subdivided κ-star with leaves in U* is a $TK_{1,\kappa}$ with all the branch vertices of the side of size κ contained in U. A *κ-comb with teeth in U* consists of a chain of paths $(R_i)_{i\in I}$ together with κ disjoint $\bigcup_{i\in I} R_i$-U-paths.

Proposition 1.2.1. *Let $l \gg k$ be natural numbers, $\mathcal{P}$ a connected path space and $U \subseteq V(\mathcal{P})$ a set of size at least l. Then $\mathcal{P}$ contains either a subdivided k-star with leaves in U or a k-comb with teeth in U.*

Proof. List l elements of U as $(u_i)_{1\leq i\leq l}$. For $1 \leq i \leq l$ inductively define finite sets of paths $\mathcal{P}_i$ starting with $\mathcal{P}_1 = \{P_1\}$, where $P_1 = \{u_1\}$. For $i > 1$ add a u_i-$\bigcup \mathcal{P}_{i-1}$-path P_i to $\mathcal{P}_{i-1}$ if $u_i \notin \bigcup \mathcal{P}_{i-1}$. Set all undefined P_i for $i \in \omega$ as empty. Let T be the graph with vertex set l and an edge between natural numbers $n < m$ if $u_m \in P_n$ or P_m ends in P_n. Note that T is a tree.

First consider the case that T contains a vertex n of degree $k' \gg k$. If P_n contains k points u_m for $m > n$, then it obviously defines a k-comb with teeth in U. Otherwise there exist $k'' \gg k$ paths P_m for $m > n$ ending in P_n. These can only meet in P_n by construction. If k of them end in the same point of P_n, this gives a subdivided k-star with leaves in U. Otherwise k of them are disjoint, which gives a subdivided comb with leaves in U.

Thus we may assume T has maximum degree less than k and so contains a path $n_1, \ldots, n_k$. Then clearly P_{n_i} is nonempty for all $1 \leq i \leq k$. Construct a path in $\mathcal{P}$ by inductively following each P_{n_i} until it hits $P_{n_{i+1}}$. Adding segments to each of the u_{n_i} from the correspond P_{n_i} if necessary gives a k-comb with teeth in U. $\square$

With a parallel proof we also get the countable equivalent.

Proposition 1.2.2. *Let $\mathcal{P}$ be a connected path space and $U \subseteq V(\mathcal{P})$ an infinite set. Then $\mathcal{P}$ contains either a subdivided $\aleph_0$-star with leaves in U or an $\aleph_0$-comb with teeth in U.*

However, our proof method does not extend to higher cardinalities, since after infinitely many steps the set of everything constructed so far might not be closed.

1.3 Menger's theorem

In this section we prove our first main result, namely Menger's theorem for dipath spaces. Since we want to use augmenting paths, we first need to introduce walks.

Let $\mathcal{P}$ be a dipath space. A *walk* in $\mathcal{P}$ is a sequence $P_1, \ldots, P_n$ of paths of $\mathcal{P}$ for some natural number n such that the maximum of P_i is equal to the minimum of P_{i+1} for all $i < n$. In some respects we can work with walks as we do with paths: for instance walks have a first point (the first point of P_1) and a last point (the last point of P_n) and if the last point of a walk V is the first point of another walk W, we can *concatenate* them by simply appending W to V. Also, we say that a walk V is a *segment* of another walk $W = P_1, \ldots, P_n$ if V has the form $Q_i, P_{i+1}, \ldots, P_{j-1}, Q_j$ where $1 \leq i \leq j \leq n$, Q_i is a segment of P_i and Q_j is a segment of P_j.

If $W = P_1, \ldots, P_n$ is a walk, we write $\widehat{W}$ for the directed completion of $\{P_1, \ldots, P_n\}$. A *trail* is a walk $P_1, \ldots, P_n$ such that there is no nontrivial path which is a segment of both P_i and P_j for some $i \neq j$.

Note that even if $\mathcal{P}$ arises from the set of arcs of a Hausdorff space, say, its walks in this sense do not necessarily match the paths of that space. Indeed, this is impossible since there may be multiple spaces with different sets of paths, but the same set of arcs.[1] However, the simple notion of walk just defined will suffice for our purposes.

Our proof of Theorem 1.1.1 closely follows the augmenting paths proof in [14]. That proof consists of two main parts: In the first part one shows that if for some set of disjoint A-B-paths $\mathcal{P}$ there is an alternating path from A to $B \setminus \bigcup \mathcal{P}$, there is a larger set of disjoint A-B-paths. This is shown by taking the symmetric difference between $\mathcal{P}$ and the alternating path and then simply analyzing its components. In the second part one shows that if no such alternating path exists, there is an A-B-separator consisting of one point on each path in $\mathcal{P}$. This separator is chosen by taking the last point on each path in $\mathcal{P}$ where an alternating path starting in A ends (or the first point if there is none), but showing that this is indeed a separator takes some work.

Perhaps surprisingly, both parts can be emulated for path spaces without major changes, as long as we use the correct definition of alternating paths. For the first part, the following lemma helps with the analysis of the components.

Lemma 1.3.1. *Let $\mathcal{P}$ be a connected, finitary dipath space with every indegree and outdegree at most one. Then $\mathcal{P}$ is a directed path or circuit.*

Proof. If $\mathcal{P}$ contains a directed circuit, then it must actually be a directed circuit, since any other path meeting it would increase the degree at some point. Thus we may assume that it does not.

Since $\mathcal{P}$ is finitary, it is the completion of paths $P_1, \dots, P_n$. We will show inductively for $1 \leq k \leq n$ that there are k of the P_i such that the directed completion of these P_i is a dipath. For $k = 1$ this is trivial, so we assume $k > 1$. By induction hypothesis we may assume without loss of generality that the directed completion $P_1, \dots, P_{k-1}$ is some directed path Q. Since $\mathcal{P}$ is connected, there must be one of the other P_i, say, P_k, which meets Q. Because of the degree conditions this can only mean that P_k ends where Q starts or vice versa (we cannot have both since this would form a directed circuit). Then the directed completion of $P_1, \dots, P_k$ is a directed path, completing the proof. □

For the rest of this section, fix a dipath space $\mathcal{P}$, sets $A, B \subseteq V(\mathcal{G})$ and a finite set $\mathcal{Q}$ of disjoint A-B-paths in $\mathcal{G}$.

We call a trail $P_1, \dots, P_n$ of $\overline{\mathcal{P}}$ an *alternating path* (with respect to $\mathcal{Q}$) if it satisfies the following conditions:

[1] Indeed, one may construct a topological star either such that there is a topological path reaching each leaf and converging to the center or such that every path reaches only finitely many leaves.

1. Its first point is contained in $A \setminus V(\mathcal{Q})$.

2. For odd i we have $P_i \in \mathcal{P}$ and P_i does not share a segment with a $Q \in \mathcal{Q}$ or its inverse.

3. For even i the path P_i is a nontrivial segment of the inverse of a $Q \in \mathcal{Q}$.

4. No point outside $\mathcal{Q}$ occurs in multiple P_i.

The *application* of an alternating path $P_1, \ldots, P_n$ to the set $\mathcal{Q}$ is the directed completion of the set consisting of the P_i for i odd and the components of each $Q \in \mathcal{Q}$ after deleting the inverses of the P_i for i even. Note that the application is finitary by definition.

Now we can show that alternating paths ending in $B \setminus V(\mathcal{Q})$ actually do improve $\mathcal{Q}$ just as for graphs.

Proposition 1.3.2. *If there is an alternating path ending in $B \setminus V(\mathcal{Q})$, there is a set of disjoint A-B-paths $\mathcal{Q}'$ with $|\mathcal{Q}'| > |\mathcal{Q}|$.*

Proof. Let T be the application of the given alternating path W to $\mathcal{P}$

Let A' be the set of initial points of W and the paths of $\mathcal{Q}$ and B' their set of final points. Now define $\mathcal{Q}'$ to be the set of components of T meeting $A' \cup B'$.

Note that T has maximum indegree and outdegree one, exactly the points of A' have indegree zero and exactly the points of B' have outdegree zero. By Lemma 1.3.1 any element of $\mathcal{Q}'$ is then a path starting in A' and ending in B'. In particular, $|\mathcal{Q}'| > |\mathcal{Q}|$. $\square$

It remains to be shown that conversely the absence of such an alternating path implies the existence of a suitable separator.

Proposition 1.3.3. *If there is no alternating path ending in $B \setminus V(\mathcal{Q})$, there is a choice of one point from each element of $\mathcal{Q}$ which meets every A-B-path.*

Proof. For every $Q \in \mathcal{Q}$ let x_Q be the supremum in P of all points v such that there is an alternating path ending in v and let X be the set of all these points. We claim that X meets every A-B-path. Let S consist of the segment of each $Q \in \mathcal{Q}$ up to x_Q.

Assume for contradiction that there is an A-B-path R avoiding X. Since R does not include an alternating path to a point of $B \setminus V(\mathcal{Q})$ or $V(\mathcal{Q}) \setminus S$, it meets S; let y be its last point in S and Z the element of $\mathcal{Q}$ containing y. Since R avoids X, $y \neq x_Z$, so there is a z on Z after y such that there is an alternating path W ending in z. Let z' be the first point in W on the segment of Z between z and y and let W' be the concatenation of the segment of W until z' with the inverse of Z between y and z'. Then W' is an alternating path ending in y. Since W' meets $V(\mathcal{Q})$ only in S and y is the last point of R on S, the segment of R from y can meet W' only outside $V(\mathcal{Q})$ and in y. If they only meet in y, let W'' be their concatenation. Otherwise, let W'' be the concatenation of W' up to meeting the segment of R starting in y for the first time with R starting from that meeting point.

Now W'' is an alternating path with respect to the set consisting of the segment of each Q up to x_Q. But if W'' had a first vertex on the complementary segments, say on some path Q, it would then include an alternating path to that point, contradicting the choice of x_Q. □

Now the theorem follows easily.

Theorem 1.3.4. *Let $\mathcal{P}$ be a dipath space, $A, B \subseteq V(\mathcal{P})$ and k a natural number. Then either there is a set of size less than k meeting every A-B-path or a set of k disjoint A-B paths.*

Proof. Assume that there is no set of k disjoint A-B-paths. Then there is a set $\mathcal{Q}$ of disjoint A-B-paths of maximal size. By Proposition 1.3.2 there is no alternating path with respect to $\mathcal{Q}$ ending in $B \setminus V(\mathcal{Q})$. But then there must be a set of size $|\mathcal{Q}| < k$ meeting every A-B-path by Proposition 1.3.3. □

While we have only looked at A-B-paths so far, the same proof will work if we replace all occurences of 'A-B-paths' in Theorem 1.3.4 with 'paths from A to B'.

1.4 Mader's theorem

Now we will pivot to our second subject, Mader's theorem. Let us start by fixing some notation. Let $\mathcal{P}$ be a path space and let $\mathcal{S}$ be a set of disjoint subsets of $V(\mathcal{P})$. We write S for $\bigcup \mathcal{S}$. Since we work in an undirected setting in this section, when we talk about segments of a path, we also mean those of its inverse. An $\mathcal{S}$-*system* is a set of disjoint paths each starting in some $S \in \mathcal{S}$ and ending in some different $S' \in \mathcal{S}$. The corresponding dual object is defined similarly to the one in graphs, but we need a closure condition to avoid problematic limit behavior. A tuple $(X, \mathcal{Y})$ is called an $\mathcal{S}$-*witness* if X and all the $Y \in \mathcal{Y}$ are disjoint subsets of $V(\mathcal{P})$, $Y \cup X$ is closed for every $Y \in \mathcal{Y}$ and every path between distinct sets of $\mathcal{S}$ meets X or has a nontrivial segment which is a subset of some $Y \in \mathcal{Y}$. We will refer to X as the witness's *separator* and to $\mathcal{Y}$ as its *cover*. The *order* $|W|$ of an $\mathcal{S}$-witness $W = (X, \mathcal{Y})$ is $|X| + \sum_{Y \in \mathcal{Y}} \lfloor \frac{|\delta Y|}{2} \rfloor$, where δY is the number of $y \in Y$ with $y \in \bigcup \mathcal{S}$ or for which there is a nontrivial path starting in y otherwise avoiding Y and X.

Lemma 1.4.1. *Let $\mathcal{Q}$ be an $\mathcal{S}$-system and $W = (X, \mathcal{Y})$ an $\mathcal{S}$-witness. If $Q \in \mathcal{Q}$ does not meet X, then there is $Y \in \mathcal{Y}$ with $|Q \cap \delta Y| \geq 2$.*

Proof. By definition of witness there is some $Y \in \mathcal{Y}$ such that Y contains a nontrivial segment of Q. Let a be the infimum of $Q \cap Y$ in P and let b be the supremum. Then $a \neq b$ by choice of Y and $a, b \in Y$ since $Y \cup X$ is closed. Further, if $a \in \mathcal{S}$, then $a \in \delta Y$ by definition and otherwise the segment of Q up to a witnesses that $a \in \delta Y$. Similarly, $b \in \delta Y$. □

This immediately implies that the order of any $\mathcal{S}$-witness is an upper bound for the order of any $\mathcal{S}$-system.

We call an $\mathcal{S}$-system $\mathcal{Q}$ and an $\mathcal{S}$-witness $(X, \mathcal{Y})$ *dual* if X consists of at most one point of every $Q \in \mathcal{Q}$ and there are b_Y for all $Y \in \mathcal{Y}$ such that all points of $\delta Y - b_Y$ lie on paths of $\mathcal{Q}$ and $|\delta Y \cap Q - b_Y|$ is 0 for all $Y \in \mathcal{Y}$ if $Q \in \mathcal{Q}$ meets X and 0 for all but one $Y \in \mathcal{Y}$ and at most 2 for that one Y [2] if $Q \in \mathcal{Q}$ does not meet X.

Lemma 1.4.2. *Let $\mathcal{Q}$ be a finite order $\mathcal{S}$-system and W an $\mathcal{S}$-witness. Then they are dual if and only if they have the same order.*

Proof. Let $W = (X, \mathcal{Y})$. First assume that $\mathcal{Q}$ and W are dual witnessed by some b_Y. Then we get $|W| = |X| + \sum_{Y \in \mathcal{Y}} \lfloor \frac{|\delta Y|}{2} \rfloor \leq |X| + \sum_{Y \in \mathcal{Y}} \lceil \frac{|\delta Y - b_Y|}{2} \rceil$.

Calculating further, we obtain $|W| \leq |X| + \sum_{Y \in \mathcal{Y}} \sum_{Q \in \mathcal{Q}} \lceil \frac{|\delta Y \cap Q - b_Y|}{2} \rceil = |X| + \sum_{Q \in \mathcal{Q}} \sum_{Y \in \mathcal{Y}} \lceil \frac{|\delta Y \cap Q - b_Y|}{2} \rceil$. Using the bounds given by $\mathcal{Q}$ and W being dual, we can deduce $|W| \leq \sum_{Q \in \mathcal{Q}} (|Q \cap X| + (1 - |Q \cap X|)) = |\mathcal{Q}|$.

Conversely, we now assume $|W| = |\mathcal{Q}|$. Let $\mathcal{Q}_0$ be the set of $Q \in \mathcal{Q}$ avoiding X and $\mathcal{Q}_1$ its complement. Let f be a function mapping each $Q \in \mathcal{Q}_0$ to some $Y \in \mathcal{Y}$ as in Lemma 1.4.1. Clearly $|X| \geq |\mathcal{Q}_1|$ and $|\delta Y| \geq 2|f^{-1}(Y)|$. But since $|\mathcal{Q}_1| + \sum_{Y \in \mathcal{Y}} \lfloor \frac{2|f^{-1}(Y))|}{2} \rfloor = |\mathcal{Q}|$, our assumption of $|W| = |\mathcal{Q}|$ implies $|X| = |\mathcal{Q}_1|$ and $|\delta Y| \leq 2|f^{-1}(Y)| + 1$.

[2]We could require exactly 2 here, but that would make it more cumbersome to write down a correct b_Y in the proof of our main theorem.

For each $Q \in \mathcal{Q}$ we choose a set A_Q of two points from $\delta f(Q) \cap Q$. Since $|\delta Y| \leq 2|f^{-1}(Y)| + 1$ for any $Y \in \mathcal{Y}$, there is at most one point from δY not contained in some A_Q for $Q \in f^{-1}(Y)$. Let b_Y be such a point if it exists and outside of δY otherwise. Now we claim that these b_Y witness that W and $\mathcal{Q}$ are dual. From $|X| = |\mathcal{Q}_1|$ it is clear that X consists of at most one point from every $P \in \mathcal{Q}$. Furthermore by choice of b_Y the set $\delta Y \cap Q - b_Y$ is empty if $f(Q) \neq Y$ and A_Q otherwise, finishing the proof. $\square$

A set $\mathcal{S}$ is *single* if it is a set of singletons of $V(\mathcal{P})$ and *bounded* if an $\mathcal{S}$-system of size k does not exist for some $k \in \mathbb{N}$. Now we can state the main theorem of this section.

Theorem 1.4.3. *Let $\mathcal{S}$ be single and bounded. Then there are a $\mathcal{S}$-system and a $\mathcal{S}$-witness which are dual.*

The remainder of this section is taken up with the proof of this theorem. For this we fix a single, bounded $\mathcal{S}$.

Let $\mathcal{Q}$ be an $\mathcal{S}$-system. We will always assume without loss of generality that the paths in $\mathcal{Q}$ have no inner points in $\bigcup \mathcal{S}$. Given $x \in \bigcup \mathcal{Q}$ we write $\mathcal{Q}(x)$ for the element of $\mathcal{Q}$ containing x.

We call a trail $P_1, \ldots, P_n$ of $\mathcal{P}$ *alternating* (with respect to $\mathcal{Q}$) if it satisfies the following conditions:

1. For odd i the path P_i does not share a segment with a $Q \in \mathcal{Q}$ or its inverse.

2. For even i the path P_i is a nontrivial segment of a $Q \in \mathcal{Q}$ or its inverse

3. No point outside $\mathcal{Q}$ occurs in multiple P_i.

4. There are no even i and j with P_i a segment of $Q \in \mathcal{Q}$ and P_j a segment of its inverse.

Figure 2 shows a hopefully instructive example. We say that an alternating trail P *arrives* in some $x \in \bigcup \mathcal{Q}$ if $x \in P$ and there is some $y \in P$ before x such that the segment from y to x in P shares no nontrivial segment with $\mathcal{Q}(x)$.

The augmenting property of alternating trails can be deduced from that of alternating paths.

Lemma 1.4.4. *Let $\mathcal{Q}$ be a finite $\mathcal{S}$-system. If there is an alternating trail for $\mathcal{Q}$ starting and ending in different points of $S \setminus V(\mathcal{Q})$, there is an $\mathcal{S}$-system of greater size.*

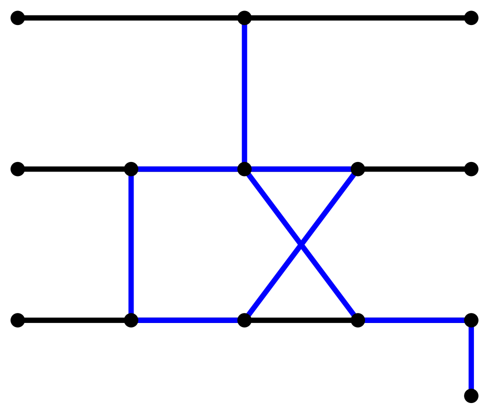

Figure 2: An example of an alternating trail (shown in blue)

Proof. Let W be such an alternating path. Without loss of generality each $Q \in \mathcal{Q}$ is not traversed backwards, otherwise reverse the path. Let A consist of the first point of W and of each $Q \in \mathcal{Q}$ and similarly let B consist of their last points. Since the first point of W does not lie on a path of $\mathcal{Q}$, W is an alternating path for A, B and $\mathcal{Q}$. Then Proposition 1.3.2 gives a set $\mathcal{Q}'$ of A-B-paths larger than $\mathcal{Q}$ which is also an $\mathcal{S}$-witness. $\square$

For the rest of this section we fix a finite $\mathcal{S}$-system $\mathcal{Q}$ of maximum size. We may assume without loss of generality that $\mathcal{Q}$ has no inner points in S. We write U for the set of those elements of S which are not endpoints of paths in $\mathcal{Q}$.

Our goal is now to construct an $\mathcal{S}$-witness of order $|\mathcal{Q}|$. For this we need a separator and a cover. Once we have found a suitable separator, the elements of its cover will be defined as the points reachable from a given starting point via alternating trails avoiding this separator. To achieve this, we will construct pairs (A, Z) with A a set of starting points and Z a separator recursively. The pairs which qualify for the construction are the so called split pairs. A pair (A, Z) of disjoint subsets of $V(\mathcal{P})$ is *split* if there is no Z-avoiding alternating trail between different points of A, Z consists of at most one point from every $Q \in \mathcal{Q}$ and there is no Z-avoiding alternating trail from some $a \in A$ to a $Q \in \mathcal{Q}$ meeting Z except if $a \in Q$.

The first condition is necessary for the disjointness of the cover we want to construct and the second and third to keep the order of the final witness low. Lemma 1.4.4 implies that $(U, \emptyset)$ is split, which will be the start of our recursion.

When working with a split pair (A, Z), we will often be given alternating trails with exactly one endpoint in A. To use the first condition for splitness to gain a contradiction, we will then need to combine two such paths. It is not clear how to do this in every case, however. For instance, two alternating trails P and R *cross* if they first meet an element of $\mathcal{Q}$ in the same point. As we will soon see, this is the only case where combining the two will prove troublesome. To combat this, points where this happens will be added to the separator in our recursion. However, we need to be a bit more selective since it may be that a point x where two alternating trails cross cannot be reached from A via alternating trails without passing through another such point, in which case it would be superfluous to add x to the separator.

Specifically, if P and R are Z-avoiding, start in different points of A and P does not cross any other Z-avoiding alternating trail starting in a diffent point of A before R the point at which P and R first cross is called a *crossing* (with respect to a given split pair (A, Z)). In particular, if P crosses any Z-avoiding alternating trail starting in a different point of A, P contains a crossing. When we talk about P and R witnessing a crossing x we always take them to stop when they first reach x.

Given alternating trails P and R we write $< P, R >$ for the union of the ground sets of P and R together with the ground set of any $Q \in \mathcal{Q}$ met by both P and R. Note that if (A, Z) is a split pair, P and R are Z-avoiding alternating trails and at least one of them starts in A, then $< P, R >$ avoids Z.

Now we can prove the first of two observations about combining alternating trails, namely that the only thing that can go wrong is the paths crossing.

Lemma 1.4.5. *Let P be an alternating trail starting at some a and R an alternating trail starting at b, which meet each other or meet a common $Q \in \mathcal{Q}$. Then P and R cross or there is an alternating trail from a to b contained in $< P, R >$.*

Proof. We will only consider the case that P and R meet, the other case is similar. We may assume that P and Q both end in their sole common point x. If P and R do not meet a common $Q \in \mathcal{Q}$, then concatenating gives the required alternating trail. So let $Q \in \mathcal{Q}$ be the first path met by P which R also meets. If P and R first meet Q in different points, then connecting the relevant segments of P and R via Q again gives the required alternating trail. Thus P and R cross in some v. $\square$

Applied to a split pair, we get the following.

Corollary 1.4.6. *Let (A, Z) be split. If P and R are Z-avoiding alternating trails starting at different points of A and they meet each other or meet a common $Q \in \mathcal{Q}$, then P crosses R.*

Proof. Assume that P does not cross R and let P' be the alternating trail obtained from applying Lemma 1.4.5 to P and R. Note that P' is Z-avoiding since P and R cannot use elements of $\mathcal{Q}$ which meet Z. This contradicts the fact that (A, Z) is split. $\square$

Given a split pair (A, Z), call a Z-avoiding alternating trail *acceptable* (with respect to (A, Z)) if it has no crossing before its final point. By definition at least one of the two paths witnessing a crossing is acceptable. We write $X(A, Z)$ for the set of all those $x \in \bigcup \mathcal{Q}$ such that there is an acceptable alternating trail from A arriving in x, but not in any other point of $\mathcal{Q}(x) - x$. This will be the set of points added to the separator of our split pair.

Lemma 1.4.5 allows us to combine alternating trails, but only if they do not cross. Our second key observation allows us to combine even crossing trails, as long as we may use one of them in either direction.

Lemma 1.4.7. *Let $Q \in \mathcal{Q}$, $a, b \in V(\mathcal{P})$ and $c, d \in V(Q)$ with $c \neq d$. Let P an be alternating trail from a to c and R an alternating trail from b to d, which avoid $Q - c$ and $Q - d$ respectively. Then there is an alternating trail from b to a contained in $< P, R >$ or an alternating trail from b to c in $< P, R >$ avoiding $Q - c$.*

Proof. Let v be the first point of R contained in $< P, P >$ and let R' be the segment of R up to v. If $v \notin \bigcup \mathcal{Q}$, then concatenating R' and the segment of the inverse of P from v to a gives an alternating trail as required. Otherwise, if P first meets $\mathcal{Q}(v)$ in a point $w \neq v$, combining R' with the segment of the inverse of P from w to a and the segment of $\mathcal{Q}(v)$ between v and w gives a suitable alternating trail. Finally, if P first meets $\mathcal{Q}(v)$ in v, then the segment of P following $\mathcal{Q}(v)$ starting in v cannot also end in v. Thus combining R' with the segment of P from v to c gives an alternating trail from b to c, which is as desired, since $\mathcal{Q}(v)$ must differ from $\mathcal{Q}(b)$ in this case. $\square$

Next we want to check that $X(A, Z)$ contains all crossings. In fact, we will show a slightly stronger statement.

Lemma 1.4.8. *Let (A, Z) be split and let x be a crossing. Then there is no Z-avoiding alternating trail starting in A first meeting $\mathcal{Q}(x)$ in a point other than x.*

Proof. Assume for a contradiction that there is a Z-avoiding alternating trail P starting in A and meeting $\mathcal{Q}(x)$ only in some $y \neq x$, which must then be its last point. Let B and C witness that x is a crossing, with C acceptable, and let b and c be their starting points.

First we will show that P must start in c. Indeed, if not by Corollary 1.4.6 P must cross C, which can happen only in x, but $x \notin P$.

Now applying Lemma 1.4.7 to B and P gives a Z-avoiding alternating trail Q starting in b. Since (A, Z) is split, Q cannot end in c, so it must end in y. But then Q first meets $\mathcal{Q}(x)$ in y, but every such alternating trail must start in $c \neq b$, a contradiction. □

Corollary 1.4.9. *Let (A, Z) be split. Then $X(A, Z)$ contains all crossings.*

In particular, Z-avoiding alternating trails which do not meet $X(A, Z)$ before their final point are acceptable, which will be the main use of this corollary. Now we have enough to prove the key lemma which will help us improve a split pair.

Lemma 1.4.10. *Let (A, Z) be split and let $x_1, x_2 \in X(A, Z)$. Then there is no Z-avoiding alternating trail starting in $\mathcal{Q}(x_1) - x_1$ and ending in $\mathcal{Q}(x_2) - x_2$ which does not otherwise meet these paths.*

Proof. Let D be an alternating trail as in the statement. We may assume that D does not meet any $Q \in \mathcal{Q}$ meeting $X(A, Z)$ except for its first and last point, otherwise we may work with an appropriate segment of D instead.

First we consider the case that x_2 is not a crossing. Let C be an acceptable alternating trail from some $a \in A$ arriving in x_2. Applying Lemma 1.4.7 to D and C gives an acceptable alternating trail arriving in $\mathcal{Q}(x_1) - x_1$ or $\mathcal{Q}(x_2) - x_2$ by Corollary 1.4.9, a contradiction to $\{x_1, x_2\} \subseteq X(A, Z)$.

We may thus assume that x_2 is a crossing witnessed by B and C, where C is acceptable and starts at a. Applying Lemma 1.4.7 to D and C gives a Z-avoiding alternating trail P starting in a and first meeting $\mathcal{Q}(x_1)$ in a point other than x_1 or first meeting $\mathcal{Q}(x_2)$ in a point other than x_2. This contradicts Lemma 1.4.8.

□

Let the set $X'(A, Z)$ consist of the endpoints of those $Q \in \mathcal{Q}$ containing points of $X(A, Z)$ except those which are contained in $X(A, Z)$ themselves. These will be the new starting points added to our split pair.

Lemma 1.4.11. *Let $\mathcal{Q}$ be an $\mathcal{S}$-system of maximum size and let (A, Z) be split. Then $(A \cup X'(A, Z), Z \cup X(A, Z))$ is split.*

Proof. First we will show that there is no $Z \cup X(A, Z)$-avoiding alternating trail between different points of $A \cup X'(A, Z)$. Since (A, Z) was split, such an alternating trail clearly cannot exist between different points of A. It also cannot exist between A and $X'(A, Z)$ since it would then be acceptable by Corollary 1.4.9, contradicting the definition of $X(A, Z)$. Any $Z \cup X(A, Z)$-avoiding alternating trail between different points of $X'(A, Z)$ contains an alternating trail as in Lemma 1.4.10, a contradiction.

To complete the proof, we assume for a contradiction that there is a $Z \cup X(A, Z)$-avoiding alternating trail P from some $a \in A \cup X'(A, Z)$ to $\mathcal{Q}(z)$ for some $z \in Z \cup X(A, Z)$ with $a \notin \mathcal{Q}(z)$. Any such path is acceptable by Corollary 1.4.9. Clearly, $a \notin A$, since this contradicts (A, Z) being split if $z \in Z$ and the definition of $X(A, Z)$ if $z \in X(A, Z)$. So $a \in X'(A, Z)$. Furthermore we have $z \notin X(A, Z)$ since otherwise P contains an alternating trail as in Lemma 1.4.10. So $z \in Z$.

Let x be the point of $X(A, Z)$ on $\mathcal{Q}(a)$ and let R be an acceptable alternating trail from A to x. Applying Lemma 1.4.7 to R and the inverse of P gives a Z-avoiding alternating trail B starting in A. Then B cannot end in $\mathcal{Q}(z)$ since (A, Z) was split, so it must end in $\mathcal{Q}(a)$ and meet that path only in a. Now it suffices to show that B is acceptable, since this will then contradict $x \in X(A, Z)$. If not, there is some crossing $y \neq x$ contained in $< P, R >$ and thus R meets $\mathcal{Q}(y)$. But R does not contain y, so this contradicts Lemma 1.4.8. $\square$

Call a pair (A, Z) *splendid* if it is split, the set $U = S \setminus V(\mathcal{Q})$ is a subset of A, for any $z \in Z$ the endpoints of $\mathcal{Q}(z)$ are contained in $A \cup Z$ and $X(A, Z) = \emptyset$. We can now easily construct such a pair by recursively applying Lemma 1.4.11.

Lemma 1.4.12. *A splendid pair exists.*

Proof. We construct an increasing sequence of pairs $(A_n, Z_n)_{n \in \mathbb{N}}$ satisfying the first three conditions for splendidness inductively, starting with $(A_0, Z_0) = (U, \emptyset)$. If (A_n, Z_n) has been defined already, we set (A_{n+1}, B_{n+1}) as $(A_n \cup X'(A_n, Z_n), Z_n \cup X(A_n, Z_n))$. By Lemma 1.4.11 (A_{n+1}, Z_{n+1}) is split. The second and third conditions are immediate by construction.

Since $\mathcal{Q}$ is finite and the Z_n meet every $Q \in \mathcal{Q}$ only once, the Z_n become eventually constant. But this means that there is some $n \in \mathbb{N}$ with $X(A_n, Z_n) = \emptyset$, so (A_n, Z_n) is splendid. $\square$

To finish the proof, we now need to construct a witness from a splendid pair (A, Z). As mentioned before, the separator will just be Z and the elements of the cover will be the sets $Y_a(A, Z)$ for $a \in A$ defined as the set of all $y \in V(\mathcal{P})$ such that there is a Z-avoiding alternating trail from a to y. If these sets are to form a form a cover, they first need to be disjoint.

Lemma 1.4.13. *Let (A, Z) be splendid and let $a, b \in A$ be different. Then $Y_a(A, Z)$ and $Y_b(A, Z)$ are disjoint.*

Proof. Assume otherwise and let $c \in Y_a(A, Z) \cap Y_b(A, Z)$. Let P be a Z-avoiding alternating trail from a to c and R a Z-avoiding alternating trail from b to c. By Corollary 1.4.6 P contains a crossing, contradicting Corollary 1.4.9. $\square$

Proving the other conditions for a witness will not be too difficult with the aid of the following lemma.

Lemma 1.4.14. *Let (A, Z) be splendid, let P be a Z-avoiding path and let $a \in A$ be arbitrary. If P meets $Y_a(A, Z)$, then $P \subseteq Y_a(A, Z)$.*

Proof. We may assume without loss of generality that the first point v of P is contained in $Y_a(A, Z)$. We will prove the statement by induction over the number of elements of $\mathcal{Q}$ met by P except its final point z. If P does not meet $\bigcup \mathcal{Q}$ before z, then for any $w \in P$ appending P up to w to an alternating trail witnessing $v \in Y_a(A, Z)$ and shortcutting if necessary gives $w \in Y_a(A, Z)$.

Otherwise let m be the maximum of all those points where P first meets some $Q \in \mathcal{Q}$ other than z. By the induction hypothesis the segment P' of P up to m is contained in $Y_a(A, Z)$. Every $Q \in \mathcal{Q}$ met by P before z is then also contained in $Y_a(A, Z)$ since we may simply extend a witnessing alternating trail along Q. Let $w \in P \setminus P'$ be arbitrary. If w is a point of some $Q \in \mathcal{Q}$ except z, then $Q \subseteq Y_a(A, Z)$ and thus $w \in Y_a(A, Z)$. Otherwise there is some last $q \in \bigcup \mathcal{Q}$ before w. Since $q \in Y_a(A, Z) \setminus X(A, Z)$ there is an acceptable alternating trail R from some $a' \in A$ arriving in $\mathcal{Q}(q) - q$. But an extension of R now shows $q \in Y_{a'}(A, Z) \cap Y_a(A, Z)$, so $a = a'$ by Lemma 1.4.13. Now combinining R with the segment of P from q to w and shortcutting if necessary gives $w \in Y_a(A, Z)$. $\square$

We are now well equipped to prove the main theorem of this section.

Proof of Theorem 1.4.3. Let $\mathcal{Q}$ be an $\mathcal{S}$-system of maximum order. Then by Lemma 1.4.12 there is a splendid pair (A, Z) for $\mathcal{Q}$. In this proof all alternating trails are tacitly assumed to be Z-avoiding. Let Y^* be the set of all points not contained in any $Y_a(A, Z)$ or Z and let $\mathcal{Y}$ consist of all the $Y_a(A, Z)$ together with Y^*.

We claim that $W = (Z, \mathcal{Y})$ is an $\mathcal{S}$-witness. By Lemma 1.4.13 the elements of $\mathcal{Y}$ are disjoint. To show that for all $Y \in \mathcal{Y}$ the set $Y \cup Z$ is closed, let $P \in \mathcal{P}$ be some path and let $P_1, \ldots, P_n$ be a cover of P with segments of P such that each P_i meets Z at most in its endpoints. To show that $P \cap (Y \cup Z)$ is complete, it now suffices to show that $P_i - Z$ is either completely contained in Y or does not meet it at all. This follows immediately by applying Lemma 1.4.14 to Z-avoiding segments of P_i covering $P_i - Z$. Finally, let P be some path between different points of S. If P meets Z, we are done. On the other hand, if P is Z-avoiding, its initial vertex is contained in some $Y \in \mathcal{Y}$. By Lemma 1.4.14 P is then completely contained in Y.

It remains to be shown that $\mathcal{Q}$ and W are dual. Note that for every $Y \in \mathcal{Y}$ we have $\delta Y = Y \cap S$ by Lemma 1.4.14 and $|Q \cap S| \leq 2$ for every $Q \in \mathcal{Q}$. We set $b_{Y_a(A,Z)} = a$ and choose b_{Y^*} arbitrarily. Since (A, Z) is splendid, each $Q \in \mathcal{Q}$ contains at most one element of Z.

If $Q \in \mathcal{Q}$ does not contain a point of Z, then Q is a subset of some $Y \in \mathcal{Y}$ by Lemma 1.4.14 and we are done. Thus we may assume that Q does contain some $z \in Z$. Let $x \in Q \cap S$. Then x is an endpoint of Q. It is enough to show that x is not contained in $Y - b_Y$ for any $Y \in \mathcal{Y}$. This is clear if $x = z$. Otherwise $x \in A$ by splendidness and thus $x \in Y_x$. But $b_{Y_x} = x$, completing the proof. $\square$

1.5 Counterexamples

Some natural-sounding alternate versions of Theorem 1.3.4 can be refuted by simple, known counterexamples. This is also shown in [49] with very similar examples originally given in [18] and [51].

Figure 3 shows a combination of such examples: a dominated ray, where the first vertex a and the end d have two extra neighbors each (b, c and e, f respectively) and b and c are also adjacent to the dominating vertex.

In this example there is no single vertex except a or d meeting every path from a to d, but there are not even two edge-disjoint paths between them. Thus

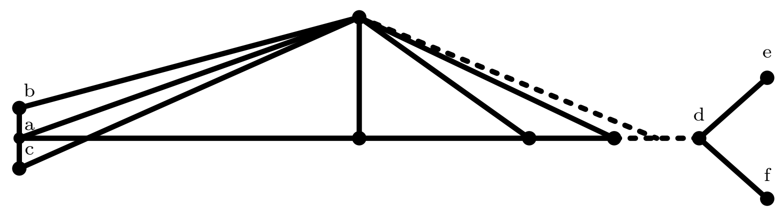

Figure 3: A combined counterexample to three Menger variations

a version of Menger for internally disjoint paths fails between a and d, one for fans between $\{a, b\}$ and d and one for edge-disjoint paths between $\{b, c\}$ and $\{e, f\}$.

One could also ask for a version for infinite cardinalities or even an Aharoni-Berger-type statement, but an example from [49] shows that a cardinality version of Theorem 1.3.4 already fails for $\aleph_0$. For this consider the space $[0, 1]^2$ and take A to be the points of the form $(0, \frac{1}{n})$ and B those of the form $(1, \frac{n-1}{n})$. In this example there are n disjoint A-B-paths for every natural number n, but not infinitely many.

One tool which is usually quite helpful in connectivity theory are spanning trees. While we can easily define a *tree* as a connected path space not containing a circuit, unfortunately not every path space contains a *spanning tree*, that is a tree with the same ground set, as a subspace. This is shown by the following example based on the infinite binary tree with every edge replaced by two subdivided edges.

Example 1.5.1. Let S be the set of finite 0-1-sequences, S^* the set of 0-1-sequences of length ω and $V = S \cup S^* \cup (S - \emptyset) \times 2$. Let $\mathcal{P}_1$ be the set of paths of the graph G on $S \cup (S - \emptyset) \times 2$ with edges of the form $(s, i)s$ for $s \in S$, $i \in 2$ or of the form $s(s', i)$ with $s, s' \in S$, $i \in 2$ and s' either $s.0$ or $s.1$. Extend the obvious partial order on S to a partial order on $V(G)$ by taking $(s, 0)$ and $(s, 1)$ to be the unique direct predecessors of s. Let A be the set of all rays of G which are ascending in this partial order and contain $s.i$ whenever they contain (s, i). For each $a \in A$ let $\hat{a}$ be obtained from a by adding the corresponding limit sequence from S^* as a final point and let $\mathcal{P}_2 = \{\hat{a}; a \in A\}$.

Let $\mathcal{P}$ be the completion of $\mathcal{P}_1 \cup \mathcal{P}_2$. Note that $\mathcal{P}$ is connected. Assume that T is a spanning tree of $\mathcal{P}$. Then for every nontrivial $s \in S$ there is an edge of the 4-cycle containing $(s, 0)$ and $(s, 1)$ in G which is not contained in T.

We will now recursively construct a ray $a \in A$, so assume the ray has been constructed up to some $s \in S$. Let $i \in 2$ be chosen such (s, i) is contained in the ray (or arbitrarily if s was the first vertex). Since T is a tree, it must avoid an edge of the 4-cycle of G containing s and $s.i$. We continue the ray using whichever of the two paths from s to $s.i$ of length two is not contained in T.

Then by construction no final segment of a is contained in T. Let x be the final point of $\hat{a}$. But since $\hat{a}$ is the only path containing x in $\mathcal{P}_1 \cup \mathcal{P}_2$, every path from S to x shares a nontrivial final segment with $\hat{a}$. In particular, no such path is contained in T, a contradiction.

Finally, we want to note that the path space given by the arcs of the closed unit disc is k-connected for every k, but given points a, b, c, d along the order of the boundary circle, there are no disjoint paths from a to c and b to d (corresponding to the definition of 2-linked for graphs). Thus high connectivity does not force linkedness for paths spaces, as it does for graphs.

1.6 Outlook

While the examples given in the prior section quite strongly constrain any further extension of Theorem 1.3.4, the situation is quite different for Theorem 1.4.3. The first question here is whether boundedness may be dropped to get a full analogue to Gallai's theorem. In particular this would imply that if there are arbitrarily many disjoint paths between points of some set S, there are also infinitely many. Beyond that it is natural to ask whether we can make the step from Gallai's theorem to Mader's theorem by allowing $\mathcal{S}$ to consist of finite (or even infinite) sets. This, however, seems to require new techniques, since neither our alternating paths approach nor emulating the induction step of [48] appear to work. We could also again enquire whether boundedness could be dropped even for these extensions. This does not work if we allow $\mathcal{S}$ to contain infinite sets, since in the example from the previous section in which there are arbitrarily many disjoint A-B-paths, but not infinitely many, we can set $\mathcal{S} = \{A, B\}$. Therefore the following two conjectures are close to the maximum we could still ask for.

Conjecture 1.6.1. *Let $\mathcal{P}$ be a path space and $\mathcal{S}$ be a set of finite subsets of $V(\mathcal{P})$. Then there are an $\mathcal{S}$-system and an $\mathcal{S}$-witness which are dual.*

Conjecture 1.6.2. *Let $\mathcal{P}$ be a path space and $\mathcal{S}$ be a bounded set of subsets of $V(\mathcal{P})$. Then there are an $\mathcal{S}$-system and an $\mathcal{S}$-witness which are dual.*

The star-comb lemma of Proposition 1.2.2 also suggests further research, since the proof method used here does not easily generalize to higher cardinalities. For infinite graphs [27] characterizes even k-connected sets of arbitrary cardinality. Successfully generalizing that article to path spaces, even in part, would thus be a large improvement. Questions about characterizing k-connected sets could also be asked for dipath spaces, but might quickly become unwieldy.

Another interesting question is whether Pym's theorem of [43] holds for path spaces. Since we need to deal with infinite sets of paths here, a positive resolution would probably require new techniques.

As an alternative to the theory developed here one could axiomatize the walks instead of the paths of a space. In this way one could preserve some more information when moving from topological spaces to 'walk spaces' and it could also allow one to define a product of such spaces which commutes with the product of topological spaces.

Chapter 2

Tree width

2.1 Introduction

While originally defined in [32] under a different name, tree decompositions of graphs only became an active field of research when they were rediscovered by Robertson and Seymour in the course of their graph minors project in [45]. An important subsequent development is the theory of abstract separation systems given in [13], where only separations and their relations are considered and all other structure is abstracted away. We will work within this framework, but since separation systems will become much more central in Chapter 3, we will cover them more fully in that chapter's introduction.

In graphs every tree decomposition induces a tree set of separations corresponding to the edges of the tree and conversely from such a tree set N one can construct a tree decomposition which induces N by having the vertices correspond to the consistent orientations of N, the edges to the unoriented separations and having the parts at each vertex be obtained as the intersection of all second components of the corresponding orientation. This turns out to be very useful, since it is often easier to construct a tree set of separations as opposed to directly constructing a tree decomposition, especially when criteria for the decomposition are naturally formulated in terms of separators, as for the upcoming examples of blocks and 2-blocks.

In Section 2.2 our goal is to find an analogue of tree decompositions for path spaces which corresponds to tree sets of separations in the same way. While a lot of the above translates to path spaces fairly well, there is one major problem: even in the simplest examples, like blocks, we will have to consider tree sets with chains of length greater than ω, which cannot represented using trees. Fortunately, it was proven in [28] that tree sets with chains of any length can be represented using tree-like spaces. Applying their work, we replace the decomposition tree with a tree-like space to obtain tree-like decompositions, which are otherwise mostly analogous to tree decompositions.

Section 2.3 covers the block decomposition, our first application, which was known long before tree decompositions. Given its simplicity, it is perhaps not surprising that this is mostly straightforward with the framework we established before. A decomposition along 2-separations would be a logical next step. For graphs such decomposition was first given in [50], again without using tree decompositions. It has been extended to infinite graphs by [44] and to infinite matroids by [2]. In Section 2.4 we use some of their techniques to obtain a similar decomposition for path spaces. Applying the two decompositions constructed so far, in Section 2.5 we extend the topological minor characterizations of tree width one and two to path spaces. We also show that, as in graphs, highly connected sets are necessary for large tree width.

The grid theorem, first proven in [46] as part of the graph minors project, states that every graph of high tree width has a large grid minor. While this is not a characterization, it does provide a qualitative minor-based criterion for high tree width. We will prove analogously that every path space which has tree width much greater than some k has a $k \times k$-grid minor. Since we have already mentioned that high tree width implies the existence of a large connected set, it might seem advantageous to follow the proof of [19], but that article contains some arguments which are not easily replicated in path spaces. However, there is another proof which makes use of these highly connected sets, namely the one in [25] based on constructing so called necklaces. Even though there were still some problems which needed to be fixed, overall this proof is surprisingly amenable to our generalization, as we will see in Section 2.6. In Section 2.7 we complete the proof of the grid theorem using the prior analysis of necklaces. From our grid theorem we then deduce in Section 2.8 that planar graphs have the Erdős-Pósa property following the proof in [46].

2.2 Tree-like decompositions

Recall that a *tree decomposition* of a graph G is a a pair $(T,(V_t)_{t\in V(T)})$, where T is a tree and $V_t \subseteq V(G)$, satisfying the following conditions:

1. $\bigcup_{t\in V(T)} V_t = V(G)$
2. For every $vw \in E(G)$ there is some $t \in V(T)$ with $\{v,w\} \in V_t$.
3. If a vertex x lies on the path of T between vertices s and t, then $V_s \cap V_t \subseteq V_x$.

When we try to give a similar definition for path spaces, we encounter two problems. The first is that the second condition explicitly mentions edges. This, however, can be resolved rather easily by using a different condition equivalent for graphs, namely that for each edge vw of T the set $V_v \cap V_w$ separates V_a and V_b for a, b in different components of $T - vw$.

The second is more fundamental: tree sets arising from path spaces will often have chain of separations of arbitrary length, but trees can only represent tree sets with no $\omega + 1$-chains. After observing this, [28] offers a solution. Instead of using a tree to represent a tree set, they use a tree-like space instead. We use this idea to define our decompositions.

As for tree decompositions in graphs, the parts of our decompositions will be obtained as an intersection of the second components of all separations belonging to some orientation and it will be useful to have a short notation for this. For any set o of separations of a path space we will write P_o for the set $\bigcap_{(A,B)\in o} B$. We will obtain the parts of our tree decompositions from orientations in this manner, just like for graphs. Note that by Lemma 0.4.6 these sets are closed.

A *tree-like decomposition* of a path space $\mathcal{P}$ is a pair $(T,(V_t)_{t\in V(T)})$, where T is a tree-like space and $V_t \subseteq V(\mathcal{P})$, satisfying the following conditions:

1. $\bigcup_{t\in V(T)} V_t = V(\mathcal{P})$
2. $V_s \cap V_t \subseteq V_x$ for all $s, t, x \in V(T)$ with $x \in L_T(s,t)$
3. For every $s, t \in V(T)$ and every edge $e \in E(L_T(s,t))$ with endpoints a and b we have that every path of $\mathcal{P}$ from V_s to V_t meets $V_a \cap V_b$.

It is *proper* if every empty V_t is a leaf and a limit point. We call each V_t a *part* of the tree-like decomposition. The *width* of a tree decomposition is the least cardinal κ such that $|V_t| < \kappa$ for all $t \in V_t$. The *tree width* of $\mathcal{P}$ is the least width of any tree-like decomposition.

In the finite case (that is if T and $\mathcal{P}$ are finite) tree-like decompositions are just tree decompositions. This is because finite tree-like spaces are just trees and the endpoints of each edge of the graph $\mathcal{P}$ cannot be separated by any of the separators, so they must be contained in a common part.

Given a tree-like decomposition $D = (T, (V_t)_{t \in V(T)})$ we define a separation $S(D, e)$ for each oriented edge e of T from v to w as $(\bigcup_{t \in C_v} V_t, \bigcup_{t \in C_w} V_t)$, where C_v and C_w are the components of $T - e$ containing v and w, respectively. Let $S(D)$ be the set of all these separations. Clearly, $S(D)$ is a tree set.

In the following, our goal is to reverse this operation. More specifically given a tree set N we want to construct a proper tree-like decomposition D with $S(D) = N$. The basis for this decomposition will be the tree-like space given by the following reformulation of [28, Theorem 4.15].

Theorem 2.2.1. *Let N be a regular tree set. Then there is a tree-like space T satisfying the following conditions:*

1. *The vertices of T are the consistent orientations of T.*
2. *The edges of T are the unoriented separations of N.*
3. *An unoriented separation e is incident with two orientations differing exactly in e*
4. *For $s, t \in N$ we have $s < t$ if and only if $L_T(v_s, v_{t^*}) \subseteq L_T(v_{s^*}, v_{t^*}) \subseteq L_T(v_{s^*}, v_t)$, where v_x for some separation x is the orientation incident with $\{x, x^*\}$ containing x.*

Let us write $T(N)$ for the tree-like space from Theorem 2.2.1 and $D(N)$ for the pair $(T(N), (P_o)_{o \in V(T(N))})$.

Proposition 2.2.2. *$D(N)$ is a tree-like decomposition of $\mathcal{P}$ with $S(D(N)) = N$.*

Proof. Let us first check that $D(N)$ is a tree-like decomposition. For the first condition, let any $x \in V(\mathcal{P})$ be given. The set $\{(A, B) : x \in B \setminus A\}$ is a partial consistent orientation. Thus, by Lemma 0.2.2 it can be extended to a consistent orientation o_x. Then we have $x \in V_{o_x}$ by definition.

For the second condition, let $o, p, x \in O(S)$ with $x \in L_{T(N)}(o, p)$ and $a \in P_o \cap P_p$ be given. Then $a \in P_{o \cup p}$ and in particular $a \in P_x$ since $x \subseteq o \cup p$.

Now let s, t, e, a, b be given as in the third condition and let P be a path from V_s to V_t. Let (A, B) be the oriented separation of e contained in b. Since e is an unoriented separation such that s and t contain different orientations of it, P must meet $A \cap B$. Thus it is enough to show that $A \cap B$ is equal to $V_a \cap V_b$. If $z \in V_a \cap V_b$ we must also have $z \in A \cap B$, since if $z \notin A$, $z \notin V_a$ and if $z \notin B$, $z \notin V_b$. Conversely let $z \in A \cap B$ and $(C, D) \in b - (A, B)$ be arbitrary. Assume $z \notin D$ for a contradiction. Then we have neither $(C, D) \leq (A, B)$ nor $(C, D) \leq (B, A)$. By nestedness (C, D) is thus larger than one of (A, B) or (B, A), but this contradicts the consistency of a or b respectively.

Finally, note that for each separation $n \in N$ we have $S(D(N), n) = n$ by construction and thus $S(D(N)) = N$. □

Since we want the decomposition to display the connectivity of the path space, it would be problematic for a part that lies between nonempty parts to be empty. Fortunately, this does not happen.

Lemma 2.2.3. *Let C be a chain of finite order separations in some connected path space $\mathcal{P}$, such that P_C and P_{C^*} are nonempty. Then for any partition (C_1, C_2) of C with $c_1 < c_2$ for all $(c_1, c_2) \in C_1 \times C_2$ the set $P_{C_1 \cup C_2^*}$ is nonempty.*

Proof. Let $a \in P_{C^*}$ and $b \in P_C$ be given and let P be a path from a to b. Clearly, P must meet $A \cap B$ for every $(A, B) \in C_1$. Let $x_{(A,B)}$ be the last vertex on P in $A \cap B$ and let x be the supremum on P of all the $x_{(A,B)}$ for $(A, B) \in C_1$. By choice of the $x_{(A,B)}$ we have $x \in B$ for all $(A, B) \in C_1$. Now it suffices to show that $x \in E$ for all $(E, F) \in C_2$.

Assume for a contradiction that $x \in F \setminus E$. Then P must meet $E \cap F$ somewhere before x, let y be the last such point. Since $y < x$, there must be some $x_{(A,B)}$ between them. But then $x_{(A,B)} \in F \setminus E$, contradicting $(A, B) < (E, F)$. □

Lemma 2.2.4. *If $\mathcal{P}$ is connected and nonempty and N consists only of finite order separations, $D(N)$ is proper.*

Proof. By Lemma 2.2.3 $L_{T(N)}(a, b)$ does not contain any c with P_c empty if P_a and P_b are nonempty. Since further N contains no trivial separation any empty part of $D(N)$ is a leaf and cannot have an adjacent edge, so it must be a limit point. □

We say that a set N of proper separations is *closed* if it contains all proper suprema of its chains. An important example is the set of all good separations in $S_k(\mathcal{P})$, as shown in Lemma 0.2.1. We write $g_k(\mathcal{P})$ for this regular tree set or just g_k if $\mathcal{P}$ is apparent.

Thus it will be useful to note a few properties of $D(N)$ for N closed. For this we first need to analyze limit separations. In graphs, since paths are finite, the supremum of a chain $(A_i, B_i)_{i \in I}$ is simply $(\bigcup_{i \in I} A_i, \bigcap_{i \in I} B_i)$. Since in our setting points in the separator can converge to other points, we cannot quite achieve this, but by applying Menger's theorem we can still control the supremum quite tightly.

Proposition 2.2.5. *Let $\mathcal{P}$ be a path space and $(A_i, B_i)_{i \in I}$ be a nonempty chain of separations of order $\leq k$ for some natural number k. Let $A = \bigcup_{i \in I} A_i$ and $B = \bigcap_{i \in I} B_i$. Then there is $X \subseteq V(\mathcal{P})$ such that $(A \cup X, B)$ is a separation of order at most k which is the supremum of C.*

Proof. There cannot be a set of $k + 1$ disjoint paths from A to B, since there would then be some (A_i, B_i) such that A_i contains the first points of all these paths. Let $\mathcal{Q}$ be a set of disjoint paths from A to B of maximal size. Without loss of generality $\mathcal{Q}$ contains (the singleton of) every point of $A \cap B$ as a path. Clearly $|\mathcal{Q}| \leq k$. By Theorem 1.3.4 there is a set S separating A from B which meets every $Q \in \mathcal{Q}$ exactly once. Then $S \subseteq B$, since if there was $s \in S \cap (A \setminus B)$, say meeting $Q \in \mathcal{Q}$, it would have to be the last point of A on Q and so a segment of Q would witness that any (A_i, B_i) with $s \in A_i \setminus B_i$ are not separations. Let $X = S \setminus A$. Then $(A \cup X, B)$ is a separation of order at most k which is clearly an upper bound for C. Let (E, F) be any other upper bound for C. It now suffices to show that $(A \cup X, B) \leq (E, F)$. Clearly, $A \subseteq E$ and $F \subseteq B$. But since all paths in Q clearly do not meet B before S, each point of S must be contained in E, completing the proof. $\square$

We call a set of separations *bounded* if every separation is comparable to some maximal element. We can now use the above result to control the unbounded parts.

Lemma 2.2.6. *Let N be closed and let $o \in O(N)$ be unbounded. If N consists only of separations of order at most k, then $|P_o| \leq k$. If o is not a leaf, then P_o is equal to the separator of the unique maximal separation of o.*

Proof. If o is not bounded, then by Zorn's Lemma there is a nonempty chain C in o whose supremum (A, B) is not contained in o. If all separations in N have order at most k, then by Proposition 2.2.5 so does (A, B). If (A, B) is improper, then o is a leaf and if (A, B) has order at most k then $|P_o| \leq |B| \leq k$.

Otherwise since o is an orientation, we have $(B, A) \in o$, implying $P_o \subseteq A \cap B$ by Proposition 2.2.5. Let $(E, F) \in o - (B, A)$ be arbitrary. By consistency we have $(E,F) \leq (B,A)$ or $(E,F) \leq (G,H)$ for some $(G,H) \in C$. This implies that (E,F) is not a maximal separation of o and that (E,F) is not greater than (B,A) and so (B,A) is the unique maximal separation of o. Further we also get $P_o = A \cap B$. If (A,B) has order at most k, then $|P_o| \leq k$. □

For a part P_o of $D(N)$ its *border* is the set of points of P_o also contained in some other part. The following observation justifies this name.

Lemma 2.2.7. *Let N be closed. Any part of $D(N)$ that is not an unbounded leaf is separated from its complement by its border.*

Proof. By Lemma 2.2.6 we may assume that o is bounded. Since any path from P_o to its complement must meet A for some $(A,B) \in o$, it must then also meet the separator of some maximal element of o and thus the border. □

Lemma 2.2.8. *Let N be closed. If $o, p, v \in O(N)$ are such that $v \in L_{T(N)}(o,p)$ and $o \neq p$, then P_v separates P_o and P_p.*

Proof. If v is a leaf, then $v \in \{o,p\}$ and we are done. Otherwise we may assume that v is bounded by Lemma 2.2.6. Then $L_{T(N)}(o,p)$ contains an edge e corresponding to a maximal element of v. By definition of tree-like decomposition the separator of $S(D,e)$ is a subset of P_v separating P_o and P_p. □

2.3 Blocks

The blocks of a graph are usually defined as the maximal connected parts not containing a cut vertex and their structure is described using a block tree containing both the blocks and the cutvertices as vertices. Equivalently this can also described via a tree decomposition with separators of size one whose parts are 2-connected or single vertices. When generalizing this to path spaces we will use the latter approach, since we have already defined tree-like decompositions.

To find an analogous decomposition of some fixed connected path space $\mathcal{P}$, we will proceed in three steps: first we characterize the good 1-separations of $\mathcal{P}$, then show that any path space with no good 1-separation is 2-connected and finally prove that any separation of such a block lifts to one of $\mathcal{P}$.

Let us start by making an observation about 1-separations.

Lemma 2.3.1. *Two 1-separations (A,B) and (C,D) can cross only if $A \cap B = C \cap D$.*

Proof. Let (A, B) and (C, D) be two crossing 1-separations and let v be the unique element of $A \cap B$ and w the unique element of $C \cap D$. Assume for a contradiction that $v \neq w$. Without loss of generality we may assume that $w \in A$ and $v \in C$. By Lemma 0.4.7 $(A \cup C, B \cap D)$ is a separation, but $(A \cup C) \cap (B \cap D) = (A \cap B \cap D) \cup (C \cap B \cap D) = (\{v\} \cap D) \cup (\{w\} \cap B) = \emptyset$. Since $\mathcal{P}$ is connected, this implies that one of its sides is empty, which must be $B \cap D$. But then $(B, A) \leq (C, D)$, a contradiction. □

We can use this to reach our first goal and characterize the good 1-separations of $\mathcal{P}$.

Proposition 2.3.2. *A proper* 1*-separation* (A, B) *of* $\mathcal{P}$ *is good if and only if* A *or* B *contains just one component of* $G - (A \cap B)$ *.*

Proof. For the forward direction let C_1 and C_2 be two components in A and D_1 and D_2 be two in B. Then moving C_2 to B and D_2 to A gives a separation crossing (A, B). For the other direction, let (A, B) be a separation satisfying the condition. By symmetry we may assume that A is the side with just one component. We know that $A \cap B$ has exactly one element, call it v. Clearly (A, B) cannot cross any separation with intersection v since $A + v$ must always be contained in one side. Now Lemma 2.3.1 gives the result. □

Similar to the known decomposition theorems for graphs, the sets P_o for some $o \in O(g_1)$ are promising candidates for our blocks: since they are not separated by any good 1-separation we can hope that they correspond to maximal 2-connected subspaces. A first step in this direction is showing that arbitrary 1-separations (not just good ones) fail to separate them.

Corollary 2.3.3. *For any* $o \in O(g_1)$ *and any* 1*-separation* (A, B) *the set* P_o *is completely contained in* A *or* B*.*

Proof. Let v be the unique point of $A \cap B$. If P_o met multiple components of $G - v$, let C and D be two. Then $(C + v, G - C)$ is a 1-separation satisfying the conditions of Proposition 2.3.2. So o must contain it or its inverse. Therefore one of C or D cannot be contained in P_o. □

In particular, this implies that any connected space with no good 1-separations is 2-connected, which was the second objective.

To show that the P_o are 2-connected, we still need show that separations of P_o induce ones of $\mathcal{P}$.

Lemma 2.3.4. *For any separation (A, B) of P_o for some set of 1-separations o there is some separation (E, F) of $\mathcal{P}$ with $A \subseteq E$, $B \subseteq F$ and $A \cap B = E \cap F$.*

Proof. Let (A, B) be a separation of P_o and assume that P is a path from some $a \in A$ to some $b \in B$ avoiding $A \cap B$. Since (A, B) is a separation in P_o, P contains some point $c \notin P_o$. Let $(C, D) \in o$ be some separation with $c \in C \setminus D$. Then the parts of P from a to c and from c to b both meet $C \cap D$, but this only has one element, a contradiction. Therefore $A \cap B$ separates $A - (A \cap B)$ and $B - (A \cap B)$ in $\mathcal{P}$ and we may arrange the components of $G - (A \cap B)$ not meeting either set arbitrarily to get a separation with the required properties. □

Having checked off all requirements, the desired statement is now immediate.

Corollary 2.3.5. *For any $o \in O(g_1)$ the subspace P_o is 2-connected.*

Proof. P_o has no proper 0-separation because this would contradict the 1-connectivity of $\mathcal{P}$ by Lemma 2.3.4. P_o has no proper 1-separation because this would contradict Corollary 2.3.3 by Lemma 2.3.4. □

With this result it now makes sense to call $D(g_1(\mathcal{P}))$ the *block decomposition* of $\mathcal{P}$ and its parts *blocks*. Our last goal here is giving a more concise characterization of the blocks. However, there is one small wrinkle here: blocks of size at most one can be contained in other blocks. To avoid this, we will characterize the maximal blocks instead. Since blocks never meet both sides of a 1-separation, these are uniquely determined as the maximal such sets by construction. Using Corollary 2.3.5, we can also show that they correspond to maximal 2-connected subspaces, as we hoped before.

Proposition 2.3.6. *The maximal blocks of $\mathcal{P}$ are exactly the maximal sets X such that the induced space on X is 2-connected.*

Proof. If X is a nonempty block, then the induced space on X is 2-connected by Corollary 2.3.5. On the other hand if the space induced by some set X is 2-connected, then there is no good 1-separation in $\mathcal{P}$ separating vertices of X and so X is contained in some block Y. □

2.4 2-Blocks

Next we will focus on 2-separations in 2-connected spaces. For graphs it is straightforward to construct a tree decomposition from the good 2-separations of the graph. While the parts of this tree decomposition will not necessarily be 3-connected, their torsos (obtained by making the separators complete) will either be 3-connected or cycles.

Using our results about blocks, in this section we will define a decomposition along the 2-separations of some fixed 2-connected path space $\mathcal{P}$. For this, as in the last section, we will proceed in three steps: first we will characterize the good 2-separations, then show that any path space with no good 2-separation is 3-connected or a circuit and finally prove a lifting lemma to show how this translates to the parts of our decomposition.

We start our search for a characterization of good 2-separations with a more general observation.

Proposition 2.4.1. *Let (A, B) be a proper k-separation in a k-connected path space $\mathcal{Q}$ for some finite k. Then any proper l-separation (C, D) of $\mathcal{Q}[A]$ for $l < k$ has a point of $A \cap B$ in $C \setminus D$ and $D \setminus C$.*

Proof. Assume not. By symmetry we may assume $A \cap B \subseteq D$. Since $\mathcal{Q}$ is k-connected, there is a path P from $C \setminus D$ to $D \setminus C$ avoiding $C \cap D$. Since (C, D) is a separation of $\mathcal{Q}[A]$, P must meet B. Since (A, B) is a separation of $\mathcal{Q}$, it follows that P meets $A \cap B$. Let P' be the segment of P up to its first point in $A \cap B$. Then P' is contained in $A \setminus (C \cap D)$, but meets both C and D, a contradiction. □

For 2-separations we obtain the following corollary.

Corollary 2.4.2. *For any proper 2-separation (A, B) the subspace $\mathcal{P}[A]$ is 1-connected.*

Proof. Let v and w be the vertices of $A \cap B$. If there were a proper 0-separation (C, D) of $\mathcal{P}[A]$, then $(C + v, D + v)$ would be a proper 1-separation of $\mathcal{P}[A]$, so by Proposition 2.4.1 both $C \setminus D$ and $D \setminus C$ must contain w, which is impossible. □

Now we can prove the desired characterization.

Proposition 2.4.3. *A proper 2-separation (A, B) is good if and only if $\mathcal{P}[A]$ or $\mathcal{P}[B]$ is 2-connected and A or B contains just one component of $\mathcal{P} - (A \cap B)$.*

Proof. Let v and w be the vertices of $A \cap B$. If both $\mathcal{P}[A]$ and $\mathcal{P}[B]$ are not 2-connected, then there are proper 1-separations (C, D) of $\mathcal{P}[A]$ and (E, F) of $\mathcal{P}[B]$. By Proposition 2.4.1 we may assume that $v \in (C \setminus D) \cap (E \setminus F)$ and $w \in (D \setminus C) \cap (F \setminus E)$. Then $(C \cup E, D \cup F)$ is a proper 2-separation crossing

(A, B). If both A and B contain at least two components of $\mathcal{P} - (A \cap B)$ exchanging one component from A with one from B gives a proper 2-separation crossing (A, B).

For the other direction, assume for a contradiction that the proper 2-separation (C, D) crosses (A, B), but (A, B) still satisfies the two conditions. Then the second condition implies that $A \cap B \neq C \cap D$. By symmetry we may assume that $\mathcal{P}[A]$ is 2-connected. Then $C \cap D$ must be contained in A. Indeed, otherwise $(C \cap A, D \cap A)$ would a separation of $\mathcal{P}[A]$ of order at most 1, which would have to be improper and thus A would be contained in C or D. Furthermore $C \cap D$ has to meet B since otherwise B would lie completely in C or D by Corollary 2.4.2. So $C \cap D$ must consist of one vertex $v \in A \cap B$ and one vertex $w \in A \setminus B$. This however contradicts Proposition 2.4.1 applied to the separation $(C \cap B, D \cap B)$ of B. $\square$

Note that the condition that one side should be 2-connected is only a restriction if the separator leaves exactly two components, otherwise, as we prove next, it is always satisfied.

Lemma 2.4.4. *Let (A, B) be a proper 2-separation. If A contains two components of $\mathcal{P} - A \cap B$, then $\mathcal{P}[A]$ is 2-connected.*

Proof. Let $A \cap B = \{v, w\}$. Assume that (C, D) is a proper 1-separation of $\mathcal{P}[A]$ with separator $\{x\}$. By Proposition 2.4.1 v and w lie on different sides of $C \cap D$ and neither is x. So x lies in one of the components of $A - \{v, w\}$. Let F be some other component. Moving all the other components in A to B gives a proper 2-separation, so by Corollary 2.4.2 the subspace defined by $F \cup \{v, w\}$ is 1-connected. Thus there is some path from v to w in this space, which then avoids x, a contradiction. $\square$

Proposition 2.4.3 and Lemma 2.4.4 together imply that it is sufficient to find any proper 2-separation with a 2-connected side to prove that $\mathcal{P}$ has a good 2-separation.

Lemma 2.4.5. *If there is a proper 2-separation (A, B) in $\mathcal{P}$ such that $\mathcal{P}[A]$ is 2-connected, then there is a good 2-separation with separator $A \cap B$ in $\mathcal{P}$.*

Proof. Let v and w be the two elements of $A \cap B$. If (A, B) is not good, then by Proposition 2.4.3 both A and B must contain at least two components. Let (C, D) be the separation with all but one component from A moved to B. Then by Lemma 2.4.4 $\mathcal{P}[D]$ is 2-connected. This implies that (C, D) is good by Proposition 2.4.3. $\square$

These tools will help us with our second step of proving that $\mathcal{P}$ is 3-connected or a cycle if it does not have a good 2-separation. To achieve this we will analyze the block decomposition of the sides of a 2-separation, following a strategy used for infinite graphs in [44, Theorem 6]. Let us first note some properties of these decompositions.

Just as the block decomposition of each side of a 2-separation of a 2-connected graph is always a path, it turns out to always be a pseudo-line in our case.

Proposition 2.4.6. *For any proper 2-separation (A, B) of $\mathcal{P}$ with $A \cap B = \{v, w\}$ the block decomposition of $\mathcal{P}[A]$ is a pseudo-line, whose endpoints are the unique blocks containing v and w, respectively.*

Proof. First note that v is contained in just one block V_o. Indeed, if not, then v lies in the separator of some proper 1-separation of $\mathcal{P}[A]$, contradicting Proposition 2.4.1. Similarly, w is contained in just one block V_p. If there was some vertex x of the tree-like space not on the pseudo-line between o and p, there must be some pseudo-line from x to o. It must contain at least one edge e not on the pseudo-line from o to p. Then the associated separation does not separate v and w, a contradiction to Proposition 2.4.1. $\square$

Moreover, in both cases no block has more than two vertices.

Lemma 2.4.7. *Let (A, B) be any proper 2-separation of $\mathcal{P}$. If any block of the block decomposition of $\mathcal{P}[A]$ has at least three vertices, then $\mathcal{P}$ has a good 2-separation.*

Proof. If A has just one block then $\mathcal{P}[A]$ is 2-connected by Corollary 2.3.5, so there is a good 2-separation by Lemma 2.4.5. Otherwise the two vertices $v, w \in A \cap B$ lie in different blocks V_o and V_p by Proposition 2.4.1.

First, let us assume that one of these blocks, say V_o, has at least three vertices. Since by Corollary 2.3.5 V_o is 2-connected, by Lemma 2.4.5 it suffices to show that there is no path from V_o to its complement not meeting the border of V_o or v. By Lemma 2.2.6 o is bounded and thus Lemma 2.2.7 implies that any such path cannot be contained in A and thus must meet w. But w is separated from V_o in $\mathcal{P}[A]$ by the border of V_o, so this is impossible.

If any other block V_x has size at least three, then similarly it suffices to show that the border of V_x separates V_x from its complement. If there were a path from V_x to its complement avoiding its border, however, it would need to meet v or w, which the border of V_x separates from V_x, a contradiction. □

Corollary 2.4.8. *If $\mathcal{P}$ has no good 2-separation and (A, B) is a proper 2-separation of $\mathcal{P}$, then there exists a linear order of A such that every $a \in A$ separates everything below a from everything above a within $\mathcal{P}[A]$.*

Proof. Proposition 2.4.6 gives us a linear order on the blocks of $\mathcal{P}[A]$. To define our order on A we set v as a minimum, w as a maximum and for any other $a_1, a_2 \in A$ we set $a_1 \leq a_2$ if the lowest block a_1 appears in is not bigger than the lowest block a_2 appears in. Let us first prove that this partial order is indeed linear by assuming that there are different a_1, a_2 which first appear in the same block V_x. Then by Lemma 2.4.7 V_x has no other points. But by Lemma 2.2.6 x has a lower neighbor z and so V_x must meet V_z, a contradiction to the choice of x.

It remains to be shown that a separates everything smaller than a from everything bigger than a. But if the first block a appears in has size 1, this is clear from Lemma 2.2.8 and if it has size 2 it follows from Lemma 2.2.6 together with the definition of tree-like decomposition. □

Using these two results, we now emulate the proof of [44, Theorem 6].

Proposition 2.4.9. *If $\mathcal{P}$ has no good 2-separation, then it is 3-connected or a cycle.*

Proof. If $\mathcal{P}$ is 3-connected then we are done. So we may assume that there is a proper 2-separation (A, B) with $A \cap B = \{v, w\}$. Now it suffices to show that $\mathcal{P}[A]$ and $\mathcal{P}[B]$ are paths from v to w and by symmetry we will only consider $\mathcal{P}[A]$.

By Corollary 2.4.2 $\mathcal{P}[A]$ is 1-connected, so there is a path from v to w in $\mathcal{P}[A]$. On the other hand any path in $\mathcal{P}[A]$ must agree with the order from Corollary 2.4.8 or its inverse, so we are done. □

For the third and final step we now want to prove a statement similar to Lemma 2.3.4. Just like in graphs and matroids, we cannot quite lift separations from the parts themselves, so we use the torso instead. We have already shown in Lemma 0.4.9 that separations of the torso lift to separations of $\mathcal{P}$, but we still need to show that for 2-separations we can preserve goodness under lifting.

Lemma 2.4.10. *For any good* 2*-separation* (A, B) *of the torso of some part* V_o*, there is some good* 2*-separation* (E, F) *of* $\mathcal{P}$ *with* $A \cap B = E \cap F$ *and* V_o *meeting both* $E \setminus F$ *and* $F \setminus E$*.*

Proof. If $\mathcal{P} - (A \cap B)$ has at least three components, then by Lemma 2.4.4 and Proposition 2.4.3, arranging the components such that one component meeting $A \setminus B$ is alone on one side and the other components are on the other side gives a good 2-separation fulfilling the requirements. So by Lemma 0.4.9 we may assume that it has two components. Let (C, D) be the unique separation of $\mathcal{P}$ with $A \subseteq C$ and $B \subseteq D$. By Proposition 2.4.3 (A, B) has a 2-connected side, say A. Then it suffices to show that C is 2-connected.

By Corollary 2.4.2 it is 1-connected, so by Proposition 2.4.6 it is enough to show that $A \cap B$ is contained in the same block, which will then be the only one. To do this, we will show that A is 2-connected in $\mathcal{P}$, not just in the torso.

Let $Y, Z \subseteq A$ be sets of size two. Then there are two disjoint Y-Z-paths P, Q in the torso. Applying Lemma 0.4.10 to P and Q gives Y-Z-paths P', Q'. Then these could meet only in some x outside of V_o, but x is separated from V_o by some $a, b \in V_o$. But since P' and Q' start and end in V_o they must then each use both a and b. But then P and Q must also meet in a and b, contradicting the fact that they are disjoint. Thus P' and Q' are disjoint and so A is 2-connected, completing the proof.

□

From this we can now deduce the desired characterization.

Theorem 2.4.11. *The torsos of the parts of the* 2*-block decomposition of* $\mathcal{P}$ *are* 3*-connected or cycles.*

Proof. The torso of any part V_o has no proper 0-separations or 1-separations since this would contradict the 2-connectivity of $\mathcal{P}$ by Lemma 0.4.9. It also has no good 2-separations by Lemma 2.4.10, since otherwise o would need to orient the lifted separation. Then it is 3-connected or a cycle by Proposition 2.4.9. □

Thus it makes sense to call $D(g_2(\mathcal{P}))$ the 2*-block decomposition of* $\mathcal{P}$.

2.5 Spaces of low tree width

One useful application of the block and 2-block decompositions is that they make it relatively easy to show that the forbidden minors for tree width one and two are K^3 and K^4, respectively, as for graphs.

Proposition 2.5.1. *If $\mathcal{P}$ has tree width at least two, it contains a TK^3.*

Proof. Without loss of generality $\mathcal{P}$ is connected. Since the block decomposition of $\mathcal{P}$ has width at least 2, $\mathcal{P}$ has a block X of size at least three. If $|X| = 3$, then $\mathcal{P}[X]$ is a circuit, which yield the desired TK^3. Otherwise $\mathcal{P}[X]$ contains two disjoint nontrivial paths P and Q. Applying Theorem 1.3.4 gives two disjoint paths between P and Q, again inducing a circuit. □

Lemma 2.5.2. *Any 3-connected path space of size at least four contains a TK^4.*

Proof. We may assume that the path space $\mathcal{P}$ has an infinite path P, since this is easy for graphs. Let P_1 and P_2 be disjoint segments of P of size at least three. By Theorem 1.3.4 there are disjoint paths Q_1, Q_2, Q_3 between P_1 and P_2. Let x_1 be whichever of the endpoints of the Q_i on P_1 comes between the other two in the order and let x_2 be analogous for P_2. Then there are three paths R_1, R_2, R_3 between x_1 and x_2 disjoint except for their endpoints. At most one of these can have length two, so we may assume that R_1 and R_2 have inner points y_1 and y_2 respectively. After deleting x_1 and x_2, $\mathcal{P}$ is still connected, so there is a path S from y_1 to y_2 avoiding x_1 and x_2. Let a be the last point of S on R_1 and let b be the first point on S after a contained in $R_2 \cup R_3$. Let S' be the segment of S from a to b. The completion of $\{R_1, R_2, R_3, S'\}$ is then a TK^4 with branch sets x_1, x_2, a, b as branch vertices. □

Proposition 2.5.3. *If $\mathcal{P}$ has tree width at least three, it contains a TK^4.*

Proof. It suffices to prove this for 2-connected $\mathcal{P}$, since we can then combine appropriate tree-like decompositions of the blocks of the components of some general path space.

First we consider the case that there is some torso of $D(g_2(\mathcal{P}))$ of size more than three that is 3-connected. Then it contains a TK^4 by Lemma 2.5.2. The paths making up this TK^4 must then come from some paths in $\mathcal{P}$, which will then form a TK^4 in $\mathcal{P}$, since they cannot leave their part via the same separation. Thus by Theorem 2.4.11 we may assume that every torso of $D(g_2(\mathcal{P}))$ of size at least three is a cycle.

Now we will define for each part o of $D(g_2(\mathcal{P}))$ of size more than three a set of separations S_o. We start by fixing some $x \in V_o$. For each $y \in V_o$ which is not equal or adjacent to x we define a separation s_y of the torso of V_o whose sides are the two intervals between x and y on the cycle. Let s'_y be a separation obtained by applying Lemma 0.4.9 to s_y and let S_o consist of all the s'_y.

Now let S be the union of all the S_o defined above and $g_2(\mathcal{P})$. We claim that S is nested. By definition the good separations do not cross any others. Furthermore each set S_o is nested since the s_y were nested and points of $V(\mathcal{P})\setminus V_o$ are always added to the side where the corresponding separator is contained. Finally, S_o and S_p for $o \neq p$ are nested because V_o is completely contained on one side of every separation of S_p and vice versa.

Then $D(S)$ is a tree-like decomposition of $\mathcal{P}$ of width at most two, completing the proof. $\square$

This suggests the following conjecture.

Conjecture 2.5.4. *If $\mathcal{P}$ has tree width at least $k \in \mathbb{N}$, then there is a finite graph G of tree width k such that $\mathcal{P}$ contains an IG.*

If true, this would show that (for finite tree width) the finite minors of $\mathcal{P}$ determine the tree width. The main theorem of the rest of this chapter, which is a version of the grid theorem, will at least imply a qualitative version of this. More precisely, we obtain that large enough tree width forces a minor of some fixed tree width $k \in \mathbb{N}$.

For this it will be useful to transform high tree width into something more tangible. While for the particular cases above we were already able to get specific minors, in this case we find a *k-connected* set, that is a set X such that for $Y, Z \subseteq X$ with $|Y| = |Z| \leq k$ there exist $|Y|$ disjoint Y-Z-paths, which will later be useful to find a minor of high tree width. While the proof of the corresponding statement for graphs uses induction, it can be adapted into our setting by using limit separations and Zorn's Lemma. A tree set has *partial width k* if all nonleaf parts have size at most $k+1$.

Proposition 2.5.5. *If $\mathcal{P}$ has no k-connected set of size at least k, then $\mathcal{P}$ has tree width at most $2k-2$.*

Proof. Let $\mathcal{X}$ be the set of all tree sets of $\leq k$-separations of $\mathcal{P}$ of partial width $2k-2$ ordered by inclusion. This defines a partial order and since the union of any chain is again a tree set with partial width $2k-2$, by Zorn's Lemma $\mathcal{X}$ has a maximal element T. Clearly, $T \neq \emptyset$. Furthermore T is closed under limits since adding any limit separation preserves nestedness and does not increase partial width.

Assume for a contradiction that T has a part D of size greater than $2k-1$. Then T contains a separation (C, D) for some set C. If $|C \cap D| < k$, adding the

$\leq m$-separation $(C+v, D)$ and its inverse for some $v \in D \setminus C$ to T again gives a tree set of partial width $2k-2$, contradicting its maximality.

Otherwise $C \cap D$ is not k-connected, let Y and Z witness this. Theorem 1.3.4 now gives us a set S of size less than $|Y|$ separating Y and Z. Write $X_1 = (C \cap D) \setminus Z \cup S$ and $X_2 = (C \cap D) \setminus Y \cup S$. Every component K of $D \setminus (C \cup S)$ is already a component of $D \setminus X_i$ for some $i \in \{1, 2\}$; let s_K be the separation $(K^c, K \cup X_i)$ for such an i and let T' be obtained from T by adding s_K and its inverse for all these K.

Since the s_K are $\leq k$- separations, to show that $T' \in \mathcal{X}$ it is enough to show that part corresponding to the orientation containing all the s_K^* has size at most $2k-1$. We calculate $|X_1 \cup X_2| \leq |(C \cap D)| + |S| \leq k + (k-1) = 2k-1$. Thus $T' \in \mathcal{X}$, contradicting the maximality of T. $\square$

2.6 Necklaces

The following section is heavily based on [25]. In that article Geelen and Joeris use necklaces to prove a generalization of the grid theorem. While we only need the basic grid theorem, their approach seems to be the easiest to adapt to our setting. As we will see in the following, many of the proofs can be adapted almost verbatim. For completeness we will write out all the proofs even where they are identical, but it may be helpful to refer to the original article for additional motivation.

We fix a connected path space $\mathcal{P}$. Given natural numbers t, s, l, n with $t \geq s$ a (t, s, l, n)-*necklace* is a pair $((B_i)_{i \in \mathbb{Z}_n}, Z)$ satisfying the following conditions:

1. The B_i are connected path spaces contained in $\mathcal{P}$ such that $V(B_i)$ avoids Z for all $i \in \mathbb{Z}_n$.

2. Z is a subset of $V(\mathcal{P})$ of size l

3. $V(B_i) \cup Z$ is finitary for all $i \in \mathbb{Z}_n$

4. For all $i, j \in \mathbb{Z}_n$ the sets B_i and B_j meet only if $i - j \in \{-1, 0, 1\}$

5. For each $i \in \mathbb{Z}_n$ the sets B_i and B_{i+1} meet in exactly t points if $i \neq n$ and in exactly s points if $i = n$

6. If $i \in \{2, \ldots, n-1\}$ there are t disjoint $B_{i-1} \cap B_i$-$B_i \cap B_{i+1}$-paths

7. If $i \in \{1, n\}$ there are s disjoint $B_{i-1} \cap B_i$-$B_i \cap B_{i+1}$-paths

8. Given $z \in Z$ and $i \in \mathbb{Z}_n$ there is a path from B_i to z contained in $B_i + z$.

A *weak* (t, s, l, n)*-necklace* is a a pair $((B_i)_{i \in \mathcal{Z}_n}, Z)$ satisfying all these conditions but the fourth. All definitions we make for necklaces are also intended for weak necklaces. For the sake of conciseness, we adopt the convention that when a necklace is called N its beads are called B_i and its set of hubs Z and when it is called N with a superscript we use these names with the same superscript added. We call the sets B_i *beads* of N and the elements of Z its *hubs*. We write $V(N)$ for the union of all beads and Z. We say that x is an *inner point* of B_i if $x \in B_i$ and $x \notin B_j$ for any $j \neq i$. The *rank* of N is $t + s + l$ and n is called its *length*.

The main difference between our definition of necklaces and the one in [25] is that while the beads in their definition are disjoint and connected by sets of edges of specified sizes, we instead have the beads meet in sets of those sizes. This is necessary because path spaces need not have any edges.

A necklace N is *supported* by a set U if every bead of N contains a point of U and supported by a necklace N' if every bead of N contains a bead of N'. We call N *nontrivial* if $t \geq 1$. Call a partition of $\mathbb{Z}_n$ into intervals $(I_1, \ldots, I_m)$ a *good partition (of* $\mathbb{Z}_n$*)* if $1 \in I_1$ and $n \in I_m$. In particular, we write $P(i_1, \ldots, i_k)$ for the good partition $([1, i_1 - 1], [i_1, i_2 - 1], \ldots, [i_{k-1}, i_k - 1], [i_k, n])$. Given a good partition $P = (I_1, \ldots, I_m)$ of $\mathbb{Z}_n$, the *contraction of* N *to* P is the necklace N' with $Z' = Z$ and with each bead B'_i consisting of the union of all those B_j with $j \in I_i$. Note that reversing the order of a necklace again gives a necklace. A sequence $y_1, \ldots, y_n$ *tracks* N if y_i is an inner point of B_i.

Lemma 2.6.1. *Let* $m \geq k \gg n$ *be natural numbers, let* N *be a* (t, s, l, m)*-necklace and let* Y *be a sequence of points in different beads of* N *of length* k*. Then there exist a* (t, s, l, n)*-necklace* N' *supported by* N *and a subsequence of* Y *tracking* N'*.*

Proof. Without loss of generality each y_i is an inner point of a bead, otherwise we can contract the necklace and take a subsequence to make this true. By Erdos-Szekeres we can now find a subsequence which occurs in the same order in N, reversing N if necessary. Now N' can be obtained by contracting N to track this subsequence. $\square$

Before we continue, let us quickly review the structure of the argument from [25]. Given a highly connected set U, we want to find a necklace of sufficient rank and length. To start with, they construct a rank 1 necklace supported by U. During the rest of the construction we always maintain the property of being supported by U. Since later arguments will only work for nontrivial necklaces, they iterate this to find either a nontrivial necklace or a necklace which is already as desired.

Once they have such a starting necklace, to improve it they first analyze paths connecting nonadjacent beads, which they call long jumps. They show that one can contract down our necklace such that there are either many or no such jumps. In the first case, they use such paths to increase the rank of the necklace, so one may assume that there are none. Now they use the high connectivity of U to find more disjoint paths through the necklace than its rank provides (after deleting the hubs) and using the structure imposed by the absence of long jumps, these can then be used to construct a necklace of higher rank. The desired result then follows by induction.

When generalizing this to path spaces, the first difficulty arises in the induction start ([25, Lemma 5.1]), which is there proven via induction on the number of edges. We will instead deduce it from Proposition 1.2.1.

Lemma 2.6.2. *Let $m \gg n$ be natural numbers and $\mathcal{P}$ a connected path space. If N is a $(0,0,0,m)$-necklace, then $\mathcal{P}$ contains either a $(0,0,1,n)$-necklace or a $(1,0,0,n)$-necklace supported by N.*

Proof. Let s and t be natural numbers with $m \gg s \gg t \gg n$ and let $B = \bigcup_{1 \leq i \leq n} B_i$. We start by inductively defining sets of paths $\mathcal{P}_i$ for $1 \leq i \leq m$ starting with $\mathcal{P}_1 = \emptyset$. At each step i we take P_i to be a C-D-path, where C is the component of $B \cup \bigcup \mathcal{Q}_{i-1}$ containing B_1 and D is the union of all other components. If P_i starts in a bead meeting $\bigcup \mathcal{P}_{i-1}$ then we add a $\bigcup \mathcal{P}_{i-1}$-$P_i$-path Q_i to its beginning, trivial if possible. Now $\mathcal{P}_i$ is obtained by adding P_i to $\mathcal{P}_{i-1}$.

If t elements of $\mathcal{P}_m$ meet one bead B_k, then let $\mathcal{T}$ be the subspace of $\bigcup \mathcal{P}_m$ in B_k and let U be the set of those points of B_k which are endvertices of paths Q_l. Applying Proposition 1.2.1 to $\mathcal{T}$ and U gives either a n-comb with teeth in U or a subdivided n-star with leaves in U. In the first case we obtain a $(1,0,0,n)$-necklace by having each bead consist of a part of the comb's spine, a segment of the relevant P_i starting at the spine and the bead reached by P_i. In the second case we obtain a $(0,0,1,d)$-necklace with the center vertex z of the star in Z and each bead consisting of the corresponding bead of N together with its path to z excepting z.

So we may assume that each bead meets less than t elements of $\mathcal{P}_m$. Let U be the set of all endpoints of P_i. Applying Proposition 1.2.1 to $\mathcal{P}_m$ and U gives either an s-comb with teeth in U or a subdivided s-star with leaves in U. By our assumption we may now find an n-comb with teeth in U or a subdivided n-star with leaves in U in which each vertex of U is contained in a different bead. Then we obtain one of the required necklaces as before. □

The construction of a nontrivial necklace in [25, Lemma 5.2] then requires no significant changes.

Lemma 2.6.3. *Let m, n and θ be natural numbers with $m \gg n$ and $m \gg \theta$ and let $U \subseteq V(\mathcal{P})$ be a θ-connected set of size m. Then $\mathcal{P}$ contains a $(0,0,\theta,n)$-necklace supported by U or a $(1,0,0,n)$-necklace every bead of which contains at least θ points of U.*

Proof. Let $m_0, \dots, m_\theta$ be natural numbers with $m = m_0 \gg \dots \gg m_\theta \gg n$. Then $\mathcal{P}$ contains a $(0,0,0,m_0)$-necklace, we choose k maximal such that $\mathcal{P}$ contains a $(0,0,k,m_k)$-necklace supported by U. If $k = \theta$, we are done.

Otherwise let N be a $(0,0,k,m_k)$-necklace supported by U and Z its set of hubs. Since U is θ-connected, all but θ vertices of U are contained in a single component C of $\mathcal{P} - Z$. In particular, by Lemma 2.6.2 C contains a $(0,0,1,m_{k+1})$-necklace or a $(1,0,0,m_{k+1})$-necklace. In the first case, adding Z to this necklace gives a contradiction to the maximality of k. In the second case, contracting N to a good partition with all intervals long enough gives a $(1,0,0,n)$-necklace such that each bead contains at least θ points of U. □

The arguments about long jumps in [25, Section 6] can also be emulated with little trouble. An (i,j)*-jump* for $i,j \in \mathbb{Z}_n$ in a necklace N is a path starting in B_i and ending in B_j with no inner vertex in $V(N)$. A *long jump* is an (i,j)-jump with $i - j \notin \{-1,0,1\}$. A necklace is called *long-jump-free* if it has no long jump and *jumpy* if it has a $(1,i)$-jump for every index i.

Lemma 2.6.4. *Let $m \gg n$ be natural numbers and N a nontrivial (t,s,l,m)-necklace. Then $\mathcal{P}$ has a (t,s,l,n)-necklace supported by N which is either long-jump-free or jumpy.*

Proof. Let n' be a natural number with $m \gg n' \gg n$ and let $m_0, \dots, m_{n'}$ be natural numbers with $m \gg m_0 \gg \dots \gg m_{n'} \gg n$. Define S to be the set of all (i,j) with $1 \leq i \leq j \leq m$ such that N has an (i,j)-jump. Given $S' \subseteq S$ we write $I(S') = \{x \in [1,m]; x \in [i,j] \forall (i,j) \in S'\})$. Let $S' \subseteq S$ be maximal such that $|I(S')| \geq |m_{|S'|}|$. We write $S' = \{(a_1,b_1), \dots, (a_k,b_k)\}$ and $I(S') = [a,b]$.

If $k < n'$, we obtain N' by contracting N to a good partition with each interval containing at least m_{k+1} points of $I(S')$. Then any long jump in N'

would correspond to a pair (i, j) with $|I(S' + (i, j))| \geq m_{k+1}$, contradicting the maximality of S'.

Now we may assume $k = n'$. Then without loss of generality we have $b_1 > \cdots > b_n$. Let N' be obtained by contracting N to $P(b_1, \ldots, b_n)$. Then the pairs $(a_i, b_i) \in S$ witness that N' is jumpy. $\square$

Lemma 2.6.5. *Let $m \gg n$ be natural numbers and let N be a nontrivial jumpy (t, s, l, m)-necklace with $s \geq 1$. Then $\mathcal{P}$ contains a $(t+1, s-1, l, n)$-necklace or a $(t, s-1, l+1, n)$-necklace.*

Proof. Let P_i be the $(1, i)$-jump for all $3 \leq i \leq m-1$, let x_i be its endvertex in B_1 and y_i its endvertex in B_i. Without loss of generality y_i is an inner vertex of B_i, otherwise contract accordingly. We may also assume that P_i and P_j never meet outside B_1 for different $1 \leq i, j \leq m-1$, otherwise we extend B_1 appropriately.

By definition of necklace there is a set of s disjoint B_m-B_1-paths in B_i. Now we want to extend these paths into disjoint connected, finitary subsets $T_1, \ldots, T_s$ of $V(B_i)$ such that all the x_i are contained in their union. We do this recursively: if $T_1, \ldots, T_n$ have been constructed so far and some x_k is not yet contained in any of them, we add a $\bigcup_{1 \leq i \leq n} T_i$-$\{x_k\}$-path to whichever T_i it starts in.

Without loss of generality we have $|I| \gg n$ for the set $I = \{3 \leq i \leq m-1; x_i \in T_s\}$. Obtain T from T_s by adding all P_i with $i \in I$.

By Lemma 2.6.3 T contains a necklace N^1 of rank 1 with length $n_1 \gg n$. Applying Lemma 2.6.1 gives a necklace N^2 of rank 1 with length $n_2 \gg n$ such that some subsequence $y_{\alpha_1}, \ldots y_{\alpha_n}$ of $y_3, \ldots, y_{m-1}$ tracks N^2. Let N^3 be obtained from N by contracting to $P(\alpha_1, \ldots, \alpha_n)$ and let N^4 be obtained from N^3 by replacing N^3_1 with $\bigcup_{1 \leq i < s} T_i \cup \bigcup_{2 \leq j < \alpha_1} B_j$. Then N is a $(t, s-1, l, n)$-necklace.

We observe that $Z^2 \cap V(N^4) = \emptyset$, $Z^4 \cap V(N^2) = \emptyset$ and further that $B^2_i \cap B^4_j$ is $\{y_{\alpha_i}\}$ if $i = j$ and empty otherwise. Let N^5 be the necklace with beads $B^5_i = B^2_i \cup B^4_i$ and hubs $Z^5 = Z^2 \cup Z_4$. We will prove that this is the required necklace. The first, second, third, fifth and eighth conditions are immediate with the observations above. For the fourth condition note that y_{α_i} is an inner point of B^2_i and of B^4_i. To prove the sixth condition, it suffices to show that y_{α_i} does not separate B^2_{i-1} and B^2_{i+1} in B^2_i if these sets meet B^2_i at all. This, however, follows from the fact that y_{α_i} is not an inner point of any path in T. The seventh condition is witnessed by $T_1, \ldots, T_{s-1}$. $\square$

Corollary 2.6.6. *Let $m \gg n$ be natural numbers and let N be a nontrivial necklace of rank r and length m. Then there exists a nontrivial necklace of length n supported by N which either has rank $r+1$ or has rank r and is long-jump-free.*

Proof. By Lemma 2.6.4 we may inductively apply Lemma 2.6.5 until the resulting necklace is either long-jump-free or a jumpy $(t', 0, l', n)$-necklace N^*. In the first case, we are done. In the second case, N^* can be made into a $(t', 1, l', n)$-necklace, which then has rank higher than N, by distributing the points of a $(1, n)$-jump appropriately. □

[25, Lemma 7.1] is an easy consequence of Menger's theorem, for graphs as for path spaces.

Lemma 2.6.7. *Let k be a natural number, $\mathcal{Q}_1$ and $\mathcal{Q}_2$ be sets of k disjoint paths and $\mathcal{C}$ a set of connected path spaces which has k disjoint elements such that every $Q \in \mathcal{Q}_1 \cup \mathcal{Q}_2$ meets every $C \in \mathcal{C}$. Then there are k disjoint paths from the initial vertices of paths in $\mathcal{Q}_1$ to the final vertices of paths in $\mathcal{Q}_2$.*

Proof. Without loss of generality $\mathcal{C}$ consists of k disjoint elements, otherwise reduce it accordingly. Let A be the set of initial vertices of paths in $\mathcal{Q}_1$ and B the set of final vertices of paths in $\mathcal{Q}_2$. By Theorem 1.3.4 it suffices to show that no set X of less than k vertices can separate A and B. But X will always miss some $Q_1 \in \mathcal{Q}_1$, some $Q_2 \in \mathcal{Q}_2$ and some $C \in \mathcal{C}$. The union of these contains an A-B-path. □

A necklace N of length n is called *linear* if it is long-jump-free and has no $(1, n)$-jump. In the proof of [25, Lemma 7.2] we encounter another obstacle: we want to find a set of paths which traverses the beads in an orderly manner. To accomplish this, they choose these paths within a minimal subgraph. We cannot proceed in this way, but we can still make sure the paths do not behave too badly via inductive applications of Lemma 2.6.7.

Lemma 2.6.8. *Let m, n, a, b, s, t be natural numbers with $m \gg n$, $a \gg s$, $b - a \gg n$ and $m - b \gg s$ and let N be a linear weak $(0, 0, l, m)$-necklace. Furthermore let $X \subseteq B_1$ and $Y \subseteq B_m$. Suppose $\mathcal{P} - Z$ contains s disjoint X-B_a-paths, s disjoint B_b-Y-paths and t disjoint B_1-B_m-paths. Then G contains a $(t, 0, l, n)$-necklace N' supported by N with $Z' = Z$ such that B'_1 contains s disjoint X-$B'_1 \cap B'_2$-paths and B'_n contains s disjoint Y-$B'_{n-1} \cap B'_n$-paths.*

Proof. The case $t \leq 1$ is trivial, so we may assume $t \geq 2$. Let $\mathcal{P}' = \mathcal{P} - Z$.

Let $\mathcal{Q}_1$ be a set of s disjoint X-B_a-paths in $\mathcal{P}'$, $\mathcal{Q}_2$ be a set of s disjoint B_b-Y-paths in $\mathcal{P}'$ and $\mathcal{Q}^*$ a set of t disjoint B_1-B_m-paths in $\mathcal{P}'$. Since N is linear we may assume that $\mathcal{Q}_1$ does not meet any B_i with $i > a$ and $\mathcal{Q}_2$ does not meet any B_i with $i < b$. Let $\mathcal{Q}_1^*$ be a subset of $\mathcal{Q}^*$ of size s and let $\mathcal{Q}_2^*$ be its complement. Let H be the union of all paths in $\mathcal{Q}_1$, $\mathcal{Q}_2$ and $\mathcal{Q}_1^*$ together with the beads $B_1, \ldots, B_a$ and $B_b, \ldots, B_m$.

Since every path of $\mathcal{Q}_1$ and $\mathcal{Q}_1^*$ meets every B_i for $i \leq a$, by Lemma 2.6.7 there exists a set $\mathcal{Q}'$ of s disjoint X-B_m-paths $\mathcal{Q}'$ in H. Because N is linear, every path of $\mathcal{Q}'$ meets every B_i for $i \geq b$. Since the same is true for every path in $\mathcal{Q}_2$, by Lemma 2.6.7 there exists a set $\mathcal{Q}''$ of s disjoint X-Y-paths in H. Let $a' = a + 2$ and $b' = b - 2$. Then the linearity of N implies that each path of $\mathcal{Q}_2^*$ contains a $B_{a'}$-$B_{b'}$-path avoiding H. Let $\mathcal{Q}$ be obtained by adding these paths to $\mathcal{Q}''$. Let X' be the set of endvertices of paths of $\mathcal{Q}$ in $X \cup B_{a'}$ and Y' its set of endvertices in $Y \cup B_{b'}$. Then $\mathcal{Q}$ is a set of t disjoint X'-Y'-paths in $\mathcal{P}'$.

Given an X'-Y'-path P we write $L_i(P)$ for the segment of P up to its first point in B_i and $R_i(P)$ for the segment of P from this point.

Let c be a natural number with $b-a \gg c \gg t$. Fix a function $\alpha : \{1, \ldots, n\} \to \mathbb{N}$ with $\alpha(1) = a'$, $\alpha(n) = b'$ and $\alpha(i) - \alpha(i-1) \gg c$ for all $2 \leq i \leq n$. We will recursively construct a set $\mathcal{Q}'$ of t disjoint X'-Y'-paths such that $R_{\alpha(i)}(P) \cap B_{\alpha(i-1)} = \emptyset$ for every $P \in \mathcal{Q}'$ and $2 \leq i \leq n$. We start the construction with $\mathcal{Q}^1 = \mathcal{Q}$. Now we describe how to construct $\mathcal{Q}^i$ for $2 \leq i \leq n$ assuming $\mathcal{Q}_{i-1}$ is already defined. Let $\mathcal{Q}_1$ consist of $L_i(P)$ for every P in $\mathcal{Q}^{i-1}$. Let $\mathcal{Q}_2$ consist of final segments of each P in $\mathcal{Q}^{i-1}$ starting at its last point in $B_{\alpha(i)-c}$. Applying Lemma 2.6.7 to $\mathcal{Q}_1$, $\mathcal{Q}_2$ and the beads from $B_{\alpha(i)-c}$ to B_i (in the path space defined by all these) gives a set of t disjoint X'-Y'-paths, which we call $\mathcal{Q}^i$. Every path P of $\mathcal{Q}^i$ satisfies $R_{\alpha(j)}(P) \cap B_{\alpha(j-1)} = \emptyset$ for $j \leq i$. Indeed, for $j < i$ this is by induction, since the beads from $B_{\alpha(i)-c}$ to B_i do not meet those before $B_{\alpha(j)}$ and for $j = i$ this is implied by the construction, since $\alpha(i) - c \gg \alpha(i-1)$. Thus $\mathcal{Q}' = \mathcal{Q}^n$ is as desired.

As a corollary we get $R_{\alpha(i)}(P) \cap B_{\alpha(j)} = \emptyset$ for $P \in \mathcal{Q}'$ and $1 \leq j < i \leq n$, since $\alpha(j) \leq \alpha(i-1)$.

For $1 \leq i \leq n$ define $L^i = L_{\alpha(i)}[\mathcal{Q}']$ and R^i similarly. Now we set $B_1' = B_{\alpha(1)} \cup L^2$, $B_n' = B_{\alpha(n)} \cup R^n$ and $B_i' = B_{\alpha(i)} \cup (L^{i+1} \cap R^i)$ for $2 \leq i \leq n-1$ and claim that the necklace N' with the B_i' as beads and $Z' = Z$ as its set of hubs is as required.

Indeed, the fourth and fifth property follow from the definition of L^i and the claim proven above, the sixth condition as well as the extra conditions required in the statement are witnessed by segments of the $\mathcal{P}_i$ and all other conditions are immediate.

$\square$

Let N be a necklace of length n and let $1 \leq c \leq d \leq n$. Then a $[c, d]$-*separation* (C, D) of $\mathcal{P}$ is one with $C \cap D = B_c \cup B_d \cup Z$, $B_i \subset C$ for $i \in [c, d]$ and $B_i \subset D$ for $i \in [1, c] \cup [d, n]$. Note that, if N is long-jump-free it has such a separation for all c, d as above. [25, Lemma 7.3] once again requires no major changes.

Lemma 2.6.9. *Let $m \gg n$ be natural numbers and let N be a nontrivial long-jump-free necklace of rank r and length m. If there exist $r+1$ disjoint B_a-B_b-paths for all $a, b \in \mathbb{Z}_n$, then there exists a nontrivial necklace of rank $r+1$ and length n supported by N.*

Proof. Let N be a (t, s, l, m)-necklace as in the statement and let c and d be natural numbers with $c \gg s$, $d - c \gg s + n$ and $m - d \gg s + n$. By assumption there are $r + 1$ disjoint B_c-B_d-paths, in particular there is a set $\mathcal{Q}$ of $t + s + 1$ such paths avoiding Z. Let (C, D) be a $[c, d]$-separation of N. Then each $Q \in \mathcal{Q}$ is contained in either C or D, let $\mathcal{Q}_1$ be the set of those contained in C and $\mathcal{Q}_2$ the set of those contained in D.

First we consider the case where $\mathcal{Q}_1$ has size greater than t. Let N^* be the linear weak necklace of rank l with bead sequence $B_c, \dots, B_d$ and hub set Z. Furthermore let $X = B_{c-1} \cap B_c$, $Y = B_d \cap B_{d+1}$. Applying Lemma 2.6.8 to these gives a $(t+1, 0, l, n)$-necklace in C supported by N^* and thus by N with $Z' = Z$ such that B'_1 contains t disjoint X-$B'_1 \cap B'_n$-paths and B'_n contains t disjoint Y-$B'_{n-1} \cap B'_n$-paths. By adding s disjoint $B_1 \cap B_n$-X-paths through N completely contained in D to B'_1 and proceeding analogously for B'_n we can turn N' into a $(t+1, s, l, n)$-necklace supported by N as required.

Otherwise $\mathcal{Q}_2$ has size greater than s. We may assume that $s < t$, otherwise we are in the first case after reindexing N. Then let N^* be the linear weak necklace of rank l with bead sequence $B_d, \dots, B_m, B_1, \dots, B_c$, and hub set Z. Furthermore let $X = B_{d-1} \cap B_d$ and $Y = B_c \cap B_{c+1}$. Applying Lemma 2.6.8 to these gives a $(s+1, 0, l, 2)$-necklace supported by N^* and thus by N with $Z' = Z$ such that B'_1 contains $s+1$ disjoint X-$B'_1 \cap B'_2$-paths and B'_2 contains $s+1$ disjoint Y-$B'_2 \cap B'_1$-paths. Then the necklace with bead sequence $B'_1, B_{c+1}, \dots, B_{d-1}, B''_2$ and hub set Z can be turned into a $(t, s+1, l, n)$-necklace supported by N by contraction. $\square$

Since [25, Theorem 7.4] only combines all the prior lemmas, our proof can do the same.

Theorem 2.6.10. *Let θ, m and n be natural numbers with $m \gg n$ and $m \gg \theta$. If $\mathcal{P}$ has a θ-connected set U of size m, then $\mathcal{P}$ contains a necklace of rank θ and length n supported by U.*

Proof. Let $g : \mathbb{N} \to \mathbb{N}$ be a function with $g(i) \gg g(i+1)$ for $1 \leq i \leq \theta - 1$, $g(i) = n$ for $i \geq \theta$ and $m \gg g(1)$. By Lemma 2.6.3 we may assume that $\mathcal{P}$ contains a $(1, 0, 0, g(1))$-necklace N such that every bead contains at least θ elements of U, otherwise we are done.

We will now prove the theorem by constructing for $1 \leq i \leq \theta$ nontrivial necklaces N^i of rank i and length $g(i)$ supported by N. Clearly $N^1 = N$ works, so suppose $i > 1$ and that N_{i-1} is already defined. By Corollary 2.6.6 we may assume that $\mathcal{P}$ contains a nontrivial long-jump-free necklace N^* of rank $i - 1$ and length $m' \gg g(i)$, otherwise we have already found our required necklace N^i. Since every bead of N^* contains θ vertices of U, there are θ disjoint paths between any two beads of N^*. Lemma 2.6.9 now gives us the required necklace N^i. □

2.7 Grid theorem

While in [25] Theorem 2.6.10 is used to deduce a generalization of the grid theorem, for our purposes it is enough to deduce the usual one.

Theorem 2.7.1. *Let $m \gg k$ be natural numbers and $\mathcal{P}$ be a path space of tree width m. Then $\mathcal{P}$ has has an $R_{k,k}$-minor.*

Proof. Let q be a natural number with $m \gg q \gg k$. By Proposition 2.5.5 $\mathcal{P}$ has a q-connected set of size at least q. Let r and n be a natural numbers with $q \gg n \gg r \gg k$. By Theorem 2.6.10 there exists a (t, s, l, n)-necklace N of rank r.

If $l \gg k$, then we can even embed a $TR_{k,k}$ with branch vertices only in Z and each path using its own bead. Since $s \leq t$, we may thus assume $t \gg k$.

For every $2 \leq i \leq n-1$ we fix t disjoint B_{i-1}-B_{i+1}-paths in B_i and enumerate them as $P_1^i, \ldots, P_t^i$ in such a way that P_x^{i-1} always meets P_x^i and thus the P_x^i for all the i form a path P_x together. For every such i let T_i be the graph on $\{1, \ldots, k\}$ with an edge between j and j' if there is a P_j^i-$P_{j'}^i$-path in B_i which does not meet any other $P_{j^*}^i$. If there are $k' \gg k$ indices i such that T_i contains a path of length k, then there are $k'' \gg k$ such indices for which these graphs contain the same path Q of length k and such that their corresponding beads are disjoint. We may partition this set of indices into sets $S_1, \ldots, S_k$ of size at least k, each associated with a different edge of Q, by putting every k-th index in order into each of them. When constructing our minor the i-th row of the grid will then correspond to P_l, where l is the i-th vertex of Q. When constructing the $IR_{k,k}$ we thus take the branch sets as suitable segments of the corresponding P_l sets and embed a path witnessing the edge of Q corresponding to S_i into every bead with an index in S_i.

Thus we may assume that there are $k' \gg t$ indices i for which T_i contains a vertex of degree at least k. Then there are some $k'' \gg k$ such indices for which this is the same vertex r, the neighborhood contains the same k vertices $j_1, \ldots, j_k$ and the beads associated to these indices are disjoint. We can embed an $IR_{k,k}$ by taking appropriate segments of the P_{j_i} for $1 \leq i \leq k$ as branch sets for the vertices of the i-th row and obtaining the disjoint paths, between segments of P_{j_y} and P_{j_z}, say, by taking a path witnessing the edge from j_y to r in one of the beads corresponding to our k'' indices and a path witnessing the edge from r to j_z in the next of these beads and connecting them via a segment of P_r. $\square$

As for graphs we immediately get another helpful result.

Corollary 2.7.2. *Let G be a finite graph. Then the class of path spaces with no IG has bounded tree width if and only if G is planar.*

Furthermore we now obtain the qualitative version of Conjecture 2.5.4 promised above.

Corollary 2.7.3. *Let $l \gg k$ be a natural number and let $\mathcal{P}$ be a path space such that every finite minor of $\mathcal{P}$ has tree width at most k. Then $\mathcal{P}$ has tree width at most l.*

2.8 Erdős-Pósa property

In this section we apply the grid theorem once more, following [46, Section 8] to obtain a result on the Erdős-Pósa property.

We say that a path space $\mathcal{Q}$ has the *EP* if for natural numbers k and l with $l \gg k$ any path space $\mathcal{P}$ contains k disjoint $I\mathcal{Q}$ or has a set of l points meeting every $I\mathcal{Q}$.

The intersection graph of a family $(X_i)_{i\in I}$ of sets is the graph I with an edge between $i \neq j$ if X_i and X_j meet.

Let us first give an outline of the proof for graphs. They start with a large graph G which does not contain k disjoint embeddings of some planar H. By the grid theorem G then has bounded tree width and they take a tree decomposition, say with tree T, witnessing this. For each component C_i of H, they then look at the set $\mathcal{A}_i$, which consists of the subtree of all parts meeting some embedding of C_i for every embedding of C_i. They then make use of the fact that intersection graphs of subtrees are chordal and prove their main lemma ([46, (8.5)]) in the language of chordal graphs.

Lemma 2.8.1. *Let G be a chordal graph and let $Z_1, \dots, Z_m$ be disjoint independent sets of G each of size k. Let $x_1, \dots, x_m$ be nonnegative integers with $x_1 + \cdots + x_m = k$. Then there exists an independent subset X of $V(G)$ which for every $1 \leq i \leq m$ satisfies $|X \cap Z_i| = x_i$.*

Translated back to subtrees this means that if every $\mathcal{A}_i$ had many disjoint elements, we could find a set of k elements from each $\mathcal{A}_i$ which are all disjoint. This would then give k disjoint embeddings of H, a contradiction. So one may assume that some $\mathcal{A}_i$ has a bounded number of elements and thus by an easy fact for trees a bounded set of vertices of T meets every element. The union of all the corresponding parts is then their desired bounded set meeting every embedding of C_i and thus H.

When translating this proof to our setting, the obvious problem is that instead of decomposition trees we have tree-like spaces and the results they cite are not known for these. Thus we will have to check them ourselves.

Given a tree-like decomposition $D = (T, (V_t)_{t\in V(T)})$ of $\mathcal{P}$ and a subset $X \subseteq V(\mathcal{P})$ we write $D(X)$ for the subgraph-like space of T induced by those parts meeting X and call it the *trace of X in D*. First we should check that the traces we need form subtree-like spaces.

Lemma 2.8.2. *Let D be a tree-like decomposition of $\mathcal{P}$ and $X \subseteq V(\mathcal{P})$ connected and simple. Then $D(X)$ is a subtree-like space.*

Proof. Let $D = (T, (V_t)_{t\in V(T)})$. Clearly $D(X)$ is connected. Since T is compact, it is enough to show that $D(X)$ is closed. We may assume without loss of generality that X is a path, since a finite union of closed sets is again closed. Thus we need to show that for $t \in T \setminus D(X)$ there is an open neighborhood of t not meeting $D(X)$. This is easy if t is an inner point of an edge, so we may assume that t is a vertex. Now it suffices to show that some edge e separates t from $D(X)$, since then adding an open interval of e to the component of $T - e$ containing t is a neighborhood as desired.

Let L be some pseudo-line in T from some $s \in D(X)$ to t. We may assume that there is a cofinal sequence of points of $D(X)$ before t, otherwise we can easily find an edge as desired. Then we can find a cofinal subsequence $(d_\alpha)_{\alpha<\beta}$ of this sequence and for each d_α a choice of one $x_\alpha \in V_{d_\alpha} \cap X$ such that $(x_\alpha)_{\alpha<\beta}$ is ascending or descending. Let $x \in X$ be the corresponding supremum or infimum respectively (with respect to X). To complete the proof we will show that this implies that $x \in V_t$, contradicting $t \notin D(X)$. Let $(A, B) \in t$. Then it is either also contained in a cofinal subsequence of the d_α and thus B must contain a cofinal subsequence of the x_α. By Lemma 0.4.4 we must then also have $x \in B$ and since (A, B) was arbitrary, $x \in V_t$. $\square$

We next prove that the relevant intersection graphs are chordal.

Lemma 2.8.3. *Let T be a tree-like space. The intersection graph of a finite family of subtree-like spaces of T is chordal.*

Proof. Assume otherwise and let $T_1, \dots, T_n$ represent an induced cycle of length at least four. Since T_2 meets both T_1 and T_3, but these do not meet each other, there is an edge e of T_2 separating T_1 and T_3. Since every T_k for $4 \le k \le n$ meets T_{k-1}, but not T_2 by induction every such T_k is contained in the component of $T - e$ containing T_3. In particular T_n does not meet T_1 contradicting our assumption. $\square$

Finally, to show that if there are no k disjoint subtree-like spaces contained in a set, we can find less than k vertices meeting them all, we need a bit of preparation. First note that the intersection of any two subtree-like spaces is either empty or again a subtree-like space. As for graphs (see for example [14]), it is useful to first show that any family with nonempty pairwise intersections has nonempty intersection.

Lemma 2.8.4. *Let T be a tree-like space and $\mathcal{T}$ a set of subtree-like spaces which pairwise meet. Then $\mathcal{T}$ has nonempty intersection.*

Proof. Since T is compact and the $\mathcal{T}$ are closed in T, it is enough to check that the intersection of any finite subset of $\mathcal{T}$ is nonempty. Indeed, we will prove this by induction for any finite set $\mathcal{X}$ of subtree-like spaces of T which pairwise meet.

If $\mathcal{X}$ has at most two elements, this is trivial, so we first consider the case where $\mathcal{X}$ has three elements X_1, X_2, X_3. Assume for a contradiction that $X_1 \cap X_2 \cap X_3$ is empty. Then there exists some edge e separating the subtree-like spaces $X_1 \cap X_2$ and X_3. Without loss of generality X_1 does not contain e. But by assumption it must meet both components of $T - e$, a contradiction.

Now assume $\mathcal{X}$ has elements $X_1, \ldots, X_k$ for $k \geq 4$. Let $X_i' = X_i \cap X_k$ for $1 \leq i \leq k-1$. The X_i' then pairwise meet by the induction hypothesis for 3. By the induction hypothesis for $k-1$, we obtain the desired result. $\square$

For the next proof we require the fact that complements of chordal graphs are perfect, which is a consequence of the weak perfect graph theorem of [39].

Lemma 2.8.5. *Let T be a tree-like space and $\mathcal{T}$ a set of subtree-like spaces. Then there are k disjoint elements of $\mathcal{T}$ or a set of less than k vertices meeting every element of $\mathcal{T}$.*

Proof. Let G be the complement of the intersection graph of $\mathcal{T}$. If G can be colored with $c < k$ colors, this gives a decomposition of T into classes $\mathcal{T}_1, \ldots, \mathcal{T}_c$ such that any two elements of each class meet. By Lemma 2.8.4 there are then $x_1, \ldots, x_c$ with x_i in the intersection of $\mathcal{T}_i$. These $c < k$ points then meet every element of $\mathcal{T}$.

Thus we may assume that G cannot be colored by less than k colors. By compactness, G then has a finite subgraph H which requires k colors. But by Lemma 2.8.3 H is a complement of a chordal graph and thus perfect. In particular, H has a clique of size k. The k subtree-like spaces representing this clique are then disjoint, as required. $\square$

Now we can follow the outline given in the beginning.

Theorem 2.8.6. *Let G be a finite planar graph. Then IG has the EP.*

Proof. W.l.o.g Let $C_1, \ldots, C_m$ be the components of G. Let $l \gg n \gg k + m$ be natural numbers and let $\mathcal{P}$ be a path space. We need to show that $\mathcal{P}$ contains k disjoint IG or l points meeting all IG. We may assume that $\mathcal{P}$ has a tree-like decomposition D of width at most m, since otherwise by Corollary 2.7.2 it would contain k disjoint IG. Let $\mathcal{A}_i$ be the set of all IC_i in G.

By Lemmas 2.8.2 and 2.8.5 for each $1 \leq i \leq m$ there are either $A_1^i, \ldots A_n^i \in \mathcal{A}_i$ with disjoint traces or a set X_i of size less than n meeting all traces of elements of $\mathcal{A}_i$.

If for some i we get the second case, let Y be the union of all the parts corresponding to X_i. Then Y meets every IC_i in $\mathcal{P}$ and thus also every IG. Since furthermore $|Y| \leq l$, we are done.

If for every i we get the first case, then let H be the intersection graph of the $(T(A^i_j))_{1 \leq i \leq m, 1 \leq j \leq n}$. By Lemma 2.8.3 H is chordal, so by Lemma 2.8.1 there exists an independent set I which contains k elements corresponding to traces of elements of $\mathcal{A}_i$ for every $1 \leq i \leq k$. Since sets with disjoint traces are disjoint, this implies that there are disjoint $(B^i_j)_{1 \leq i \leq m, 1 \leq j \leq k}$ with B^i_j an IC_i for all i, j. Combining these gives k disjoint IG, as desired. $\square$

2.9 Outlook

Having provided decompositions along separations of order one and of order two, a decomposition along 3-separations would be a natural next step. For graphs [29] provides such a decomposition, which seems like it could plausibly be translated to path spaces. This could also help resolve Conjecture 2.5.4 for $k = 4$. However, it seems less clear how one could resolve the conjecture for higher values of k or even in the general case.

For graphs, the theory of tree decompositions is strongly linked to the theory of well-quasi orderings. Unfortunately this tool seems unhelpful for path spaces. Because it is possible to construct a pseudo-line from any linear order, not even the class of paths is well-quasi ordered by the minor relation.

Chapter 3

Limit-closed separation systems

3.1 Introduction

Tangles were first introduced by Robertson and Seymour in [47] as a witness for high-tree width. One of the important ingredients in their graph minor project was a tangle-tree theorem, that is a theorem giving a tree decomposition separating all the tangles of a graph.

As it turns out these features of tangles, both being an obstruction and admitting a tangle-tree theorem, can formulated more abstractly. This is the core of the theory of abstract separation systems formulated in [13]. In this setting tangle-tree theorems are now usually formulated more generally in terms of profiles, as seen for instance in [15]. While that article gives a very general such theorem for finite separation systems, for infinite separation systems there is no general tangle-tree theorem.

When searching for a tangle-tree theorem for path spaces, we must thus look to some special property of the set of separations of a path space. Unfortunately this set is neither the inverse limit of finite minors, as would be needed for the tools of [17], nor does every separation cross only finitely many others as would be needed for [20].

One property which the separation system obtained from a path space does have, however, is closure under limits. For separation systems with this property it is possible to use maximal separations as in the finite case.

In Section 3.2 we use this to find a tree set distinguishing profiles of an abstract separation system using a two step process: first we consider a set of equivalence classes of separations and find a tree set of such equivalence classes and then choose representatives for these classes. Section 3.3 then applies this result inductively to get a full tree of tangles for a separation system with an order function.

To further understand these limit-closed separation systems we want to translate the theory of flowers to this setting. First introduced in [42] specifically for the case of 3-separations of matroids, flowers give a way to further organize the structure of crossing separations. They were then generalized to arbitrary orders in [3] and [10] then related them to tangles and gave a tangle-tree theorem, in which all the nodes not containing tangles would display flowers.

In Section 3.4 we give our definition of flowers. Since our separations are not bipartitions and we need to allow for infinite index sets our definition ends up more complex than the ones from the literature. We also define a weakening, pseudoflowers, which will be necessary as a limit object.

Section 3.5 shows that if there is a separation which interacts badly with a pseudoflower, we can extend the pseudoflower to remedy this. Given a chain of pseudoflowers Section 3.6 then constructs a limit pseudoflower. Our starting point for this is an inverse limit, which we then carefully repair to make it a pseudoflower. The existence of these limits then implies the existence of maximal pseudoflowers. Section 3.7 gives a few short notes on how these maximal pseudoflowers interact with the tree sets constructed in the beginning of the chapter.

3.2 Equivalence classes

Given a separation system S in some universe, a *profile* (as defined for example in [15]) is a consistent orientation P of S such that for $s, t \in P$ we do not have $(r \vee s)^* \in P$.

Let S be a separation system in some universe U and $\mathcal{P}$ a set of regular profiles of S. Then we define a map $f_{\mathcal{P}} : S \to 2^{\mathcal{P}}$ via $f_{\mathcal{P}}(s) = \{P \in \mathcal{P} | s^* \in P\}$. We consider $2^{\mathcal{P}}$ as a separation system with inclusion as the order.

Lemma 3.2.1. *The map $f_{\mathcal{P}}$ is a homomorphism of separation systems which respects $\vee$ and $\wedge$.*

Proof. Since profiles orient every separation, $f_{\mathcal{P}}$ respects the involution. Furthermore if $s \leq t$ and $s^* \in P$ for some profile P, then we cannot have $t \in P$ by consistency, so $f_{\mathcal{P}}(s) \leq f_{\mathcal{P}}(t)$. Now it suffices to prove that for $s, t \in S$ with $s \vee t \in S$ an arbitrary $P \in \mathcal{P}$ contains $s \vee t$ if and only if it contains both s and t. The forward implication follows by consistency, the backwards one by the profile property. $\square$

The fibers of $f_{\mathcal{P}}$ are exactly the equivalence classes obtained by regarding two separations as equivalent if they are oriented the same way by every profile in $\mathcal{P}$. Assuming a structurally submodular S comparing these classes via their images under $f_{\mathcal{P}}$ gives the same order as comparing them via their elements. In fact, a slightly weaker condition suffices. Call S *weakly $\mathcal{P}$-submodular (with respect to $\mathcal{P}$)* if whenever $f_{\mathcal{P}}(s) \leq f_{\mathcal{P}}(t)$ at least one of $s \vee t$ and $s \wedge t$ is contained in S.

Proposition 3.2.2. *Let S be weakly $\mathcal{P}$-submodular with respect to $\mathcal{P}$. Then for $A, B \in \operatorname{im}(f_{\mathcal{P}})$ we have $A \leq B$ if and only if there are $a \in f_{\mathcal{P}}^{-1}(A), b \in f_{\mathcal{P}}^{-1}(B)$ with $a \leq b$.*

Proof. The backward direction is immediate by Lemma 3.2.1. For the forward direction choose $a \in f_{\mathcal{P}}^{-1}(A), b \in f_{\mathcal{P}}^{-1}(B)$. At least one of $a \wedge b$ and $a \vee b$ is contained in S by weak submodularity. If $a \wedge b \in S$, then by Lemma 3.2.1 it is contained in $f_{\mathcal{P}}^{-1}(B)$ and thus forms the required pair together with b. Similarly, if $a \vee b \in S$, then $a \vee b \in f_{\mathcal{P}}^{-1}(A)$ and it forms the required pair with a. $\square$

We want to use the transformation $f_{\mathcal{P}}$ to find a distinguishing set for $\mathcal{P}$. We will proceed in two steps: First we will look for an abstract distinguishing set in the image and then look for separations of S to represent them.

Let us start by stating our objective for the first step more formally. We will say that $A \in \operatorname{im}(f_{\mathcal{P}})$ *separates* $P, Q \in \mathcal{P}$ if $P \in A$ and $Q \notin A$ or vice versa. Then we are looking for some tree set $T \subseteq \operatorname{im}(f_{\mathcal{P}})$ such that different elements of $\mathcal{P}$ are always separated by some element of $\operatorname{im}(f_{\mathcal{P}})$. Since structural submodularity of S translates to $\operatorname{im}(f_{\mathcal{P}})$, for finite $\mathcal{P}$ standard techniques easily show that this condition is enough to reach our goal. If $\mathcal{P}$ is infinite these standard methods are not sufficient, but there is a useful idea, seen for instance in [2], which may help. That idea is taking only the *good* separations, that is those not crossed by any other separation, for our tree set.

In the following we will show that, given certain conditions, the set $T(S, \mathcal{P})$ of good separations of $\operatorname{im}(f_{\mathcal{P}})$ (except $\emptyset$ and $\mathcal{P}$) does indeed meet our demands.

Since $T(S, \mathcal{P})$ is a tree set by definition, all that needs to be shown is that $T(S, \mathcal{P})$ separates any two distinct elements of $\mathcal{P}$. When dealing with finite separation systems, it is sometimes useful to consider maximal separations. If we want to use this trick in the infinite case, we encounter some difficulties. First of all, profiles may not even have maximal elements, which would render our strategy impossible. Thus we require our profiles to be *closed*, meaning that any chain in the profile has a supremum in the universe which is contained in the profile. This ensures that each profile has a maximal element, but even these maximal elements may not have the nice properties which we are used to, say when we have an order function.

Thus we need one more condition, which emulates some of the additional structure provided by an order function. We call S *orderly* (with respect to $\mathcal{P}$) if for any $s, t \in S$ such that both $\{s, t\}$ and $\{s^*, t^*\}$ are subsets of (possibly different) elements of $\mathcal{P}$ we have $s \vee t \in S$ and $s \wedge t \in S$.

To give an example, it will follow from Lemma 3.3.2 that the set of proper k-separations in a k-connected path space is orderly.

Lemma 3.2.3. *If S is orderly and every $P \in \mathcal{P}$ is closed, all different $P, Q \in \mathcal{P}$ are separated by some $t \in T(S, \mathcal{P})$.*

Proof. Let C be a chain in S such that each element of C is contained in P and not in Q which is maximal with these properties. Since P is closed, C has a supremum s, which is contained in P. Furthermore, by consistency, we have $s \notin Q$. Now it suffices to prove that $f_{\mathcal{P}}(s)$ is good. If not, there exists some $x \in S$ such that the images of s and x cross. But since S is orderly, this would imply $s \vee x \in S$ and this separation could have been added to C. □

Now let us consider the second step. Starting with a regular tree set T in $\mathrm{im}(f_{\mathcal{P}})$, we want to find a nested set of preimages (or equivalently one isomorphic to T). Once again we want to use maximal separations and so need the same conditions as before. Closure guarantees that each equivalence class has a maximal separation and orderliness that it is even greatest.

Lemma 3.2.4. *If every $P \in \mathcal{P}$ is closed, for every $t \in T$ the set $f_{\mathcal{P}}^{-1}(t)$ has a greatest element.*

Proof. Let $X = f_{\mathcal{P}}^{-1}(t)$ and let C be a nonempty chain in X. Since $\mathcal{P} \setminus t$ is nonempty and each profile in that set is closed, C has a supremum s with $f_{\mathcal{P}}(s) \supseteq t$. Conversely there cannot be any $Q \in t$ with $s \in Q$, since we would then have $C \subseteq Q$ by consistency, a contradiction. Thus $s \in X$. This implies that any element of X lies below some maximal element.

Now let a and b be two maximal elements of X. Since S is orderly, $a \vee b \in S$ and then also $a \vee b \in X$. By maximality of a and b these must then both be equal to $a \vee b$ and we have $a = b$. Thus X only has a single maximal element which must then be greater than all other separations. □

Let m be the function mapping each $t \in T$ to the greatest element of $f_{\mathcal{P}}^{-1}(t)$. We would like to use m to choose the representatives, however, this is not quite possible, since for $t \in T$ the separations $m(t)$ and $m(t^*)$ are usually not inverses. Fortunately, this is no great obstacle. If we simply choose one of these two possible unoriented separations, the only thing that could go wrong is that for $s \leq t$ we choose $m(s)$ and $m(t^*)^*$ as the representing separations. It thus suffices to choose exactly opposite to a consistent orientation, which always exists. So we fix a consistent orientation o of T and define the function m_o by mapping $s \in T$ to $m(s^*)^*$ for $s \in o$ and to $m(s)$ for $s \notin o$. Let $\hat{T}$ be the image of T under m_o.

Corollary 3.2.5. *Let S be orderly and every $P \in \mathcal{P}$ closed. Then $f_{\mathcal{P}}$ restricts to an isomorphism between $\hat{T}$ and T.*

Proof. Clearly, m_o is the inverse map of $f_{\mathcal{P}}$ restricted to the image of m_o. Since $f_{\mathcal{P}}$ is a homomorphism by Lemma 3.2.1, it suffices to show that m_o is, too. So let $s, t \in T$ be such that $s \leq t$. We need to show that $m_o(s) \leq m_o(t)$. If $t \notin o$, we have $m_o(t) = m(t)$. Since S is orderly, we have $m(t) \vee m_o(s) \in S$. By Lemma 3.2.1 $f_{\mathcal{P}}(m(t) \vee m_o(s)) = f_{\mathcal{P}}(m(t)) \cup f_{\mathcal{P}}(m_o(s)) = t \cup s = t$. But since $m(t)$ is a greatest element, we must have $m_o(s) \leq m(t) = m_o(t)$. So we may assume $t \in o$ and by consistency also $s \in o$. Then we have $m_o(s) = m(s^*)^*$ and $m_o(t) = m(t^*)^*$. Since S is orderly, we have $m(s^*) \vee m(t^*) \in S$ and calculating with Lemma 3.2.1 as before we get $f_{\mathcal{P}}(m(s^*) \vee m(t^*)) = s^*$. Again by choice of $m(s^*)$ we must have $m(s^*) \geq m(t^*)$ and thus $m_o(s) \leq m_o(t)$. □

This completes the second step. Overall, we have now proven the following theorem.

Theorem 3.2.6. *Let S be a regular separation system orderly with respect to a set $\mathcal{P}$ of closed profiles. Then there is a tree set T with the following properties:*

1. *Any two different elements $\mathcal{P}$ are distinguished by some element of T.*

2. *Any element of T distinguishes some elements of $\mathcal{P}$.*

3. *In the set $P \cap T$ for every $P \in \mathcal{P}$ every separation lies below some maximal separation.*

3.3 Tangle-tree theorem

Since path space separations form a submodular universe, we would like build on this work to take into account the order function.

If U is a submodular universe, a *k-profile* is a profile of the separation system of all separations of U with order less than k. We say that such a k-profile P is *robust* if for every $s \in P$ and $t \in U$ such that $s^* \wedge t$ and $s^* \wedge t^*$ both have order less than $|s|$, at least one of the two is not contained in P. When we refer to a profile (of U) in this context we always mean a k-profile of U for some $k \in \mathbb{N}$.

The goal of this section is to prove the following theorem.

Theorem 3.3.1. *Let U be a submodular universe and let $\mathcal{P}$ be a collection of regular robust closed profiles. Then there is a tree set T that efficiently distinguishes $\mathcal{P}$ such that every $t \in T$ efficiently distinguishes two profiles in $\mathcal{P}$.*

Note that the tree set T constructed in the proof of Theorem 3.3.1 additionally has the property that for every $P \in \mathcal{P}$, every element of $T \cap P$ is less than or equal to a maximal element of $T \cap P$.

We are going to recursively construct tree sets T_k whose union then efficiently distinguishes $\mathcal{P}$. To do this, let $\mathcal{P}_k$ be the set of profiles of order at most k that are induced by elements of $\mathcal{P}$. For the construction of a tree set T_{k+1} distinguishing the elements of $\mathcal{P}_{k+1}$ we will employ the following proof strategy in two steps, that is common in proofs of tree-of-tangles-theorems, for example in [15]. First, for a k-profile $P \in \mathcal{P}_k$, Theorem 3.2.6 is applied to the set of all $k+1$-profiles in $\mathcal{P}_{k+1}$ whose induced k-profile is P and to a carefully chosen subuniverse of U to obtain a tree set T_P. Second, it will be shown that T_k together with all the tree sets T_P is a tree set that efficiently distinguishes $\mathcal{P}_{k+1}$.

More precisely, we are going to show by induction that for every $k \in \mathbb{N}$ there is a tree set T_k with the following properties:

- T_k is a subset of S_k.
- Every element of T_k distinguishes two profiles in $\mathcal{P}_k$ efficiently.
- T_k distinguishes $\mathcal{P}_k$ efficiently.

- For every $Q \in \mathcal{P}_k$, every element of $Q \cap T_k$ is less than or equal to a maximal element of $Q \cap T_k$.

As the only 0-profile is the empty set we define T_0 to be the empty set. Assume that T_k is already defined and we now want to define T_{k+1}. For each k-profile $P \in \mathcal{P}_k$ let $\mathcal{Q}_P$ be the set of all $k+1$-profiles in $\mathcal{P}_{k+1}$ whose induced k-profile is P, N_P the set of maximal elements of $P \cap T_k$ and U_P the set of all separations in U towards which all separations of N_P point. Note that U_P is a universe and closed under taking infima and suprema of chains of bounded order, and that as a result the set of k-separations of U_P is structurally submodular and closed under taking suprema of chains of bounded order. Then every profile in $\mathcal{Q}_P$ induces a closed $k+1$-profile of U_P.

We will want to apply the results of the previous section, in particular Theorem 3.2.6, to U_P and the profiles of U_P induced by $\mathcal{Q}_P$. In order to do so, we need to show that the separation system S_k of U_P is orderly with respect to $\mathcal{Q}_P$. That follows from the following, slightly more general statement.

Lemma 3.3.2. *Let U be a submodular universe, k a non-negative integer and $\mathcal{P}$ a set of $k+1$-profiles that all have the same induced k-profile. Then the set of k-separations of U is orderly with respect to $\mathcal{P}$.*

Proof. Assume s and t are k-separations of U and P and Q are elements of $\mathcal{P}$ such that $\{s,t\} \subseteq P$ and $\{s^*,t^*\} \subseteq Q$. By submodularity, one of $s \vee t$ and $s \wedge t$ is also a k-separation, assume without loss of generality that $s \vee t$ is a k-separation. Then $s \vee t$ is also contained in P.

First consider the case that $s \vee t$ is not contained in Q. Then $s \vee t$ distinguishes P and Q, and thus has order exactly k. Hence by submodularity also $s \wedge t$ is a k-separation and we are done.

So consider the case that $s \vee t \in Q$. As a separation system with a degenerate separation does not have a profile, s is not contained in Q. Thus by consistency $s \le s \vee t \in Q$ implies $s = (s \vee t)^*$. Similarly $t = (s \vee t)^*$, so $s = t$ and the lemma holds. □

Now we will show that indeed all profiles in $\mathcal{Q}_P$ can be distinguished by separations of U_P.

Lemma 3.3.3. *All profiles in $\mathcal{Q}_P$ are distinguished by k-separations of U_P.*

Proof. Let $f_{\mathcal{Q}_P}$ be defined as in the previous section. It suffices to show that if r is a separation such that $f_{\mathcal{Q}_P}(r)$ is neither $\emptyset$ nor $\mathcal{Q}_P$, then there is an element

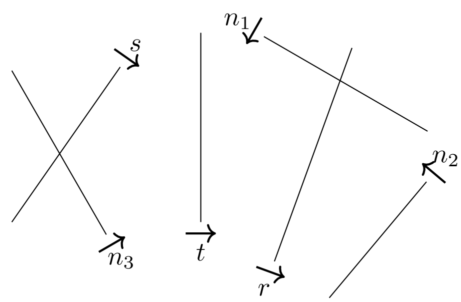

Figure 4: The separation s is less than the separations n_1^* and n_2^*, which is equivalent to $n_1 \leq s^*$ and $n_2 \leq s^*$, and the separation t is additionally bigger than n_3.

of U_P that has the same image under $f_{\mathcal{Q}_P}$ as r. Let X be the set of separations whose image under $f_{\mathcal{Q}_P}$ is $f_{\mathcal{Q}_P}(r)$. Several times in this proof we will use the fact that $f_{\mathcal{Q}_P}$ preserves infima and suprema, and in particular that X is closed under taking suprema and infima.

If N_P is empty, then r itself is a separation of U_P and we are done, so assume otherwise. By Lemma 3.2.4, X has a smallest element s. Let N be the set of all elements n of N_P such that $n \leq s^*$. Define X' to contain all elements $x \in X$ with $n \leq x^*$ for all $n \in N$. Just as in the proof of Lemma 3.2.4, the fact that the profiles in $\mathcal{Q}_P$ are closed and Zorn's Lemma imply that X' has a maximal element t. See Fig. 4 for an illustration.

We will show that all elements of N_P point towards t. In order to do so let n be an element of N_P. If there is $q \in X$ such that $q \leq n^*$, then $s \wedge q \in X$. By minimality of s this implies $s = s \wedge q$, so $s \leq q \leq n^*$ and thus $n \in N'$. Hence n points towards t. The other case is where there is no such separation q, in particular $t \wedge n^*$ is not a candidate for q. If $t \wedge n^*$ is a $\leq k$-separation, then it is contained in X. So the fact that $t \wedge n^*$ is not contained in X implies that it is not a $\leq k$-separation. Because $\mathcal{P}$ is robust and consistent, this implies that $t \vee n$ is a $\leq k$-separation. Then $t \vee n \in X$ and all $n' \in N'$ point towards $t \vee n$, so by maximality of t in X' we have $t \vee n = t$ and hence $n \leq t$. So also in this case n points towards t. □

So in order to distinguish the elements of $\mathcal{Q}_P$, it suffices to distinguish their intersections with U_P. We will do that by applying Theorem 3.2.6 to the k-separations of U_P and the profiles induced by $\mathcal{Q}_P$. Call the obtained tree set T_P. Let T_{k+1} be the union of T_k and all tree sets T_P where P is a k-profile contained in $\mathcal{P}_k$.

Lemma 3.3.4. *T_{k+1} is a tree set that distinguishes $\mathcal{P}_{k+1}$ efficiently, and in which every separation distinguishes some elements of $\mathcal{P}_{k+1}$.*

Proof. Let P be a k-profile in $\mathcal{P}_k$. As for every separation n in $P \cap T_k$ there is an element n' of N_P such that $n \le n'$, T_k is nested with every separation in U_P and thus with T_P. Also, if P and P' are distinct k-profiles in $\mathcal{P}_k$, then they are distinguished by some separation n of T_k, so they are also distinguished by some separation of N_P which then witnesses that every separation in $T_{P'}$ is nested with every separation in T_P. So T_k is nested. Also by construction every element of T_{k+1} distinguishes two elements of $\mathcal{P}_{k+1}$ and thus is neither trivial nor co-trivial nor degenerate.

In order to show that every two elements of $\mathcal{P}_{k+1}$ are distinguished efficiently, let P and Q be two such elements that can be distinguished. If P and Q can be distinguished by a $k-1$-separation, then they induce distinguishable elements of $\mathcal{P}_k$ which are thus efficiently distinguished by some separation t of T_k. Then t also efficiently distinguishes P and Q. So we are left with the case that P and Q cannot be distinguished by a k-separation, which implies that they are both $k+1$-profiles and induce the same k-profile P'. But then P and Q are distinguished by a separation in $T_{P'}$, and that separation distinguishes them efficiently. $\square$

So in order to complete the proof, we only have to prove the following statement.

Lemma 3.3.5. *For every $Q \in \mathcal{P}_{k+1}$, every element of $Q \cap T_{k+1}$ is less than or equal to a maximal element of $Q \cap T_{k+1}$.*

Proof. Let $Q \in \mathcal{P}_{k+1}$ and let $s \in Q \cap T_{k+1}$. If $Q \in \mathcal{P}_k$, then $Q \cap T_{k+1} = Q \cap T_k$ and we are done, so assume otherwise. So Q is a $k+1$-profile, denote its induced k-profile by Q'. If s is not contained in $T_{Q'}$, then it is less than or equal to an element of $N_{Q'}$, and every element of $N_{Q'}$ is either maximal in $T_{k+1} \cap Q$ or less than a separation in $T_{Q'}$. So it suffices to consider the case that $s \in T_{Q'}$. But in this case by Theorem 3.2.6 s is less than or equal to a maximal element of $T_{Q'} \cap Q'$ which then also is a maximal element of $Q \cap T_{k+1}$. $\square$

Proof of Theorem 3.3.1. Let T be the union of all the T_k as defined in this section. Then T is a tree set which distinguishes all distinguishable profiles in $\mathcal{P}$ efficiently, and every separation in it distinguishes two elements of $\mathcal{P}$ efficiently. $\square$

3.4 Pseudoflowers

As mentioned in the introduction, we now want to translate the theory of flowers to path spaces. Specifically the property we will use is that of Proposition 2.2.5. Thus we call a universe of vertex separations *limit-closed* if for any natural number $k > 0$ and chain $((A_i, B_i))_{i \in I}$ of separations of order at most k this chain has a supremum in U which has order at most k and the form $(A \cup X, B)$ where A is the union of the A_i and B is the intersection of the B. Then the universe of path space separations of finite order is clearly a limit-closed universe. For the rest of this chapter we fix any such a limit-closed universe $\mathcal{U}$ with ground set V.

Since flowers are based on cyclic orders, we will first give some background on cyclic orders.

3.4.1 Cyclic orders

We only give a rough outline here, additional details are given in Appendix A.1. Of course, we start with a definition.

Definition 3.4.1. [41] A set of triples Z in $S \times S \times S$ is a *cyclic order* of the set S if it has the following four properties:

- (cyclic) $\forall a, b, c \in S : (a, b, c) \in Z \Rightarrow (b, c, a) \in Z$
- (antisymmetric) $\forall a, b, c \in S : (a, b, c) \in Z \Rightarrow (c, b, a) \notin Z$
- (linear) $\forall a, b, c \in S$ pairwise distinct $: (a, b, c) \notin Z \Rightarrow (c, b, a) \in Z$
- (transitive) $\forall a, b, c, d \in S : (a, b, c) \in Z \wedge (a, c, d) \in Z \Rightarrow (a, b, d) \in Z$.

So formally a cyclic order is a different set than its underlying ground set, but in this paper (except in the appendix) this distinction is not made. Note that what is called a cyclic order here is sometimes (also in [41]) called a linear cyclic order or a complete cyclic order, with a cyclic order not necessarily being linear. The distinction is not made here as all cyclic orders under consideration are linear.

For two elements a and b of S, the set of elements $c \in S$ satisfying $(a, c, b) \in Z$ is denoted by (a, b). The sets $(a, b) + a$, $(a, b) + b$ and $(a, b) + a + b$ are denoted by $[a, b)$, $(a, b]$ and $[a, b]$ respectively.

When necessary to resolve ambiguities, we may add the cyclic order in which these are taken as a subscript. A subset I of S is an *interval* if for all $s, t \in I$

either $[s,t] \subseteq I$ or $[t,s] \subseteq I$. We call an interval *non-trivial* if it is neither the empty set nor all of S. Clearly all subsets of S of the form (a,b), $(a,b]$, $[a,b)$ or $[a,b]$ are intervals of Z. Similarly to linear orders, an element s is the *successor* of a in S if $b \notin [a,s]$ for all for all other elements b. Also, an element p is the *predecessor* of a if $b \notin [p,a]$ for all other elements b. A *neighbor* of a is an element that is a predecessor or successor of a.

The definition of a cyclic order implies that if Z is a cyclic order on a set S and (s,s',t) is an element of Z then s, s' and t are pairwise distinct. ([40, Lemma 1.4])

We will use cyclically ordered sets of a special kind as the index set of our k-pseudoflowers:

Definition 3.4.2. Let I be a cyclically ordered set with at least two elements. A cyclically ordered set C is a *cycle completion* of I if I is a subset of C, the cyclic order on I is the one induced by C, and every non-trivial interval of I can uniquely be written as $[v,w] \cap I$ for elements v, w of $C \setminus I$.

Intuitively, the cycle completion of a cyclic order arises from that cyclic order as follows: Let I be some cyclically ordered set. If one envisions I as boxes arranged in a circle according to the cyclic order (see also Fig. 5), then it is possible to cut up the circle at two places without cutting through boxes, thereby dividing the set of boxes into two intervals. If I is finite, then every one of these "cut points" is between two boxes. So I and the set of possible cut points form together another cyclically ordered set. If I is infinite, then not every cut point is between two boxes, but still I and the set of possible cut points form together a cyclically ordered set, the cycle completion of I.

Cycle completions can be formalised in several ways, and doing so via cuts, as is done in the appendix, is only one possibility of many.

The cycle completion has several properties which we will use throughout the article and which are intuitively clear. Also, although we did not find the exact construction of the cycle completion described in the literature, several very similar constructions are. For these reasons, we will only state the properties needed in this section, and the proofs can be found in the appendix, together with a description of how to obtain the cycle completion from similar constructions.

The first of the facts about cycle completions which we need is that cycle completions exist and are essentially unique. In order to compare different cyclic orders, we will use *monotone* maps, that is maps f between cyclic orders where $f(v) \in (f(u), f(w))$ implies $v \in (u,w)$, and *isomorphisms of cyclic orders*,

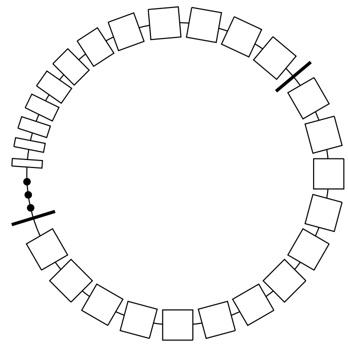

Figure 5: A set (whose elements are indicated by boxes) which is cyclically ordered (indicated by the arrangement of the boxes on a circle) and two “cutting points” dividing the cyclically ordered set into two intervals.

bijective maps f where $f(v) \in (f(u), f(w))$ is equivalent to $v \in (u, w)$. Note that if the image of some monotone map has at least three elements, then the preimage of every interval is also an interval.

Lemma 3.4.3 (see Lemmas A.1.13 and A.1.17)**.** *For every cyclically ordered set I with at least two elements there is a cycle completion. If C and C' are two cycle completions of I, then there is an isomorphism of cyclic orders $F : C \to C'$ such that the restriction of F to I is the identity.*

The fact that the cycle completion of a cyclically ordered set I is unique up to isomorphism justifies calling it “the” cycle completion of I and denoting it as $C(I)$. We will also call the elements of $C(I) \setminus I$ *cutpoints*, a terminology partly inspired by the term cut (see appendix).

Lemma 3.4.4 (see Lemma A.1.14)**.** *For a cyclically ordered set I with at least two elements let v and w be distinct elements of $C(I)$. Then $[v, w] \cap I$ is a non-trivial interval of I.*

Lemma 3.4.5 (see Lemma A.1.17)**.** *Let I and I' be cyclically ordered sets with at least two elements and let $f : I' \to I$ be a surjective monotone map. Then there is a unique surjective monotone map $F : C(I') \to C(I)$ whose restriction to I' is f.*

In particular this implies that the isomorphism from Lemma 3.4.3 is unique.

Lemma 3.4.6 (see Lemma A.1.15)**.** *Let I and I' be cyclically ordered sets with at least two elements. Let $F : C(I') \to C(I)$ be a surjective monotone map with*

$F(I') \subseteq I$. Then for every $v \in C(I') \setminus I$ there is an element w of $C(I')$ such that $F^{-1}(v) = \{w\}$. This element w satisfies $w \in C(I') \setminus I'$. Also $F(I') = I$.

So for every surjective monotone map $F : C(I') \to C(I)$ which satisfies $F(I') \subseteq I$ there is an injective monotone map $\hat{f} : C(I) \setminus I \to C(I') \setminus I'$ such that for all $v \in C(I) \setminus I$ the equation $F^{-1}(v) = \{\hat{f}(v)\}$ holds.

Lemma 3.4.7 (see Lemma A.1.16)**.** *Let $F : C(I') \to C(I)$ be a surjective monotone map with $F(I') = I$, and $\hat{f} : C(I) \setminus I \to C(I') \setminus I'$ the injective monotone map with $F \circ \hat{f}(v) = v$ for all $v \in C(I) \setminus I$. Assume that I has at least two elements. Then for all $v, w \in C(I) \setminus I$, $F^{-1}([v, w]) = [\hat{f}(v), \hat{f}(w)]$ and $F^{-1}((v, w)) = (\hat{f}(v), \hat{f}(w)$.*

3.4.2 Definition and basic properties

Now we are ready to define our main object, which we call k-pseudoflower. This definition is derived from a definition of k-flowers in matroids in [11]. A k-flower is, essentially, a neat way to encode a collection of separations that interact in a nice way. The definition we will give is a lot more complicated than the original one, and there are mainly two reasons for that: First, a separation of a matroid is a bipartition of the ground set, and thus it is possible to encode several bipartitions in a partition of the ground set. But in our context we are working with vertex separations, and elements of separators have to be contained in several of the sets corresponding to partition classes. Second, we work in infinite structures and still want to obtain maximal k-pseudoflowers. Thus we have to make limit processes work, and as a chain of separations of order k may have a limit separation whose order is less than k, we have to allow that separations do not have order exactly k, but at most k. Making limit processes work also requires to put elements of separators between petals, sometimes in addition to putting them into petals. Also, every finite cyclic order is determined by its cardinality, but that is not the case for infinite cyclic orders, and thus we have to allow arbitrary cyclic orders as index sets and cannot work with an easily described representative.

Definition 3.4.8. Let I be a set of size at least 2 with a cyclic ordering and $(P_v)_{v \in C(I)}$ a family of vertex sets. Define $X = V \setminus \bigcup_{v \in C(I)} P_v$ and for all $v, w \in C(I) \setminus I$ let $V(v, w) = \bigcup_{z \in [v,w]} P_z \cup X$. Then $(P_i)_{i \in C(I)}$ is a *k-pseudoflower* if

- For every $v \in C(I) \setminus I$ we have $|P_v| = \frac{k-|X|}{2}$;

- For all distinct $v, w \in C(I) \setminus I$ the pair $S(v,w) = (V(v,w), V(w,v))$ is a separation of order at most k and $V(v,w) \cap V(w,v) = P_v \cup P_w \cup X$;
- For every $i \in I$ we have $P_{p(i)} \cup P_{s(i)} \subseteq P_i$, where $p(i)$ and $s(i)$ are the predecessor and successor of i in $C(I)$ respectively; and
- for all $i, i' \in I$,

$$\text{if } |P_i| = (k - |X|)/2 \text{ and } P_i = P_{i'} \text{ then } i = i'. \tag{$*$}$$

A k-pseudoflower is *finite* if the index set is finite. A *k-flower* is a k-pseudoflower in which all separations $S(v,w)$ are proper separations with order exactly k.

For $v \neq w \in C(I) \setminus I$ the set $V(v,w)$ is the *interval set* of $[v,w]$ and the separation $S(w,v)$ is the *interval separation* of $[v,w]$. For a non-trivial interval $I' \subseteq I$ let v and w be the unique elements of $C(I) \setminus I$ such that $I' = I \cap [v,w]$. Then the *interval set* $V(J)$ of J is $V(v,w)$ and the *interval separation* $S(J)$ of J is $S(w,v)$. For each index $i \in I$ the *petal* of i is its interval set $V(\{i\})$ and the *petal separation* of i is its interval separation $S(\{i\})$. As is usual for maps defined on power sets, we shorten $V(\{i\})$ and $S(\{i\})$ to $V(i)$ and $S(i)$. Note that $V(i) = P_i \cup X$ for all $i \in I$. A k-pseudoflower is called a *k-pseudoanemone* if the sets P_v for $v \in C(I) \setminus I$ are all empty and a *k-pseudodaisy* otherwise. Note that the number $(k - |X|)/2$, which is the size of the sets P_v with $v \in C(I) \setminus I$, indicates how far away a k-pseudodaisy is from being a k-pseudoanemone. In this sense it plays a very similar role to the difference between the local connectivity of adjacent petals and the local connectivity of non-adjacent petals as used, for example, in [4]. Also note that in a k-pseudoanemone, all the interval separations have X as their separator. In the case where the limit-closed universe comes from a graph G, this means that the set of petals corresponds to a partition of the components of $G - X$, and that X has exactly k elements.

In a graph, k-pseudoanemones correspond to partitions of the components of $G - X$. There are no infinite k-pseudodaisies in graphs, but they do exist e.g. in graph-like spaces.

Lemma 3.4.9. *Let Φ be a k-pseudoflower on index set I. Let a, b, c and d be elements of $C(I) \setminus I$ such that a, b and c are pairwise distinct and $b \in [a,c]$. Also assume that if $d \in \{a,b,c\}$, then $d = a$, and if $d \notin \{a,b,c\}$ then $d \in [c,a]$.*

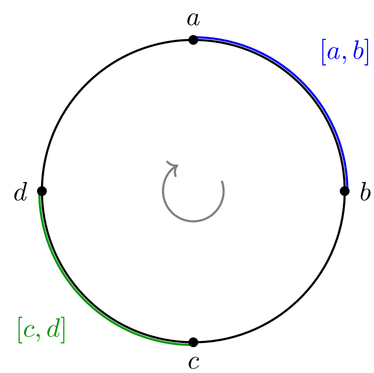

Figure 6: A cyclic order with two intervals. This figure depicts some of the notation used in Lemmas 3.4.9 and 3.6.19.

Then $V(a,b) \cap V(c,d) = (P_a \cap P_d) \cup (P_b \cap P_c) \cup X$ *and* $V(a,c) \cap V(b,d) = V(b,c) \cup (P_a \cap P_d)$.

Proof. Clearly $V(a,b) \cap V(c,d) \subseteq V(a,c) \cap V(c,a) = P_a \cup P_c \cup X$. Similarly $V(a,b) \cap V(c,d)$ is a subset of $P_a \cup P_b \cup X$, $P_d \cup P_b \cup X$ and $P_d \cup P_c \cup X$. Together these subsetrelations imply

$$\begin{aligned} & V(a,b) \cap V(c,d) \\ & \quad \subseteq (P_a \cup P_b \cup X) \cap (P_a \cup P_c \cup X) \cap (P_d \cup P_b \cup X) \cap (P_d \cup P_c \cup X) \\ & \quad = (P_a \cap P_d) \cup (P_b \cap P_c) \cup X. \end{aligned}$$

As also $(P_a \cap P_d) \cup (P_b \cap P_c) \cup X \subseteq V(a,b) \cap V(c,d)$, these two sets are equal. Thus

$$\begin{aligned} V(a,c) \cap V(b,d) &= V(b,c) \cup (V(a,b) \cap V(c,d)) \\ &= V(b,c) \cup (P_a \cap P_d) \cup (P_b \cap P_c) \cup X \\ &= V(b,c) \cup (P_a \cap P_d). \end{aligned}$$

□

Corollary 3.4.10. *If* $P_a \cap P_d = \emptyset$ *then* $S(b,c) = S(a,c) \wedge S(b,d)$.

Lemma 3.4.11. *Let* $(P_i)_{i \in C(I)}$ *be a* k*-flower on index set* I. *Then* $P_v \cap P_w = \emptyset$ *for all* $v \neq w \in C(I) \setminus I$.

Proof. By definition of a k-flower, the set $P_v \cup P_w \cup X$, which equals $V(v,w) \cap V(w,v)$, has exactly k many elements. Also P_v and P_w both have exactly $(k - |X|)/2$ many elements, so they must be disjoint. □

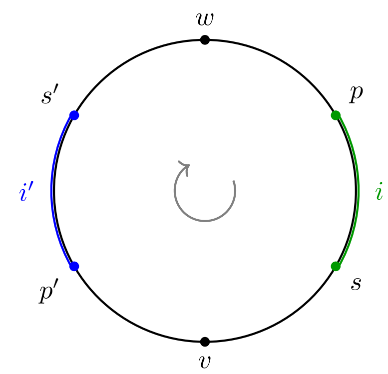

Figure 7: Some notation from the proof of Lemma 3.4.11.

Lemma 3.4.12. *For all* $x \in \bigcup_{v \in C(I) \setminus I} P_v$ *the set* $\{w \in C(I) \setminus I : x \in P_w\}$ *is an interval of* $C(I) \setminus I$.

Proof. Assume otherwise. Then there are pairwise distinct $a, b, c, d \in C(I) \setminus I$ such that $b \in [a, c]$, $d \in [c, a]$, $x \in P_b \cap P_d$ and $x \notin P_a \cup P_c$. But then $x \in V(a, c) \cap V(c, a)$, which contradicts the fact that $x \notin P_a \cup P_c \cup X$. $\square$

3.4.3 Concatenations of pseudoflowers

A separation (A, B) is *displayed* by a k-pseudoflower if it is an interval separation of that k-pseudoflower. We are now going to define a partial order on the set of k-pseudoflowers which has the property that if $\Phi \leq \Psi$ then all separations displayed by Φ are also displayed by Ψ.

Definition 3.4.13. A k-pseudoflower Φ' *extends* another k-pseudoflower Φ, written $\Phi \leq \Phi'$, if the sets X and X' coincide and there is a surjective map $F : C(I') \to C(I)$ respecting the cyclic ordering such that $F(I') = I$ and $V(v, w) = X \cup \bigcup_{z \in F^{-1}([v,w])} P'_z$ for all $v, w \in C(I) \setminus I$. If Φ' extends Φ, then Φ is a *concatenation* of Φ'.

Given a map $F : C(I') \to C(I)$ that witnesses that $\Phi \leq \Phi'$, by Lemma 3.4.6 there is for every $v \in C(I) \setminus I$ a unique $w \in C(I')$ with $F(w) = v$. Furthermore $w \in C(I') \setminus I'$, and $P_v = P_w$ by definition of $\Phi \leq \Phi'$. This allows us to identify $F^{-1}(C(I') \setminus I')$ with $C(I') \setminus I'$, and thus view Φ as a substructure of Φ'. (See Section 3.6.2 for a more detailed explanation of how this identification can be done for a whole chain of k-pseudoflowers.) Note that the witness of $\Phi \leq \Phi'$ might not be unique, and that the identification of cutpoints depends on the witness. Conversely, given a finite set D of cutpoints of a k-pseudoflower Φ' on index set I', we can take the subflower $\Phi'(D)$ of Φ' whose cutpoints correspond to D as follows: Let I be a cyclic order with $|D|$ many elements, and let $\hat{f} : C(I) \setminus I \to D$ be an isomorphism of cyclic orders. Extend the inverse of $\hat{f}$ to a surjective monotone map $F : C(I') \to C(I)$ with $F(I') \subseteq I$. For $v \in C(I) \setminus I$ there is a unique $d \in D$ with $F^{-1}(v) = d$, define $P_v = P_d$. For $i \in I$ let p and s be the predecessor and successor of i in $C(I)$ respectively, and define $P_i = \bigcup_{z \in F^{-1}([p,s])} P_z$. Then $(P_z)_{z \in C(I)}$ is a k-pseudoflower, and F witnesses $(P_z)_{z \in C(I)} \leq \Phi'$. Note that we need this construction only for finite D, but it works for all sets of cutpoints D that can be expressed as $C(I) \setminus I$ for some cyclically ordered set I.

Lemma 3.4.14. *Let Φ be a k-pseudoflower that extends a k-flower with three petals. Then for every $z \in C(I)$ there is $v \in C(I) \setminus I$ such that $P_z \cap P_v = \emptyset$. Furthermore, for every $i \in I$ and all nonempty $Y \subseteq P_i$ the set $\{z \in C(I) : Y \subseteq P_z\}$ is a non-trivial interval of $C(I)$.*

Proof. As Φ extends a k-flower with three petals, $C(I) \setminus I$ contains three distinct elements v_1, v_2 and v_3 and P_{v_1}, P_{v_2} and P_{v_3} are pairwise disjoint. By renaming if necessary assume that $z \in [v_1, v_2]$ and $v_3 \in [v_2, v_1]$. Then $P_z \cap P_{v_3} \subseteq V(v_1, v_2) \cap V(v_2, v_1) = P_{v_1} \cup P_{v_2} \cup X$. As every P_y with $y \in C(I)$ is disjoint from X, this implies that $P_z \cap P_{v_3} = \emptyset$.

Let $Y \subseteq P_i$ and for every $y \in P_i$ let C_y be the interval of those $z \in C(I)$ with $y \in P_z$. Let v be an element of $C(I) \setminus I$ such that $P_i \cap P_z = \emptyset$. As all C_y contain i but none contains v, the intersection of the C_y is a non-trivial interval of $C(I)$. □

We already stated that given two k-pseudoflowers $\Phi \leq \Psi$, there might be several maps witnessing the relation. It may not be obvious from the definition, but if Φ is a sufficiently meaningful k-pseudoflower, then two witnessing maps can only differ in very restricted ways. The rest of this section contains an analysis of the exact ways in which they can differ.

For k-pseudoflowers $\Phi \leq \Phi'$ such that Φ extends a k-flower with three petals let I'_N be the set of those $i \in I'$ such that there are witnesses of $\Phi \leq \Phi'$ that map i to different elements of I.

Lemma 3.4.15. *Let $i' \in I'_N$. Then the set of images of i' under witnesses of $\Phi \leq \Phi'$ contains only two elements of I, and those have a common neighbor $v(i')$ in $C(I)$ that satisfies $P'_i = P_{v(i')}$.*

Proof. In order to show that $P'_{i'} = P_v$ for some $v \in C(I) \setminus I$, let i_1 and i_2 be possible images of i' under witnesses of $\Phi \leq \Phi'$. Then in particular $P'_{i'} \subseteq P_{i_1}$ and $P'_{i'} \subseteq P_{i_2}$. By Lemma 3.4.14 there is $v_1 \in C(I) \setminus I$ such that $P_{v_1} \cap P_{i_1} = \emptyset$. Without loss of generality $v_1 \in [i_1, i_2]$, and as $i_1 \neq i_2$ there is $v_2 \in [i_2, i_1] \setminus I$. So

$$P'_{i'} \subseteq P_{i_1} \cap P_{i_2} \subseteq V(v_1, v_2) \cap V(v_2, v_1) = P_{v_1} \cup P_{v_2} \cup X,$$

implying that $P'_{i'} \subseteq P_{v_2}$ and thus by the first and third conditions of k-pseudoflowers that $|P'_{i'}| = |P_{v_2}| = (k - |X|)/2$.

If there is $i \in I$ with $P_i = P'_{i'}$, then every witness F of $\Phi \leq \Phi'$ has to map some i'_F to i, and $P'_{i'_F} = P_i = P'_{i'}$ then implies by $(*)$ that $i_{F'} = i'$. Thus in this case every witness of $\Phi \leq \Phi'$ maps i' to i, contradicting $i' \in I'_N$.

Let F be some witness of $\Phi \leq \Phi'$, then $P'_{i'} \subseteq P_{F(i')}$. By Lemma 3.4.14 there is $z_1 \in C(I) \setminus I$ such that $P_{F(i')} \cap P_{z_1} = \emptyset$, and the set $C := \{z \in C(I) : P'_{i'} \subseteq P_z\}$ is a non-trivial interval of $C(I)$. There is $z'_1 \in C(I') \setminus I'$ with $F(z'_1) = z_1$.

Assume that there is $i \in C \cap I$ that is (in the linear order of the interval C) neither the biggest not the smallest element of $C \cap I$. Then the two neighbors s and p of i in $C(I)$ are also contained in C, thus $P_s = P_p = P'_{i'} \neq P_i$. Hence there is $x \in P_i$ with $x \notin P_z$ for all other $z \in C(I)$, and there is $z_2 \in C(I')$ with $x \in P_{z_2}$. Now $F(z_2) = i$. If $i' \in (z'_1, z_2)$ then $F(i') \in (z_1, i]$ and otherwise $i' \in (z_2, z'_1)$, implying $F(i') \in [i, z_1]$. Both $](z_1, i] \cap C$ and $[i, z_1) \cap C$ are intervals of C.

As F was chosen as an arbitrary witness of $\Phi \leq \Phi'$, this implies that the subset of $C \cap I$ contains an interval of at most two elements that contains all images of i' under witnesses of $\Phi \leq \Phi'$. As $i' \in i'_N$ this implies that there are exactly two such images, and that they are neighbors in I and hence have a common neighbor $v(i')$ in $C(I)$. Then $v(i') \notin I$, and $P'_{i'} \subseteq V(v_1, v(i')) \cap V(v(i'), v_1) = X \cup P_{v_1} \cup P_{v(i')}$. As $P'_{i'}$ is disjoint from X and P_{v_1} and is at least as big as $P_{v(i')}$, this implies $P'_{i'} = P_{v(i')}$. $\square$

Lemma 3.4.16. *For every $i \in I$ there is some $i' \in I' \setminus I'_N$ such that every witness of $\Phi \leq \Phi'$ maps i' to i.*

Proof. Let I'_i be the set of those elements of I' that are mapped to i by some witness of $\Phi \leq \Phi'$. If I'_i contains only one element, then the lemma holds, so assume otherwise. Assume for a contradiction that all elements of I'_i are contained in I'_N. Denote the neighbors of i in $C(I)$ by s and p, then by Lemma 3.4.15 every element $i' \in I'_i$ satisfies either $P_{i'} = P_s$ or $P_{i'} = P_p$, hence there are exactly two such elements i_1 and i_2. Because both i_1 and i_2 are contained in I'_N, there are witnesses F_1 and F_2 of $\Phi \leq \Phi'$ with $F_j^{-1}(i) = \{i_j\}$. Thus $P_i \subseteq P_{i_1} \cap P_{i_2} = P_p \cap P_s$, so $P_i = P_s = P_{i_1} = P_{i_2}$, a contradiction. □

Lemma 3.4.17. *For every $v \in C(I) \setminus I$ that is not of the form $v(i')$ for some $i' \in I'_N$ there is a unique $v' \in C(I') \setminus I'$ that is mapped to v by all witnesses of $\Phi \leq \Phi'$.*

Proof. As Φ extends a k-flower with three petals, there is $i \in I$ such that $P_v \cap P_i = \emptyset$ and v is not a neighbor of i. Denote the predecessor and successor of i in $C(I)$ by p and s respectively. In particular $p \neq v \neq s$. Let i' be an element of I' that is mapped to i by every witness that $\Phi \leq \Phi'$ and denote its successor in $C(I')$ by s'.

Assume that there are witnesses F_1 and F_2 of $\Phi \leq \Phi'$ such that $F_1^{-1}(v) = v_1$ and $F_2^{-1}(v) = v_2$ and $v_1 \in (s, v_2)$. Then for $j \in \{1, 2\}$, $F_j^{-1}(s) \in [s', v_j]$ and thus $V(s', v_2) \subseteq V(p, v)$ and $V(v_1, s') \subseteq V(v, s)$. Hence every element of $V(v_1, v_2)$ is contained in $V(i) \cup V(s, v)$ and in $V(v, s)$. But $V(v, s) \cap (V(i) \cup V(s, v)) = V(i) \cup P_v$ and $V(v_1, v_2) \cap (V(i) \cup P_v) \subseteq V(v_1, v_2) \cap V(v_2, v_1) = X \cup P_v$, hence every element of $V(v_1, v_2)$ is contained in $X \cup P_v$. So $P_{\hat{i}} = P_v$ for all $\hat{i} \in [v_1, v_2]$ and in particular there can be at most one such $\hat{i}$.

As $v_1 \neq v_2$, some such $\hat{i}$ exists. Then $F_2(\hat{i}) \in [s, v]$ and, because $P_{\hat{i}} \cap P_i = P_v \cap P_i = \emptyset$, $F_1(\hat{i}) \in [v, p]$. So $\hat{i} \in I'_N$ and $v(\hat{i}) = v$, contradicting the assumptions about v. □

Let F be a witness that $\Phi \leq \Phi'$ and f the restriction of F to I'. Let $f' : I' \to I$ be obtained from f by choosing, for every $i' \in I'_N$, some neighbor of $v(i')$ in $C(I)$ as the image of f' and mapping every other element of I' to its image under f. Let $F' : C(I') \to C(I)$ be the unique surjective map respecting the cyclic order whose restriction to I' is f' (see Lemma 3.4.5).

Lemma 3.4.18. *The map F' is well-defined.*

Proof. It suffices to show that f' is surjective and respects the cyclic order. As for every $i \in I$ there is some $i' \in I'$ that is mapped to i by all witnesses of $\Phi \leq \Phi'$ and is in particular not contained in I'_N, f' is indeed surjective.

In order to show that f' respects the cyclic order, consider elements i_1, i_2 and i_3 of I' such that $f'(i_2) \in (f'(i_1), f'(i_3))$. In order to show that $i_2 \in (i_1, i_3)$ it suffices to consider the case that f' differs from f in at most three elements of I, and for that it is enough to show that f' respects the cyclic order if it maps only one element of I' differently than f. So assume that $f = f'$ except that $f'(i_1)$ is the successor of $f(i_1)$, the other cases are symmetric. Let i'_1 be an element of I' that is mapped to $f'(i_1)$ by all witnesses of $\Phi \leq \Phi'$. In particular $f(i'_1) = f'(i'_1)$, and $f(i_1)$ is the predecessor of $f(i'_1)$. If $f(i_1) \neq f(i_3)$, then $i_2 \in (i_1, i_3)$ because f respects the cyclic order. Otherwise $f^{-1}(f(i_1))$ is an interval of I' containing both i_1 and i_3 but neither i_2 nor i'_1, and $f'^{-1}(f'(i_1))$ is an interval of I' containing both i_1 and i'_1 but neither i_2 nor i_3. Together with $i_2 \in (i_1, i_3)$ this implies that $i_1 \in (i_3, i'_1)$ and hence that $i_1 \in (i_3, i_2)$. □

Lemma 3.4.19. *The map F' is a witness of $\Phi \leq \Phi'$.*

Proof. Let v and w be distinct elements of $C(I) \setminus I$. If $F'^{-1}([v, w])$ differs from $F^{-1}([v, w])$, then it does so in elements i of I'_N with $v(i) = v$ or $v(i) = w$ and one of their neighbors. Thus

$$X \cup \bigcup_{z \in F^{-1}([v,w])} P'_z = X \cup \bigcup_{z \in F'^{-1}([v,w])} P'_z.$$

□

3.4.4 Tangles and profiles in pseudoflowers

The theory presented in this chapter will be presented for two very similar objects: Profiles and tangles. The main results hold for the more general object, namely profiles. Thus they also hold for tangles, and in this case the proof for tangles is shorter. As the term tangle is more widely known than the term profile, the main results of this paper will be presented for both tangles and profiles, with the main differences to be found in this subsection. In the rest of the chapter, we will simply work with a set $\mathcal{P}$ of $k+1$-profiles, and as every $k+1$-tangle is also a $k+1$-profile the reader is free to think of a set of tangles.

Given a graph, a k-profile of that graph is defined to be a profile of S_k, the separation system of separations whose order is less than k. So it is obvious what the definition of a k-profile of a graph-like space, a path space or a limit-closed universe of vertex separations should be: A profile of S_k, which for the universe is a subsystem. Similarly, a k-tangle of a limit-closed universe of vertex separations is a tangle of S_k.

Readers interested in the technical details might note that if the universe is the universe of vertex separations of a graph, then this definition of a k-tangle is not exactly the same as the definition of a k-tangle of a graph. The two definitions are very close, however, and every k-tangle of the universe is also a k-tangle of the graph. Although Lemma 3.4.20 is formulated for k-tangles of limit-closed universes of vertex separations, the proof also holds for k-tangles of graphs.

Given a k-pseudoflower Ψ on index set I a $k+1$-profile P *is located by* Ψ if there is some $v \in C(I) \setminus I$ such that either $\forall w \in C(I) \setminus I - v : S(v,w) \in P$ or $\forall w \in C(I) \setminus I - v : S(w,v) \in P$.

For tangles we always have this property.

Lemma 3.4.20. *Let Φ be a k-pseudoflower and let T be a $k+1$-tangle. Then T is located by Φ.*

Proof. Let I be the index set of Φ and let $u \in C(I) \setminus I$. Let V' be the set of all elements w of $(C(I) \setminus I) - u$ such that $S(u,w) \in T$. If V' is empty or $(C(I) \setminus I) - u$, then u witnesses that T is located by Φ, so assume otherwise. As T is consistent, V' is an interval of $C(I) \setminus I$ and thus of the form $(u,v) \cap C(I) \setminus I$ or $(u,v] \cap C(I) \setminus I$ for some $v \in C(I) \setminus I$. First consider that case that V' is of the form $(u,v) \cap C(I) \setminus I$ for some $v \in C(I) \setminus I$. Then $S(u,v)$ is not an element of T and thus $S(v,u) \in T$. Hence for all $w \in (v,u)$ we have $S(v,w) \in T$ by consistency of T. Also for all $w \in (u,v)$ we have that $S(u,w) \in T$ so because also $S(v,u) \in T$ the tangle property implies $S(v,w) \in T$. So in this case we are done. Now consider the case that V' is of the form $(u,v]$ for some $v \in C(I) \setminus I - u$. Then $S(u,v) \in T$, so also $S(w,v) \in T$ for all $w \in [u,v)$ by consistency. Furthermore for all $w \in (v,u)$ we have $S(u,w) \notin T$, so $S(w,u) \in T$ and thus $S(w,v) \in T$ by the tangle property. So in this case we are done as well. ☐

For general $k+1$-profiles, the statement holds with an additional assumption that the k-pseudoflower is sufficiently large.

Lemma 3.4.21. *Let Φ be a k-pseudoflower which extends a k-flower with at least four petals and let P be a $k+1$-profile. Then P is located by Ψ.*

Proof. Let I be the index set of Ψ. Let s, t, x and y be elements of $C(I) \setminus I$ such that $\Phi(\{s,t,x,y\})$ is a k-flower with four petals and $t \in [s,x]$ and $y \in [x,s]$. Then P contains an orientation of $S(s,x)$ and an orientation of $S(t,y)$, and by Corollary 3.4.10 and Lemma 3.4.11 P also contains the supremum of those two orientations. So P contains the inverse of a petal separation of $\Phi(\{s,t,x,y\})$, without loss of generality assume that P contains $S(t,s)$.

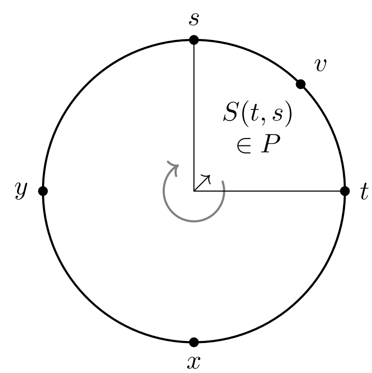

Figure 8: $S(t,s)$ is an interval separation of a k-pseudoflower and is contained in a $k+1$-profile P. See also Lemma 3.4.21.

In the interval $[s,t]$ let v be the supremum of $\{w \in [s,t) \setminus I \colon S(t,w) \in P\}$. There are three cases: $v = t$, $S(v,x) \in P$ and $S(x,v) \in P$. In the first case we have $\forall w \in C(I) \setminus I - v : S(v,w) \in P$ and are done.

Consider the second case: $v \neq t$ and $S(v,x) \in P$. Then $S(v,w) = S(v,x) \vee S(t,w)$ for all $w \in [s,v)$ by Corollary 3.4.10 and Lemma 3.4.11 and thus $S(v,w) \in P$. Hence $\forall w \in C(I) \setminus I - v : S(v,w) \in P$.

In the last case, $v \neq t$ and $S(x,v) \in P$. Again $S(t,v) = S(t,s) \vee S(x,v)$ and thus $S(t,v) \in P$. If t is the successor of v in $C(I) \setminus I$, then $S(t,v) \in P$ implies that $\forall w \in C(I) \setminus I - v : S(w,v) \in P$. Otherwise let $w \in (v,t) \setminus I$. By the definition of v, P does not contain $S(t,w)$, so $S(t,w) = S(t,s) \vee S(y,w)$ implies that $S(y,w)$ is not an element of P. Thus P contains $S(w,y)$, and $S(w,v) = S(w,y) \vee S(x,v)$ implies that $S(w,v) \in P$. As this is true for all $w \in (v,t) \setminus I$, $\forall w \in C(I) \setminus I - v : S(w,v) \in P$ holds. $\square$

For the remainder of this chapter, instead of differentiating between tangles and profiles, we will always work more generally with profiles located by the relevant flowers. Since, as noted above, Lemma 3.4.20 also works for tangles of graphs, our approach encompasses this case as well.

An equivalence class E of $\leq k$-separations is *displayed* by a k-pseudoflower Φ if an element of E is displayed by Φ. We are interested in k-pseudoflowers displaying as many equivalence classes as possible. For that, we also consider the following relation on k-pseudoflowers. Note that it is not a partial order as it is not antisymmetric.

Definition 3.4.22. Let Φ and Φ' be two k-pseudoflowers. Define $\Phi \preccurlyeq \Phi'$ if and only if every relevant separation displayed by Φ is equivalent to a separation displayed by Φ'. If $\Phi \preccurlyeq \Phi' \preccurlyeq \Phi$, then the k-pseudoflowers are *equivalent*.

3.5 Extending pseudoflowers

Two separations *properly cross* if all of their corner separations are contained in profiles of $\mathcal{P}$. In this section we will start with a separation properly crossing a petal of a pseudoflower, and our goal is to remedy this by extending the pseudoflower.

Let Ψ be a k-pseudoflower and i a petal of Ψ with predecessor p and successor s. A separation (C, D) properly crossing $S(i)$ is *anchored* at $v \in C(I) \setminus I$ if $(C, D) \wedge S(i) = S(v, p)$ and $(D, C) \wedge S(i) = S(s, v)$.

Lemma 3.5.1. *Let Φ be a k-pseudoflower distinguishing at least three elements of $\mathcal{P}$ located by Φ. Let (C, D) be a separation of order k which properly crosses some petal separation $S(i)$ of Φ. Then there is a separation (C', D') properly crossing $S(i)$ such that*

- *$(C, D) \vee S(i) = (C', D') \vee S(i)$ and $(D, C) \vee S(i) = (D', C') \vee S(i)$.*
- *(C', D') or its inverse is anchored at some $v \in C(I) \setminus I$.*

Proof. Let $P_1 \in \mathcal{P}$ be a profile which contains both (C, D) and $S(i)$. If Φ distinguishes two profiles which contain both (C, D) and $S(i)^*$ then let $P_2 \in \mathcal{P}$ be a profile which contains both (D, C) and $S(i)^*$. Then some profile $P_3 \in \mathcal{P}$ which contains both (C, D) and $S(i)^*$ is distinguished from P_2 by Φ. If Φ does not distinguish any profiles which contain both (C, D) and $S(i)^*$, then let P_3 be some such profile. As Φ distinguishes at least three profiles, it distinguishes some profile P_2 from both P_1 and P_3. Because P_2 is distinguished from P_1, it contains $S(i)^*$. Because P_2 is also distinguished from P_3, it does not contain (C, D), so it contains (D, C). In both cases, P_2 contains $S(i)^*$ and (D, C), while P_3 contains $S(i)^*$ and (C, D). Furthermore Φ distinguishes P_2 and P_3.

Denote the predecessor of i in $C(I)$ by p and the successor by s. As Φ distinguishes P_2 and P_3, there is a cutpoint $v \in C(I) \setminus I$ such that $S(v, s)$ distinguishes P_2 and P_3. Because both P_2 and P_3 contain $S(i)^*$, $S(v, p)$ also distinguishes P_2 and P_3, and P_2 contains $S(v, s)$ if and only if it contains $S(v, p)$.

From here on consider the case that P_3 contains $S(v, p)$, the other case is symmetric. Define $(C'', D'') = (C, D) \vee S(v, p)$. Because P_3 contains both (C, D) and $S(v, p)$ and P_2 contains both $S(p, v)$ and (D, C), the separation

(C'', D'') has order k and distinguishes P_2 and P_3. Then $(C'', D'') \vee S(i) = (C, D) \vee S(v, p) \vee S(i) = (C, D) \vee S(i)$ and in particular $(C'', D'') \vee S(i)$ is the same as $(C, D) \wedge S(i)$. Let P_4 be a profile that contains both $S(i)$ and (D, C). Then $(D'', C'') \leq (C, D)$ implies that (D'', C'') is contained in P_4. The profile P_3 contains both (C'', D'') and $S(i)^*$ and the profile P_4 contains both (D'', C'') and $S(i)$, so $(D'', C'') \vee S(i)$ is a separation of order k. Also $(D'', C'') \vee S(i) \leq (D, C) \vee S(i)$. Because $V(v, p) \cup V(s, p) = V$, we have $D \cup V(s, p) = (D \cap V(p, v)) \cup V(s, p)$. The set $D \cup V(s, p)$ is the left side of $(D, C) \vee S(i)$ and $(D \cap V(p, v)) \cup V(s, p)$ is the left side of the separation $(D'', C'') \vee S(i)$. Also, the right side of $(D, C) \vee S(i)$ is a subset of the right side of $(D'', C'') \vee S(i)$. Because these two separations have the same order, they are equal.

Define $(D', C') = (D'', C'') \vee S(s, v)$. Just as in the previous paragraph (C', D') is a separation of order k distinguishing P_2 and P_3. Also $(D', C') \vee S(i) = (D'', C'') \vee S(i)$ and $(C', D') \vee S(i) = (C'', D'') \vee S(i)$. Furthermore P_1 contains (C', D') and $S(i)$, and P_3 contains (D', C') and $S(i)^*$, so $(C', D') \wedge S(i)$ has order k. Because $S(v, p) \leq (C', D') \leq S(v, s)$, also

$$(V(v, p) \cup P_s, V(p, v)) = S(v, s) \wedge S(i) \geq (C', D') \wedge S(i) \geq S(v, p) \wedge S(i) = S(v, p),$$

which, as $(C', D') \wedge S(i)$ has the same order as $S(v, p)$, implies that $(C', D') \wedge S(i) = S(v, p)$. Similarly $(D', C') \wedge S(i) = S(s, v)$. $\square$

Lemma 3.5.2. *Let Ψ be a k-pseudoflower and i a petal of Ψ with predecessor p and successor s. If a k-separation (C, D) properly crossing $S(i)$ is anchored at $v \in C(I) \setminus I$, then $C \cap P_s = \emptyset$, $D \cap P_p = \emptyset$,$X \subseteq C \cap D$ and $(C \cap D) \setminus V(p, s) = P_v$.*

Proof. Since P_s is contained in the separator of $S(i)$, but does not meet $V(v, p)$ it cannot meet C. Similarly $(C, D) \wedge S(i) = S(s, v)$ implies that P_p does not meet D. The third statement is immediate from the fact that X occurs in the separators of both $S(s, v)$ and $S(v, p)$.

For the last equation, note that $(C, D) \wedge S(i) = S(v, p)$ implies both $C \cap V(s, p) = V(v, p)$ and $D \cup V(p, s) = V(p, v)$. Taking the intersection gives $(C \cap D \cap V(s, p)) \cup (C \cap V(s, p) \cap V(p, s)) = V(v, p) \cap V(p, v)$. Deleting $V(p, s)$ on both sides and simplifying gives the result. $\square$

Lemma 3.5.3. *Let Φ be a k-pseudoflower distinguishing at least three elements of $\mathcal{P}$ located by Φ. Let (C, D) be a separation of order k which properly crosses* *some petal separation $S(i)$ of Φ and is anchored at $v \in C(I) \setminus I$. Let p be and s be the predecessor and successor of i in $C(I)$ respectively. Then there is an extension Φ' of Φ, witnessed by F, such that*

- *For all $i' \in I \setminus \{i\}$ there is only one element contained in $F^{-1}(i')$.*
- *$F^{-1}(i)$ contains exactly three elements i_1, m and i_2. By switching the names of i_1, m and i_2 if necessary we can assume that m is the successor of i_1 and the predecessor of i_2.*
- *$P_m = (C \cap D) \setminus V(s, p)$.*
- *The interval set of i_1 is $C \cap V(i)$ and the interval set of i_2 is $D \cap V(i)$.*
- *(C, D) is an interval separation of Φ'.*

Proof. Let Φ' be obtained from Φ by replacing $i \in I$ with i_1, m and i_2 in order and setting $P_m = (C \cap D) \setminus V(s, p)$, $P_{i_1} = C \cap P_i$ and $P_{i_2} = D \cap P_i$.

We will now prove that Φ' is a k-pseudoflower. The third condition is trivial. By Lemma 3.5.2 we know that $C \cap D$ consists of exactly P_m, X and P_v, so $|P_m| = k - |P_v| - |X| = \frac{k-|X|}{2}$.

Assume for a contradiction that $(*)$ fails. We may assume without loss of generality that $|P_{i_1}| = \frac{k-|X|}{2}$. It follows that $P_m = P_{i_1}$ and thus $P_{i_1} \subseteq P_{i_2}$. But then $(D, C) \vee S(i)^*$ is cotrivial with witness (C, D) and therefore is contained in every profile, contradicting the fact that $S(i)$ properly crosses (C, D).

To show that Φ' is a k-pseudoflower, it thus remains to show the second condition. Let $x, y \in C(I)$ be arbitrary. Because Φ was a k-pseudoflower, w.l.o.g. $y = m$. Furthermore, since P_{i_1} and P_{i_2} meet only in P_m, we have $V(x, m) \cap V(m, x) = P_x \cup P_m \cup X$.

Thus it is enough to show that $(V(x, m), V(m, x))$ is a separation. Without loss of generality x occurs in the interval in the interval from m to v in $C(I)$, the other case is analogous. Note that, as (C, D) is anchored at v, $S(v, p) \leq (C, D) \leq S(v, s)$ and thus $V(v, p) \subseteq C \subseteq V(v, s)$. So

$$V(x, p) \cup C = V(x, p) \cup (C \cap V(i)) = V(x, m)$$

and similarly $V(p, x) \cap D = V(m, x)$. Hence $S(v, p) \vee S(C, D) = S(x, m)$, implying that $S(x, m)$ is indeed a separation.

Thus Φ' is a k-pseudoflower. By construction, (C, D) appears as the interval separation $S(m, v)$. □

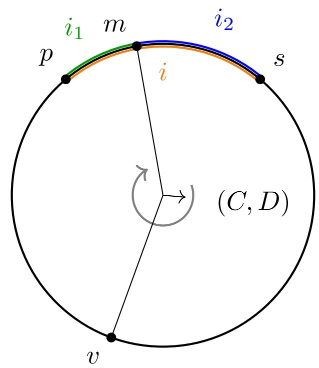

Figure 9: The the index i is replaced by the indices i_1 and i_2, creating a new cutvertex m. This allows us to display the separation (C, D), see also Lemma 3.5.3.

3.6 Maximal pseudoflowers

Let $(\Phi_j)_{j \in J}$ be a $\leq$-chain of k-pseudoflowers which distinguish at least three elements of $\mathcal{P}$ that are located by the Φ_j. In this section we will prove that there is a k-pseudoflower which is an upper bound of the chain $(\Phi_j)_{j \in J}$. The proof works for both k-pseudodaisies and k-pseudoanemones, but can be shortened a lot if the Φ_j are k-pseudoanemones as follows: The careful construction of the sets P_v with $v \in V_N$ in Section 3.6.3 becomes redundant as all those sets are necessarily empty. For the same reason, in Section 3.6.5 there is no need to derive the sets P_v with $v \in C(I \setminus I_N) \setminus (I \setminus I_N)$ from their counterparts in $C(I) \setminus I_N$. Furthermore, it is possible to carry the following stronger version of $(*)$ through the limit process: That no petal should be empty. This stronger version then makes all witnesses of comparability of k-pseudoflowers unique, making the search for a compatible choice of witnesses unnecessary. Also, the definition of I' can be made easier, as it simply consists of those $i \in I$ with $P_i \neq \emptyset$, and $C(I')$ is more obviously related to C. Lastly, in k-anemones the condition of Corollary 3.4.10 that $P_a \cap P_d = \emptyset$ always holds, which means that the proof in Section 3.6.4 that all interval separations of the newly constructed k-pseudoflower are indeed interval separations can be streamlined.

3.6.1 Finding compatible witnesses

The first step in constructing an upper bound of a chain of k-pseudoflowers is to construct the index set. For that we would like to take the inverse limit of the index sets, and for that we need compatible witnesses of the comparability of the k-pseudoflowers in the chain. Phrased more formally, we need for every $j \leq l \in J$ a witness F_{lj} of $\Phi_j \leq \Phi_l$ such that for $j \leq l \leq m \in J$ the concatenation of F_{lj} with F_{ml} is F_{mj}.

In Section 3.4.3 we showed that, given sufficiently large k-pseudoflowers $\Phi \leq \Phi'$ on index sets I and I', only few elements i of I' have several possible images under witnesses of $\Phi \leq \Phi'$, and that picking some possible image for every such i gives again a witness of $\Phi \leq \Phi'$. So defining the F_{lj} amounts to finding a way of making all these choices in a way that is compatible over the whole chain of k-pseudoflowers. For that, we first show how these choices interact in a chain of only three k-pseudoflowers.

So for the next two lemmas let $\Phi_1 \leq \Phi_2 \leq \Phi_3$ be k-pseudoflowers on index sets I_1, I_2 and I_3 respectively such that Φ_1 extends a k-flower with three petals. For some $i_3 \in I_3$ we are now interested in possible images of i_3 under witnesses of $\Phi_2 \leq \Phi_3$, their images under witnesses of $\Phi_1 \leq \Phi_2$ and how those relate to witnesses of i_3 under witnesses of $\Phi_1 \leq \Phi_3$.

Lemma 3.6.1. *If i_3 has two possible images i_1 and i_1' under $\Phi_1 \leq \Phi_3$, and there is $i_2 \in I_2$ whose possible witnesses under $\Phi_1 \leq \Phi_2$ are also i_1 and i_1', then all witnesses of $\Phi_2 \leq \Phi_3$ map i_3 to i_2.*

Proof. By Lemma 3.4.16 there is some i_3' whose unique image under $\Phi_2 \leq \Phi_3$ is i_2. Then i_3' has i_1 and i_1' as possible images under $\Phi_1 \leq \Phi_3$ and $P_{i_3} = P_{i_3'}$ by Lemma 3.4.15, so $i_3 = i_3'$. $\square$

Lemma 3.6.2. *If i_3 has two possible images i_1 and i_1' under $\Phi_1 \leq \Phi_3$, i_1' successor of i_1 in I_1, and i_3 has two possible images i_2 and i_2' under $\Phi_2 \leq \Phi_3$, i_2' successor of i_2 in I_2, then every witness of $\Phi_1 \leq \Phi_2$ maps i_2 to i_1 and i_2' to i_1'.*

Proof. See Fig. 10 for a depiction of the notation used in this proof. As a concatenation of witnesses is again a witness, all witnesses of $\Phi_1 \leq \Phi_2$ map i_2 to one of i_1 and i_1', and similarly for i_2'. If one element of I_2 had both i_1 and i_1' as possible images under $\Phi_1 \leq \Phi_2$, then by the previous lemma i_3 could not have both i_2 and i_2' as possible images under $\Phi_2 \leq \Phi_3$. Hence no element of I_2 has both i_1 and i_1' as possible images under witnesses of $\Phi_1 \leq \Phi_2$, and in particular both i_2 and i_2' have a unique image under such witnesses. Let $\hat{i}_1$ be an element of I_1 that is neither i_1 nor i_1', and let $\hat{i}_3$ be an element of I_3 that is mapped to $\hat{i}_1$ by all witnesses of $\Phi_1 \leq \Phi_3$.

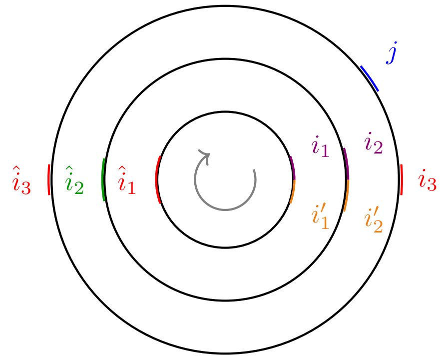

Figure 10: Three compatible k-pseudoflowers $\Phi_1 \leq \Phi_2 \leq \Phi_3$. The outer circle depicts Φ_3 and three elements of its index set, the inner circle depicts Φ_1.

Assume for a contradiction that both i_2 and i'_2 have the same image in I_1 under witnesses of $\Phi_1 \leq \Phi_2$, and that this image is i'_1 (the other case is symmetric). Let j be an element of I_3 that is mapped to i_2 by all witnesses of $\Phi_2 \leq \Phi_3$. As no witness of $\Phi_2 \leq \Phi_3$ maps $\hat{i}_3$ to i_2 or i'_2 and there is some witness that maps i_3 to i'_2, we have $j \in (\hat{i_3}, i_3)$. Also, by Lemma 3.4.18 (and because $j \neq i_3$) there is a witness of $\Phi_1 \leq \Phi_3$ that maps j to i'_1, i_3 to i_1 and $\hat{i_3}$ to $\hat{i_1}$. This implies $j \in (i_3, \hat{i}_3)$, a contradiction.

So i_2 and i'_2 have distinct unique images under witnesses of $\Phi_1 \leq \Phi_2$. If $\hat{i}_2$ denotes the image of $\hat{i}_3$ under some witness of $\Phi_2 \leq \Phi_3$, then $\hat{i}_2 \in (i'_2, i_2)$ and $\hat{i}_1 \in (i'_1, i_1)$ implies that indeed all witnesses of $\Phi_1 \leq \Phi_2$ map i_2 to i_1 and i'_2 to i'_1. $\square$

Now we want to use these two lemmas to define the maps F_{lj}. For all indices $j, l \in J$ with $j \leq l$ define a map $f_{lj} : I_l \to I_j$ as follows: For $i \in I_l$, if there are indices $j = j_0 < j_1 < \ldots < j_n = l$, $n \geq 1$, in J such that for $i_n = i$ and for every $m \leq n$ every witness of $\Phi_{j_{m-1}} \leq \Phi_{j_m}$ maps i_m to the same element i_{m-1} of I_{m-1}, then let $f_{lj}(i) = i_0$ (first case). Otherwise there are two possible images i_j and i'_j of i under witnesses of $\Phi_j \leq \Phi_l$ such that i'_j is the successor of i_j. In this case let $f_{lj}(i) = i'_j$ (second case).

Lemma 3.6.3. *For all $j \leq l \in J$ and all $i \in I_j$, $f_{lj}(i)$ is well defined.*

Proof. In order to show that $f_{lj}(i)$ is well-defined, it suffices to consider the case where $f_{lj}(i)$ is defined via the first case. Assume that $j = j_0 \leq \ldots j_n = l$ in J and $j = l_0 \leq \ldots \leq l_m = l$ are two chains of elements of J via which $f_{lj}(i)$ could be defined. We will show by induction on $n + m$ that the value of $f_{lj}(i)$ is the same in both cases. If one of n and m is 0, then $j = l$ and the claim holds, so assume otherwise. Furthermore, if $j_{n-1} = l_{m-1}$, then it suffices to apply the induction hypothesis to the chains $j_0 \leq \ldots \leq j_{n-1}$ and $l_0 \leq \ldots \leq l_{m-1}$. So assume that $j_{n-1} < l_{m-1}$, the other case is symmetric. Let r be an integer such that $l_{r-1} \leq j_{n-1} \leq l_r$, and i_r the image of i in I_{l_r} under the chain of the l_s. If i_r has several images under witnesses of $\Phi_{j_{n-1}} \leq \Phi_{l_r}$, then i also has several images under witnesses of $\Phi_{j_{n-1}} \leq \Phi_l$. So i_r has a unique image i_s under witnesses of $\Phi_{j_{n-1}} \leq \Phi_{l_r}$, and that image is also the unique image of i under witnesses of $\Phi_{j_{n-1}} \leq \Phi_l$. Then $j_0 \leq \ldots \leq j_{n-1} \leq l_r$ and $l_0 \leq \ldots \leq l_r$ are chains on which, because $n + r \leq n + m - 1$, the induction hypothesis can be applied. Thus the lemma holds. $\square$

Lemma 3.6.4. *For all $j \leq l \leq m$ in J and $i \in I_m$, $f_{lj} \circ f_{ml}(i) = f_{mj}(i)$.*

Proof. First consider the case that $f_{mj}(i)$ is defined via the first case. Let $j_0 \leq \ldots j_n$ be a sequence via which $f_{mj}(i)$ could have been defined, with $j_{r-1} \leq l \leq j_r$. If i_r has two possible images i_l, i'_l under witnesses of $\Phi_l \leq \Phi_{j_r}$, then $f_{ml}(i)$ is one of those. Also, every witness of $\Phi_{r-1} \leq \Phi_l$ maps both i_l and i'_l to i_{r-1}, implying that both $f_{lj}(i_l)$ and $f_{lj}(i'_l)$ are defined via the first case and are equal to $f_{mj}(i)$. So in this case $f_{lj} \circ f_{ml}(i) = f_{mj}(i)$. If i_r has a unique image i_l under witnesses of $\Phi_j \leq \Phi_r$, then all witnesses of $\Phi_{r-1} \leq \Phi_l$ map i_l to i_{r-1}, and thus $f_{ml}(i)$ is defined via the first case and also $f_{lj}(i_l)$ is defined via the first case. As $f_{mj}(i)$ is well defined, this implies that $f_{lj} \circ f_{ml}(i) = f_{mj}(i)$.

So assume that $f_{mj}(i)$ is defined via the second case, and that i_j and i'_j are the two possible images of i under witnesses of $\Phi_j \leq \Phi_m$. If there is $i_l \in I_l$ such that both i_l and i'_l are images of witnesses of $\Phi_j \leq \Phi_l$, then by Lemma 3.6.1 every witness of $\Phi_l \leq \Phi_m$ maps i to i_l. Thus $f_{mj}(i) = f_{ml}(i_l) = f_{ml} \circ f_{lj}(i)$. So assume that there is no $i_l \in I_l$ that has i_j and i'_j as possible images under witnesses of $\Phi_j \leq \Phi_l$. By renaming assume that i'_j is the successor of i_j. Then by Lemma 3.6.2 there are i_l and i'_l in I_l such that every witness of $\Phi_j \leq \Phi_l$ maps i_l to i_j and i'_l to i'_j. Also, i'_l is the successor of i_l, and every witness of $\Phi_l \leq \Phi_m$ maps i to i_l or i'_l. Then $f_{ml}(i)$ cannot be defined via the first case, because that would imply that $f_{mj}(i)$ would also be defined via the first case. So $f_{mj}(i) = i'_j = f_{lj}(i'_l)$ and $f_{ml}(i) = i'_l$. $\square$

For $j \leq l \in J$ let F_{lj} be the unique extension of f_{lj} to a surjective map $C(I_l) \to C(I_j)$ respecting the cyclic order with $F_{lj}(I_l) = I_j$, which exists by Lemma 3.4.5.

$$
\begin{tikzcd}
C(I_m) \arrow[d, "F_{ml}"] \arrow[dd, bend right=60, "F_{mj}"'] & I_m \arrow[l, hook'] \arrow[d, "f_{ml}"'] \arrow[dd, bend left=60, "f_{mj}"] \\
C(I_l) \arrow[d, "F_{lj}"] & I_l \arrow[l, hook'] \arrow[d, "f_{lj}"'] \\
C(I_j) & I_j \arrow[l, hook']
\end{tikzcd}
\qquad
\begin{tikzcd}
C(I) \arrow[dr, "\Pi_l"] \arrow[ddr, "\Pi_j"'] & & & I \arrow[lll, hook'] \arrow[dl, "\pi_l"'] \arrow[ddl, "\pi_j"] \\
& C(I_l) \arrow[d, "F_{lj}"] & I_l \arrow[l, hook'] \arrow[d, "f_{lj}"'] & \\
& C(I_j) & I_j \arrow[l, hook'] &
\end{tikzcd}
$$

Figure 11: To the left: The maps $(F_{lj})_{j\leq l\in J}$ extend the compatible maps $(f_{lj})_{j\leq l\in J}$ and are therefore compatible witnesses that the k-pseudoflowers are comparable are compatible. To the right: The maps $(\Pi_j)_{j\in J}$ defined via the projections $(\pi_j)_{j\in J}$ are compatible with the maps $(F_{lj})_{j\leq l\in J}$.

Lemma 3.6.5. *The maps F_{lj} are witnesses of $\Phi_j \leq \Phi_l$ and are compatible in the sense that for $j \leq l \leq m \in J$ the concatenation of F_{lj} and F_{ml} is F_{mj}.*

Proof. As f_{lj} can be obtained from the restriction of any witness of $\Phi_j \leq \Phi_l$ to I_l by changing the images of some elements of I_j to another image they can have under witnesses of $\Phi_j \leq \Phi_l$, by Lemma 3.4.19 F_{lj} is also a witness of $\Phi_j \leq \Phi_l$. Furthermore the restriction of $F_{lj} \circ F_{ml}$ to I_m is equal to f_{mj}, and $F_{lj} \circ F_{ml}$ is surjective, respects the cyclic order and satisfies $F_{mj}(I_m) = I_j$. Thus, by uniqueness of the extensions, $F_{mj} = F_{lj} \circ F_{ml}$. □

Corollary 3.6.6. *There are witnesses F_{lj} of $\Phi_j \leq \Phi_l$, one for all pairs of indices $j \leq l \in J$, such that for all $j \leq l \leq m \in J$ the concatenation of F_{lj} and F_{ml} is F_{mj}.* □

3.6.2 Inverse limits of pseudoflowers

As we showed in the last section, there is a compatible family of witnesses F_{lj} of $\Phi_j \leq \Phi_l$, one for every pair of indices $j \leq l \in J$. For the following construction we fix one such family. Then there is an inverse limit I of the I_j with projections $(\pi_j)_{j\in J}$. Because all F_{lj} respect the cyclic order, the inverse limit also has a cyclic order which is respected by the projections. Further, because all the F_{lj} are surjective, so are the projections. We will construct a partition $\Psi := (P_i)_{i\in C(I)}$ that is a upper bound of the chain $(\Phi_j)_{j\in J}$ and is a k-pseudoflower except that it may violate $(*)$. So far we can only define P_i with $i \in I$: It is the intersection of all sets $P_{\pi_j(i)}$ with $j \in J$.

We will now relate $C(I)$ to the cyclically ordered sets $C(I_j)$. By Lemma 3.4.5 every projection $\pi_j : I \to I_j$ can be extended uniquely to a surjective map $\Pi_j : C(I) \to C(I_j)$ that respects the cyclic order. For $j \leq l \in J$, the restriction of $F_{lj} \circ \Pi_l$ to I_l is π_j, and hence $F_{lj} \circ \Pi_l = \Pi_j$. By Lemma 3.4.6, the cutpoints of the $C(I_j)$ can be identfied with each other and with cutpoints of $C(I)$ as follows: Given $j \leq l \leq n \in J$ and $v \in C(I_j) \setminus I_j$, let u be the unique element of I_l with $F_{lj}(u) = v$ and let u' be the unique element of I_n with $F_{nl}(u') = u$. Then u' is also the unique element of I_n with $F(nj)(u') = v$. Also, if w is the unique element of $C(I)$ with $\Pi_l(w) = u$, then w is the unique element of $C(I)$ with $\Pi_j(w) = v$. So it is well-defined to identify v with w and with all cutpoints that get mapped to v via some F_{lj}. We do this identification for all cutpoints of all $C(I_j)$ and denote the set of cutpoints of $C(I)$ that are identified with a cutpoint of some $C(I_j)$ by V_F. Then for $v \in V_F$, P_v is the same for all $j \in J$ where v is a cutpoint, so we take that vertex set also to be P_v for Ψ. Also note that, for $v, w \in V_F$, $V(v, w)$ does not depend on the k-pseudoflower Φ_j with respect to which it is defined. But $V(v, w)$ taken in Ψ is not yet defined, and when it is we will first have to show that it is equal to $V(v, w)$ taken in some Φ_j.

In order to properly distinguish here, for cutpoints v and w of $C(I)$ we introduce the notation

$$V'(v, w) = X \cup P_v \cup P_w \cup \bigcup_{z \in [v,w] \cap (I \cup V_F)} P_z$$

and

$$\hat{V}(v, w) = X \cup \bigcup_{z \in [v,w] \cap C(I)} P_z$$

while the notation $V(v, w)$ is reserved for the case $v, w \in V_F$ and the value taken in some Φ_j. Note that up to now, in many cases V' and $\hat{V}$ are not yet well-defined because P_z is not yet defined for cutpoints $z \notin V_F$. Those sets P_z are going to be defined later.

For $v, w \in V_F$, we can see that $V'(v, w)$ and $V(v, w)$ are the same:

Lemma 3.6.7. *For all distinct $v, w \in V_F$ we have $V(v, w) \setminus X = \bigcup_{z \in [v,w]} P_z$ where the interval is taken in $I \cup V_F$.*

Proof. Let us first show that $V(v, w) \setminus X$ is a subset of $\bigcup_{z \in [v,w]} P_z$. For this, let $u \in V(v, w) \setminus X$. If also $u \in V(w, v)$ then $u \in P_v \cup P_w$ and we are done, so assume that $u \notin V(w,v)$. If there is $z \in V_F$ such that $u \in P_z$, then $u \notin V(w,v)$ implies that $z \notin [w,v]$ and thus $z \in [v,w]$. If there is no $z \in V_F$ such that $u \in P_z$, then for all $j \in J$ there is a unique $i_j \in I_j$ such that $u \in P_{i_j}$. In this case also $F_{lj}(i_l) = i_j$ for all $j \leq l \in J$. So there is $i \in I$ such that $\Pi_j(i) = i_j$ for all $j \in J$, and $u \in P_i$. Again $u \notin V(v,w)$ implies $i \in [v,w]$.

To prove the other inclusion, let $j \in J$ be sufficiently large such that $C(I_j)$ contains v, w and z if $z \in V_F$. Then Φ_j witnesses that $P_z \subseteq V(v,w)$. Furthermore all P_z are disjoint from X. $\square$

Corollary 3.6.8. $X = V(G) \setminus \bigcup_{z \in I \cup V_F} P_z$.

Corollary 3.6.9. *For all distinct $v, w \in V_F$ we have $V(v,w) = V'(v,w)$.*

Corollary 3.6.10. *For all distinct $v, w \in V_F$ the pair $(V'(v,w), V'(w,v))$ is a separation of order at most k with separator $P_v \cup P_w \cup X$.*

3.6.3 Completing the index set

Let V_N be the set of cutpoints of $C(I)$ that are not contained in V_F. In this subsection we will define the values P_z with $z \in V_N$.

Lemma 3.6.11. *Every $v \in V_N$ has a unique neighbor in $C(I)$, which is an element of I.*

Proof. As $v \in V_N$, no Π_j maps v to a cutpoint. Let i be the element of I with $\pi_j(i) = \Pi_j(v)$ for all $j \in J$. If it exists, let z be a neighbor of v in $C(I)$. For all $j \in J$, $\Pi_j(z)$ is $\Pi_j(v)$ or one of its neighbors in $C(I_j)$. Also no cutpoint can be a neighbor of another cutpoint by Lemma 3.4.4, so $z \in I$ and $z = i$. Hence if v has a neighbor in $C(I)$ then that neighbor is i. As i is the only element of I with $\Pi_j(i) = \Pi_j(v)$ for all $j \in J$, i is indeed a neighbor of v in $C(I)$. $\square$

Now we want to define P_v for $v \in V_N$. By Lemma 3.6.11 there is a unique $i \in I$ which is a neighbor of v in C. For every $j \in J$ let u_j be the predecessor and w_j the successor of $\Pi_j(i)$ in $C(I_j)$. Let $z \in V_F$ be a cutpoint such that for all sufficiently large $j \in J$ both $S(z, w_j)$ and $S(z, u_j)$ distinguish two elements of $\mathcal{P}$. In particular, for all sufficiently large indices j, P_z is disjoint from P_{u_j} and P_{w_j} and thus P_z is disjoint from $V(u_j, w_j)$.

If v is the predecessor of i, then let (Y,Z) be the supremum of the set $\{S(z,u) \colon u \in (z,v) \cap V_F\}$ and otherwise let (Y,Z) be the infimum of the set $\{S(z,w) \colon w \in (v,z) \cap V_F\}$. In both cases (Y,Z) has order at most k because the

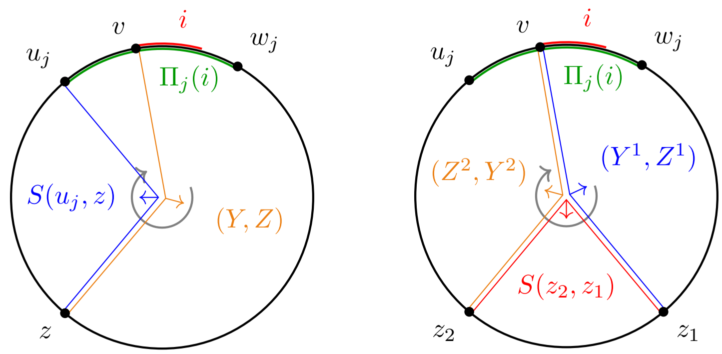

Figure 12: To the left: In the case that $v \in V_N$ is the predecessor of its neighbor i in $C(I)$, the set P_v is defined via the limit of the separations $S(z, u_j)$ for some suitable $z \in V_F$.
To the right: P_v does not depend on the choice of z, see also Lemma 3.6.15.

order function is limit-closed. Define $P_v := (Y \cap Z) \setminus (P_z \cup X)$. As $S(z, u_j) \leq (Y, Z) \leq S(z, w_j)$ for all sufficiently large $j \in J$, $Y \cap Z$ contains P_z. Also (Y, Z) distinguishes two elements of $\mathcal{P}$ so it has order k. Hence P_v has $(k - |X|)/2$ many elements.

Lemma 3.6.12. *$P_v \subseteq V(u_j, w_j)$ for all $j \in J$.*

Proof. It suffices to show the claim for all $j \in J$ such that z is a cutpoint of $C(I_j)$. As $S(z, u_j) \leq (Y, Z) \leq S(z, w_j)$ by definition, $Y \cap Z$ is contained in $V(z, w_j) \cap V(u_j, z)$, which equals $V(u_j, w_j) \cup P_z$ by Lemma 3.4.9. As P_v is a subset of $Y \cap Z$ and is disjoint from P_z, the lemma holds. □

Corollary 3.6.13. $P_v \subseteq P_i$.

Proof. As every Φ_j is a k-pseudoflower, for sufficiently large $j \in J$ we have

$$V(u_j, w_j) = X \cup P_{\Pi_j(i)} \cup P_{u_j} \cup P_{w_j} = X \cup P_{\Pi_j(i)}.$$

Because P_v is disjoint from X, this implies $P_v \subseteq \bigcap_{j \in J} P_{\Pi_j(i)} = P_i$. □

Corollary 3.6.14. *For $i' \in I$, if u is a neighbor of i' in $C(I)$ then $P_u \subseteq P_{i'}$.*

Proof. If $u \in V_N$, then i' is the unique neighbor of u and $P_u \subseteq P_{i'}$ by Corollary 3.6.13. If $u \in V_F$, then $P_u \subseteq P_{\Pi_j(i')}$ for all $j \in J$ for which u is a cutpoint, also implying $P_u \subseteq P_{i'}$. □

Lemma 3.6.15. *The set P_v does not depend on the choice of z.*

Proof. Let z_1 and z_2 be two possible choices for z, assume $z_1 \in [v, z_2]$. Denote the sets defined by z_n by a superscript index n. This proof is for the case that v is the predecessor of i, the proof of the other case is symmetric. Let l be sufficiently large that both $S(z_1, w_l)$ and $S(z_2, u_l)$ are defined in Φ_l and distinguish two elements of $\mathcal{P}$.

As $S(z_1, u_j) \geq S(z_2, u_j)$ for all $j \geq l$, also $(Y^1, Z^1) \geq (Y^2, Z^2)$. Furthermore for all $j \geq l$,

$$S(z_1, u_l) \vee (Y^2, Z^2) \geq S(z_1, u_l) \vee S(z_2, u_j) = S(z_1, u_j).$$

So $(Y^1, Z^1) \leq S(z_1, u_l) \vee (Y^2, Z^2) \leq (Y^1, Z^1)$ and thus the two separations are equal. In particular $P_v^1 \subseteq Z^1 \subseteq Z^2$ and $P_v^1 \subseteq Y^2 \cup V(z_1, u_l)$. Because also $P_v^1 \subseteq V(u_l, w_l)$, together we have

$$P_v^1 \subseteq V(u_l, w_l) \cap (Y^2 \cup V(z_1, u_l)) = (V(u_l, w_l) \cap Y^2) \cup (X \cup P_{u_l}) \subseteq Y^2$$

where the equality holds by Lemma 3.4.9. Hence $P_v^1 \subseteq Y^2 \cap Z^2$. Also $P_v^1 \subseteq V(u_j, w_j)$ for all sufficiently large $j \in J$ implies that $P_v^1 \cap V(z_1, z_2) = \emptyset$ and thus $P_v^1 = P_v^2$ as those two sets have the same size. $\square$

Lemma 3.6.16. $(Y, Z) = (V'(z, v), V'(v, z))$.

Proof. We are going to show the lemma in the case that v is the predecessor of i, the other case is symmetric.

We have $V'(z, v) = P_v \cup \bigcup_{t \in [z,v] \cap V_F} V(z, t) \subseteq Y$. Assume there is $u \in Y \setminus V'(z, v)$, then $u \notin V(z, u_j)$ for all sufficiently large $j \in J$. So $u \in Z$ and thus $u \in P_z \cup P_v \cup X$, a contradiction. So $Y = V'(v, z)$.

We have $V'(v, z) \subseteq V'(u_j, z) = V(u_j, z)$ for all sufficiently large $j \in J$ and thus $V'(v, z) \subseteq Z$. Let $u \in Z \setminus (X \cup P_v)$ and let $w \in I \cup V_F$ be some element such that $u \in P_w$. If $w \in [v, z]$, then also $u \in V'(v, z)$. Otherwise $w \in [z, v]$, so $w \in [z, u_j]$ for some $j \in J$ and thus $u \in V(z, u_j) \subseteq Y$. In this case we have $u \in Y \cap Z = P_v \cup P_z \cup X \subseteq V'(v, z)$. Thus $Z \subseteq V'(v, z)$, and these two sets are equal. $\square$

Corollary 3.6.17. *$(V'(z, v) \cup P_v, V'(v, z) \cup P_v)$ is a separation with separator $P_v \cup P_z \cup X$.*

3.6.4 Interval separations are indeed k-separations

The goal of this section is to show that the interval separations of the limit we are constructing have the properties required in the definition of pseudoflowers.

Lemma 3.6.18. *For all distinct $v, w \in V_F \cup V_N$ we have $V'(v, w) = \hat{V}(v, w)$.*

Proof. $V'(v, w) \subseteq \hat{V}(v, w)$ is clear by definition of $V'(v, w)$. In order to show the reverse inclusion, let t be an element of $[v, w]$. and $u \in P_t$. If $t \in \{v, w\}$ or $t \in I \cup V_F$, then clearly $P_t \subseteq V'(v, w) \cup P_v \cup P_w$. Otherwise t has a neighbor $i \in I$, and $i \in [v, w]$. Then by Corollary 3.6.13, $P_v \subseteq P_i \subseteq V'(v, w)$. □

Lemma 3.6.19. *(The cyclic order is illustrated in Fig. 6). Let a, b, c and d be elements of $V_F \cup V_N$ such that*

- *$b \in [a, c]$ and $d \in [c, a]$*
- *$(V'(a, c), V'(c, a))$ is a separation with separator $X \cup P_a \cup P_c$*
- *$(V'(b, d), V'(d, b))$ is a separation with separator $P_b \cup P_d \cup X$.*
- *$P_a \cap P_d = \emptyset$.*

Then $(V'(b, c), V'(c, b))$ is a separation with separator $P_b \cup P_c \cup X$.

Proof. We show that

$$(V'(b, c), V'(c, b)) = (V'(a, c), V'(c, a)) \wedge (V'(b, d), V'(d, b)).$$

$V'(c, b) = V'(c, a) \cup V'(d, b)$ and $V'(b, c) \subseteq V'(a, c) \cap V'(b, d)$ clearly hold. Let $u \in V'(c, b) \setminus (P_b \cup P_c \cup X)$. Because $P_a \cap P_d$ is empty, either $u \in V'(c, a) \setminus (P_a \cup P_c \cup X)$ or $u \in V'(d, b) \setminus (P_b \cup P_d \cup X)$. In the first case, because $(V'(a, c), V'(c, a))$ is a separation with separator $P_a \cup P_c \cup X$, u is an element of $V'(c, a) \setminus V'(a, c)$. In the second case we similarly get $u \in V'(d, b) \setminus V'(b, d)$. Hence in both cases $u \notin V'(a, c) \cap V'(b, d)$, showing that $V'(a, c) \cap V'(b, d) = V'(b, c)$. Furthermore, $V'(b, c) \cap V'(c, b) \subseteq P_b \cup P_c \cup X$, and so the inclusion must be an equality. □

Lemma 3.6.20. *For all distinct $v, w \in V_F \cup V_N$ the pair $(\hat{V}(v, w), \hat{V}(w, v))$ is a separation with separator $X \cup P_v \cup P_w$.*

Proof. By Lemma 3.6.18 it suffices to show that $(V'(v, w), V'(w, v))$ is a separation with the correct separator for all distinct elements v and w of $V_F \cup V_N$. If both v and w are contained in V_F then this is true by Corollary 3.6.10. Consider first the case that exactly one is contained in V_F, by switching the names we may assume $w \in V_F$. If w is a suitable candidate for z in the definition of P_v, then $(V'(w,v), V'(v,w))$ is a separation with separator $P_v \cup P_w \cup X$ by Corollary 3.6.17. So assume otherwise. We will consider the case that for all sufficiently large $j \in J$, $S(w, u_j)$ does not distinguish elements of $\mathcal{P}$, the case where $S(w, w_j)$ does not distinguish elements of $\mathcal{P}$ is symmetric.

Let $z \in V_F$ from which P_v might have been defined. Then by choice of z there is $t \in V_F$ such that $S(t,z)$ distinguishes elements of $\mathcal{P}$ and such that $t \in [v,z]$ if $w \in [z,v]$ and $t \in [z,v]$ if $w \in [v,z]$. Also $(V'(z,v), V'(v,z))$ is a separation with separator $P_z \cup P_v \cup X$ by Corollary 3.6.17 and we already saw that $(V'(w,t), V'(t,w))$ is a separation with separator $P_w \cup P_t \cup X$. Furthermore $S(t,z) = (V'(t,z), V'(z,t))$, and because $S(t,z)$ distinguishes elements of $\mathcal{P}$ and thus has connectivity k this implies $P_t \cap P_z = \emptyset$. Now we can apply Lemma 3.6.19. If $t \in [v,z]$ then we apply Lemma 3.6.19 for $a = z$, $b = w$, $c = v$ and $d = t$. Otherwise we apply the lemma for $a = t$, $b = v$, $c = w$ and $d = z$. In both cases $(V'(w,v), V'(v,w))$ is a separation with separator $P_v \cup P_w \cup X$.

Now assume that both v and w are contained in V_N. Because every Φ_j distinguishes at least three elements of $\mathcal{P}$, by swapping the names of v and w if necessary we may assume that there are t and u in V_F such that $t \in [v,w]$ and $u \in [t,w]$ and $S(t,u)$ distinguishes elements of $\mathcal{P}$. We already showed that $(V'(v,u), V'(u,v))$ and $(V'(t,w), V'(w,t))$ are k-separations with separator $P_v \cup P_u \cup X$ and $P_t \cup P_w \cup X$ respectively. Furthermore $P_t \cap P_u = \emptyset$ because $S(t,u)$ distinguishes elements of $\mathcal{P}$, thus we may apply Lemma 3.6.19 to $a = u$, $b = w$, $c = v$ and $t = d$ and we are done. $\square$

3.6.5 Deletion of redundant petals

Φ is now nearly a k-pseudoflower, the only property missing is $(*)$. In order to fix that, we are going to delete troublesome elements from I.

Let $\mathcal{Y}$ be the set of subsets of V of size $(k - |X|)/2$ that are of the form P_i for some $i \in I$. In order to ensure that $(*)$ holds, it suffices to delete, for every $Y \in \mathcal{Y}$, all but one index $i \in I$ with $P_i = Y$. For that, we first define one element we want to keep: If there is $j \in J$ and $i_j \in I_j$ whose petal is Y, then we want to define i_Y to be an element of I with $\pi_j(i_Y) = i_j$. If that does not exist, then we pick i_Y arbitrarily with $P_{i_Y} = Y$.

Actually if i_Y is defined via the first case, then i_Y is uniquely determined:

Lemma 3.6.21. *Let $Y \in \mathcal{Y}$ and i_j be an index element of some Φ_j whose petal* *is Y. Then there is a unique $i \in I$ with $\pi_j(i) = i_j$, and if $P_{i_l} = Y$ for some $l \in J$ and $i_l \in I_l$, then $\pi_l(i) = i_l$.*

Proof. Let $l \in J$ with $j \leq l$. By Lemma 3.4.16 there is some $i_l \in I_l$ that is mapped to i_j by all witnesses of $\Phi_j \leq \Phi_l$, in particular $F_{lj}(i_l) = i_j$. Thus P_{i_l} is contained in P_{i_j} and has size at least $(k - |X|)/2$. But P_{i_j} also has size $(k - |X|)/2$, so $P_{i_l} = P_{i_j} = Y$. By ($*$), i_l is the only element of I_l with $P_{i_l} = Y$, and thus i_l is the only element of I_l with $F_{lj}(i_l) = i_j$.

So if i and i' are elements of I with $\pi_j(i) = \pi_j(i') = i_j$, then $\pi_l(i) = \pi_l(i')$ for all $l \geq j$, implying that $i = i'$. By construction of I there is some $i \in I$ with $\pi_j(i) = i_j$, thus there is a unique such element. $\square$

The previous lemma implies that if there is some index i_j of some Φ_j with $P_{i_j} = Y$, then there is $i \in I$ such that all indices of any Φ_j whose petal is Y are the image of i under π_j.

Obtain I' from I by deleting, for every element Y of $\mathcal{Y}$, all $i \in I - i_Y$ with $P_i = Y$. Obtain C from $C(I)$ by deleting $I \setminus I'$ and, for all $i \in I \setminus I'$, the predecessor of i in $C(I)$.

The following lemma implies that the elements of I' are close to being "dense" in $C(I)$:

Lemma 3.6.22. *Let w and w' be distinct elements of V_F such that for some $Y \in \mathcal{Y}$, $P_z = Y$ for all $z \in [w, w'] \setminus I$. Then for some $i' \in [w, w']$ there is $j \in J$ with $P_{\pi_j(i')} = Y$ or there is $y \in P_{i'}$ such that $y \notin P_{i''}$ for all $i'' \in I - i'$.*

Proof. Let $j \in J$ be large enough that $C(I_j)$ contains both w and w'. Then there is $i_j \in [w, w'] \cap I_j$, where the interval is taken in $C(I_j)$, and some $i' \in I$ with $\pi_j(i') = i_j$. For the predecessor p and the successor s of i_j in $C(I_j)$, $P_s = P_p = Y$. So $P_{i'}$ is a subset of P_{i_j}. If $P_{i_j} = Y$, then the lemma holds, so assume otherwise. Let y be an element of P_{i_j} that is not contained in Y. Then s and p are contained in $[w, w']$ and $V(p, s)$ does not depend on whether it is taken in $(P_z)_{z \in C(I)}$ or in Φ_j. Hence there is $z \in [w, w']$ such that $y \in P_z$, and z has to be contained in I. So $y \in P_z$ for some $z \in [w, w'] \cap I$, implying that the predecessor p' and successor s' of z in $C(I)$ both have Y as vertex set. Then the separating set of the separation $S(z)$ is Y, implying that y is only contained in P_z and in no $P_{i''}$ with $i'' \in I - z$. $\square$

So every element of $C(I) \setminus C$ is close to some element of C:

Lemma 3.6.23. *Let $z \in C(I) \setminus C$. Then there is $v \in C \setminus I'$ such that $[z, v]$ is finite, no element of $[z, v)$ is contained in C, and $Pz' = P_v$ for all $z' \in [z, v)$.*

Proof. If $z \in I$, then let $i_0 = z$, otherwise let i_0 be the successor of z. Recursively, if i_l is defined and has a successor in I that is not contained in I', then define i_{l+1} to be that successor.

Assume for a contradiction that i_2 is defined. For any distinct $i, i' \in I$, some projection π_j maps i and i' to different images. So there is $w \in V_F$ such that $w \in [i, i']$. Thus the common neighbor w_0 of i_0 and i_1, and the common neighbor w_1 of i_1 and i_2, are both contained in V_F. Now the previous lemma implies that i_1 is contained in I', a contradiction.

If i_1 is defined, then let v be the successor of i_1 in $C(I)$. Otherwise let v be the successor of i_0 in $C(I)$. Thus, if v is the predecessor of some element of I, then that is also contained in I'. For $i \in I \setminus I'$, and the predecessor p and successor s of i in $C(I)$, $P_p = P_i = P_s$. In particular $P_z = P_{i_0} = P_v$, and if i_1 is defined then also $P_v = P_{i_1} = P_{w_0}$ for the common neighbor w_0 of i_0 and i_1 in $C(I)$. □

Lemma 3.6.24. *C is a cycle completion of I'.*

Proof. In order to show that the identity on I' can be extended uniquely to a bijective monotone map from C to $C(I')$ it suffices to show that all non-trivial intervals of I' can be uniquely written as $[v, w] \cap I'$ for elements v and w of $C \setminus I'$. In order to do so, let I'' be a non-trivial interval of I', and let $i \in I' \setminus I''$. Let $\hat{I}$ be the set of all $\hat{i}$ for which there are $i_1, i_2 \in I''$ such that the interval $[i_1, i_2]$, taken in I, contains $\hat{i}$ but not i. Then $\hat{I}$ is an interval of I that contains I'' as a subset and that does not contain i. So $\hat{I}$ can be written as $[v, w] \cap I$ for unique elements v and w of $C(I) \setminus I$. Also, $I'' = \hat{I} \cap I' = [v, w] \cap I'$.

In order to find v and w that are contained in C, first consider v. If v is the predecessor of some $i' \in I$, then i' is contained in I'' and thus in I'. So v cannot be the predecessor of an element of $I \setminus I'$, implying that v is contained in C. But w could indeed be contained in $C(I) \setminus C$. If it is, then apply Lemma 3.6.23 to w and obtain w'. Otherwise let $w = w'$. Then $\hat{I} = [v, w] \cap I' = [v, w'] \cap I'$.

In order to show that v and w' are unique, it suffices to show for any distinct $v, w \in C \setminus I'$ that $[v, w] \cap I'$ is non-empty. As in the proof of Lemma 3.6.23, if $[v, w]$ contains enough elements (at least 7) then it contains two elements of V_F and thus an element of I'. Otherwise, $[v, w]$ is finite and v is the predecessor in $C(I)$ of some $i \in [v, w] \cap I$. Because $v \in C$, this implies that $i \in I'$. □

Let $\tilde{F} : C(I') \to C$ be a bijective monotone map whose restriction to I' is the identity. This map then identifies $C(I')$ with C. Denote the family $(P_z)_{z \in C(I')}$ where each P_z is equal to $P_{\tilde{F}(z)}$, by Φ'. As Φ satisfies all properties of a k-pseudoflower with the possible exception of ($*$), also Φ' satisfies all properties of a k-pseudoflower with the possible exception of ($*$).

Lemma 3.6.25. *For all elements v and w of $C \setminus I'$ we have*

$$\bigcup_{z \in [v,w]_C} P_z = \bigcup_{z \in [v,w]_{C(I)}} P_z$$

where the intervals are taken in C and $C(I)$, respectively.

Proof. It is clear that the left-hand side is a subset of the right-hand side. In order to show the other direction, let $u \in [v,w]_{C(I)}$. If $u \in C$ then $u \in [v,w]_C$, so assume otherwise. Apply Lemma 3.6.23 to u and obtain $z' \in C \setminus I'$. Then $w \notin [u, z')$, as $w \in C$, so $z' \in [v,w]$. As $P_u = P_{z'}$, the lemma holds. □

Lemma 3.6.26. *Φ' is a k-pseudoflower such that $\Phi_j \leq \Phi'$ for all $j \in J$.*

Proof. Recall that by $\tilde{F}$, $C(I')$ is identified with C. For all $v, w \in C(I') \setminus I'$,

$$\begin{aligned} \bigcup_{z \in [v,w]_{C(I')}} P_z &= \bigcup_{z \in [v,w]_C} P_z && \text{by Lemma 3.6.25} \\ &= V'(v,w) \setminus X && \text{by Lemma 3.6.18} \end{aligned}$$

so by Lemma 3.6.20 $S(v,w)$ taken in Φ' is a separation with separator $P_v \cup P_w \cup X$ and also $X = V \setminus \bigcup_{z \in C(I')} P_z$. As every P_v with $v \in C(I) \setminus I$ has size $(k - |X|)/2$, this is also true for every $v \in C(I') \setminus I'$. Hence every $S(v,w)$ has order at most k. If some $i \in I'$ has a neighbor v in $C(I')$, then $P_v \subseteq P_i$ by Corollary 3.6.14. And by definition of I', Φ' satisfies ($*$). Thus Φ' is a k-pseudoflower.

In order to show that $\Phi_j \leq \Phi'$ for all $j \in J$, we first show that the restriction of π_j to I' is surjective. For that, let $i_j \in I_j$ and let p and s be the predecessor and successor of i_j in $C(I_j)$. Then p and s are contained in V_F, and thus by Lemma 3.6.22 there is $i \in [v,w] \cap I'$, and $\pi_j(i) = i_j$. So the restriction of π_j to I' with codomain I_j is a surjective monotone map, and there is a unique monotone surjective extension $F_j : C(I') \to C(I_j)$ with $F_j(I') = (I_j)$.

Let v and w be distinct elements of $C(I_j) \setminus I_j$. If $v \in C$, then let $v' = v$. Otherwise there is, by Lemma 3.6.23, an element $v' \in C \setminus I'$ such that $[v, v'] \cap C =$

v' and $P_v = P_{v'}$. Define w' similarly. Then

$$V(v,w) = \hat{V}(v,w) = \hat{V}(v',w') = V(\tilde{F}^{-1}(v'), \tilde{F}^{-1}(w'))$$

so F_j witnesses that $\Phi_j \leq \Phi'$. $\square$

So in this section it was shown that if $(\Phi_j)_{j\in J}$ is a $\leq$-chain of k-pseudoflowers which distinguish at least three elements of $\mathcal{P}$ that they locate, then there is a k-pseudoflower which is an upper bound of the chain $(\Phi_j)_{j\in J}$. In particular, if Ψ is a k-pseudoflower which distinguishes at least three elements of $\mathcal{P}$ that it locates, then the set of k-pseudoflowers Φ' with $\Psi \leq \Phi'$ every $\leq$-chain has an upper bound. Thus the following theorem follows by Zorn's Lemma:

Theorem 3.6.27. *Let Φ be a k-pseudoflower which distinguishes at least three elements of $\mathcal{P}$ that it locates. Then there is a $\leq$-maximal k-pseudoflower Ψ such that $\Phi \leq \Psi$.*

In this theorem, the condition that Φ distinguishes at least three elements of $\mathcal{P}$ that it locates can be weakened to Φ distinguishing any three profiles or tangles that it locates. As every k-pseudoflower locates every k-tangle, and every k-pseudoflower extending a k-flower with four petals locates every k-profile, the theorem can be specialized to the following versions:

Theorem 3.6.28. *Let Φ be a k-pseudoflower which distinguishes at least three k-profiles and extends a k-flower with four petals. Then there is a $\leq$-maximal k-pseudoflower Ψ such that $\Phi \leq \Psi$.*

Theorem 3.6.29. *Let Φ be a k-pseudoflower which distinguishes at least three k-tangles. Then there is a $\leq$-maximal k-pseudoflower Ψ such that $\Phi \leq \Psi$.*

3.6.6 Existence of $\preccurlyeq$-maximal pseudoflowers

Call an element P of $\mathcal{P}$ *closed* if whenever $(S_j)_{j\in J}$ is a chain of elements of P, then its supremum (which exists because their order is bounded) is contained in $\mathcal{P}$. If $\mathcal{P}$ only contains closed elements, then $\leq$-maximal k-pseudodaisies are also $\preccurlyeq$-maximal, and thus there are $\preccurlyeq$-maximal k-pseudodaisies.

Theorem 3.6.30. *Let Φ be a k-pseudoflower that is $\leq$-maximal and distinguishes three elements of $\mathcal{P}$ that it locates. If all elements of $\mathcal{P}$ are closed, then Φ is also $\preccurlyeq$-maximal.*

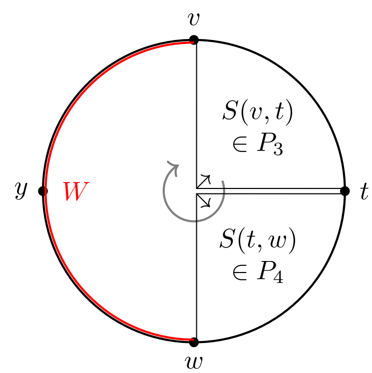

Figure 13: Some of the notation used in the proof of Theorem 3.6.30.

Proof. See Fig. 13 for a depiction of some of the notation. Let Φ be a k-pseudoflower that distinguishes three elements of $\mathcal{P}$ that it locates and which is not $\preccurlyeq$-maximal among all k-pseudoflowers, as witnessed by Φ'. Let P_1 and P_2 be two elements of $\mathcal{P}$ which are distinguished by Φ' but not by Φ.

By Lemma 3.4.20 there is a cutpoint $v \in C(I) \setminus I$ such that either $S(v,w) \in P_1$ for all $w \in C(I) \setminus I - v$ or $S(w,v) \in P_1$ for all $w \in C(I) \setminus I - v$. We will assume that $S(v,w) \in P_1$ for all $w \in C(I) \setminus I - v$, the other case is symmetric. Because Φ distinguishes at least three elements of $\mathcal{P}$ there is an interval separation $S(v,w)$ of Φ whose inverse is contained in two elements P_3 and P_4 of $\mathcal{P}$ which are distinguished and located by Φ. As Φ distinguishes and locates P_3 and P_4, and distinguishes them from P_1, there is $t \in]v,w[$ such that $S(v,t)$ distinguishes P_3 and P_4. By swapping the names of P_3 and P_4 if necessary, we may assume that $S(t,v) \in P_3$. Again by the profile property also $S(t,w)$ distinguishes P_3 and P_4 with $S(w,t) \in P_4$.

The k-pseudoflower Φ' distinguishes P_1, P_2, P_3 and P_4 pairwise, so there is an interval separation (C,D) of Φ' which is contained in P_3 but not P_4 and which distinguishes P_1 and P_2. By swapping the names of P_1 and P_2 if necessary we may assume that (C,D) is contained in P_1 and P_3 and that its inverse is contained in P_2 and P_4.

If v has a predecessor v' in $C(I) \setminus I$, then (C,D) properly crosses $S(v,v')$ and by Lemmas 3.5.1 and 3.5.3 there is an extension Φ'' of Φ which distinguishes P_1 and P_2. Assume for a contradiction that v has no predecessor in $C(I) \setminus I$. Let W be the interval $(w,v) \setminus I$. If Φ is a k-pseudoanemone, then $S(t,v)$ is the supremum of the separations $S(t,x)$ with $x \in W$. Because P_1 is closed and contains all $S(t, x)$ with $x \in W$ but not $S(t, v)$, Φ cannot be a k-pseudoanemone and thus has to be a k-pseudodaisy.

For all $x \in W$ let S_x be the separation $((C, D) \vee S(t, x)) \wedge S(t, v)$. All three k-separations (C, D), $S(t, x)$ and $S(t, v)$ are contained in P_3, not contained in P_4 and have order at most k, so S_x has order at most k. Denote the unique supremum of $(S_x)_{x \in W}$, which exists by limit-closedness, by (A, B). As the union of all sets $V(v, x)$ with $x \in W$ is the whole ground set V, there is some $y \in W$ such that $V(v, y)$ contains $A \cap B$. But all profiles are closed, so (A, B) is still contained in P_1 and thus properly crosses $S(v, y)$. Hence applying Lemma 3.5.2 to the concatenation of Φ on vertices v, t and y shows that $A \cap B$ contains a vertex not in $V(v, y)$, a contradiction to the choice of y. $\square$

Theorem 3.6.31. *Let Φ be a k-pseudodaisy that distinguishes and locates at least three elements of $\mathcal{P}$. If every element of $\mathcal{P}$ is closed, then there is a $\preccurlyeq$-maximal k-pseudoflower Ψ with $\Phi \leq \Psi$.*

3.7 Flower tree sets

Now that we have shown the existence of maximal flowers, we want to examine how they can help in constructing decompositions. Recall that earlier in this chapter we introduced a function $f_{\mathcal{P}}$ and then proceeded in two steps: first we chose an abstract tree set in the image of $f_{\mathcal{P}}$ and then we selected from the fibers of $f_{\mathcal{P}}$ a concrete set of separations.

In this section, whenever we speak of a maximal flower, this is meant with respect to $\preccurlyeq$. Using our results we could now simplify the first step by starting with the set X of all relevant separations occurring as a petal separation of a maximal k-pseudoflower.

Lemma 3.7.1. *$f_{\mathcal{P}}(X)$ is a tree set which separates any different $P, Q \in \mathcal{P}$.*

Proof. Let (A, B) and (C, D) be separations of X. There is a maximal flower Φ which has (A, B) as a petal separation. If (C, D) properly crossed (A, B), this would contradict the maximality of Φ by Lemma 3.5.3. Thus $f_{\mathcal{P}}(A, B)$ is nested with $f_{\mathcal{P}}(C, D)$. Thus $f_{\mathcal{P}}$ is a tree set.

Furthermore for different $P, Q \in \mathcal{P}$ there is some (A, B) distinguishing them. Then (A, B) is displayed by some maximal flower Φ. Since P and Q are closed, there are petal separations $(C, D) \in P$ and $(E, F) \in Q$. These must be different, however, since otherwise P and Q could not orient (A, B) differently. Thus the image of (C, D) under $f_{\mathcal{P}}$ separates P and Q. $\square$

Furthermore maximal flowers can also be located in the resulting tree set. Given a maximal flower Φ the complements of $f_{\mathcal{P}}(A, B)$ for all petal separations (A, B) of Φ extend to a unique consistent orientation, since every $P \in \mathcal{P}$ lives in some petal. This also shows that there is no profile in the intersection of the right sides of o. While in the finite cases, one can show that each node of the tree set constructed corresponds to either one of our profiles or a flower, this is not quite true for us, since there may be limit nodes which contain profiles which are not closed and thus not contained in $\mathcal{P}$

3.8 Outlook

One natural question to ask is whether all three conditions in Theorem 3.2.6, namely closure, orderliness and regularity, are necessary. Dropping closure would clearly require a completely different approach if it is possible, the same is probably also true for orderliness, since Lemma 3.2.3 will not work without it. In [22] we will show that using the notion of essential core from [13] it is possible to drop regularity.

A similar question can be asked for Theorem 3.3.1. Here, however, it is unclear whether the theorem holds without regularity since taking the essential core does not preserve the property of being a universe.

There are also two additional properties which one might wish for from a tree set as in Theorem 3.3.1. The first is being minimally distinguishing. In the finite case this is very easy to achieve by simply dropping unnecessary separations, but doing this in the infinite case could easily leave no separations at all. The second property is canonicity, that is invariance under isomorphisms, discussed for example in [15]. Our current proof strategy does not establish canonicity, since the consistent orientation in each step is chosen canonically.

As we have seen in Section 3.7, we cannot quite achieve a tree set in which every node contains either one of our profiles or displays a flower, analogously to [10]. Perhaps, however, one might show that each of the nodes containing neither contains a non-closed profile which is a limit of our set of profiles

Chapter 4

Ubiquity

4.1 Introduction

A graph G is called ubiquitous (with respect to subgraphs) if every graph containing arbitrarily many disjoint copies of G also contains infinitely many disjoint copies of G. This notion was introduced by Halin in [31], who proved that rays have this property. However, this is not particularly common, so interest has recently focused more on ubiquity with respect to topological minors and minors (defined analogously), for instance in a series of articles starting with [6].

Ubiquity, however, is not only meaningful in graphs, but also, for instance, in topological spaces and by extension path spaces. In these settings, it becomes more natural to slightly reformulate the definition of ubiquity. Instead of talking about ubiquity with respect to a specific relation, one can call a class $\mathcal{C}$ of path spaces ubiquitous if any path space containing arbitrarily many disjoint elements of $\mathcal{C}$ contains infinitely many disjoint elements of $\mathcal{C}$.

Many ubiquity proofs are like Halin's original proof in [31] in that they build up copies step-by-step during a recursive process. However, this approach is problematic in path spaces since we might end up in the same limit point during the construction of two supposedly disjoint copies. To illustrate the new method we use instead, in Section 4.2 we first show an easier result implied by our main theorem, namely the ubiquity of certain theta-like spaces, which we call ladders. In particular this gives a short proof that S^1 is ubiquitous in topological spaces.

The first step of our proof method is somewhat reminiscent of [6, Lemma 5.4]:

we recursively delete all ladders whose deletion still leaves infinitely many ladders. If this does not stop, we are done. Otherwise we fix a set of three ladders, such that deleting any of them leaves only some finite number r of ladders. But taking some other huge set of disjoint ladders, we show that between two of our three ladders we can construct a set of more than r disjoint ladders not meeting the third.

Next we will extend this proof method to prove the main theorem of this chapter, which we state here in a slightly weaker version.

Theorem 4.1.1. *Let G be a finite, planar graph. Then the class of IG is ubiquitous.*

Section 4.3 gives some prepatory lemmas for Theorem 4.1.1, before we complete the proof in Section 4.4.

Ubiquity is also sometimes considered with respect to edge-disjointness. While this is not immediately sensible for path spaces, in Section 4.5 we define a notion of edge-ubiquity and show that cycles have this property. In Section 4.6 we give some counterexamples for topological minors. In particular we show that the class of TK^5 is not ubiquitous and that there is some tree G such that TG is not ubiquitous.

4.2 Ladders

In this section we will prove the ubiquity of a special class of path spaces called s-ladders. While this is implied by the main theorem of this chapter, it is instructive to isolate the main idea of our argument from the more complicated details that arise in the general case. We also remark that the proofs in this section do not use the full properties of path spaces: instead of compatibility it is enough to require that every path P has a segment that is an A-B-path whenever A, B are paths meeting P.

We start by observing a Ramsey-type property of linear orders which we will apply quite a few times to clean up large sets of paths, possibly interacting badly. Let $k \leq \aleph_0$ and l a positive natural number. An *l-system* is a set of l linear orders on the same set, which we call its *ground set*. The *size* of an l-system is the cardinality of its ground set. In this section L will always be an l-system and X its ground set. A set is *sorted in L* if its order is the same up to reversing in any element of L. For any $k \leq \aleph_0$ and finite l let $R(k, l)$ be the smallest cardinal

$\kappa \leq \aleph_0$ such that for any l-system L of size at least κ there exists some set of size k sorted in L, if such a cardinal exists.

Let $g(k) = (k-1)^2 + 1$. Then the following is a consequence of the Erdős–Szekeres theorem for finite k and follows by a well-known Ramsey-type argument for $k = \aleph_0$.

Lemma 4.2.1. $R(k, 2) = g(k)$

We can use this to deduce bounds on our Ramsey numbers.

Proposition 4.2.2. $R(k, l) \leq g^{l-1}(k)$

Proof. Clearly, $R(k, 1) = k$. For $l \geq 2$ we prove this assertion by induction on l for all k simultaneously. The base case is Lemma 4.2.1. Now let L be some l-system for $l > 2$. Let $a \in L$ be arbitrary. Since $R(g(k), l-1) \leq g^{l-2}(g(k)) = g^{l-1}(k)$, by the induction hypothesis there exists some set Y of size $g(k)$ sorted in $L - a$. Let b be induced by an element of $L - a$ on Y. Using the induction hypothesis once more, we obtain a subset Z of Y of size k sorted in $\{a, b\}$. Since Y was sorted in $L - a$, the set Z is sorted in L. □

In particular, all these $R(k, l)$ exist.

We now define s-ladders for a given positive natural number s; refer to Fig. 14 for an example. Let P and Q be disjoint paths, let p and q be s-tuples consisting of different points in order on P and Q respectively and for every $1 \leq i \leq s$ let L_i be a P-Q-path with endpoints p_i and q_i. If the L_i are all disjoint, we call the set $\{L_i; 1 \leq i \leq s\} \cup \{P, Q\}$ an s-ladder. The L_i are called its *rungs* and P and Q its *sides*. A *part* of an s-ladder is either a rung or a side.

Next we apply the discussion of linear orders above to prove what will turn out to be the key lemma of our proof.

Lemma 4.2.3. *Let $\mathcal{P}$ be a path space, $P, Q \in \mathcal{P}$ and $\mathcal{Q}$ a set of at least $R(sk, 2)$ disjoint P-Q-paths. Then $\mathcal{P}$ contains a set of k disjoint s-ladders.*

Proof. Let P' and Q' be linear orders on $\mathcal{Q}$ induced by their endpoints in P and Q respectively. Since $|\mathcal{Q}| \geq R(sk, 2)$ there is a set Y of size sk sorted in P' and Q'. List the elements of Y as $(y_i)_{0<i<sk}$. For each $1 < j < k$ we obtain a ladder L_j whose rungs are the paths in Y and whose sides are the segments defined by the endpoints of y_{sj} and $y_{(s+1)j-1}$ in P and Q respectively. Clearly, all the L_j are disjoint. □

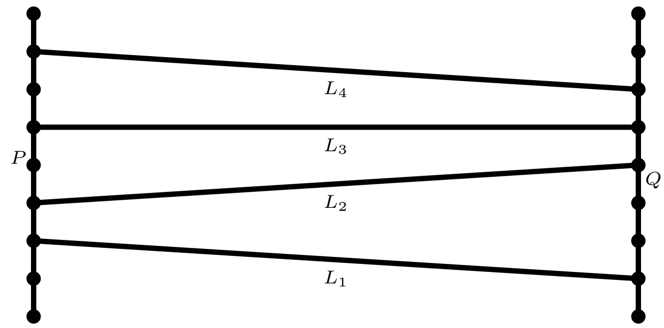

Figure 14: An example of a 4-ladder.

We call an s-ladder *redundant* if there are arbitrarily many disjoint s-ladders in $\mathcal{P}$ disjoint from it.

Theorem 4.2.4. *s-ladders are ubiquitous in path spaces.*

Proof. Assume that $\mathcal{P}$ contains arbitrarily many disjoint s-ladders. Recursively delete any redundant s-ladder in $\mathcal{P}$ as long as there is one. It is enough to show that this never terminates, since then we will have found an infinite set of disjoint s-ladders. So we assume for a contradiction that the process terminates after only finitely many steps, without loss of generality we may then assume that $\mathcal{P}$ has no redundant s-ladders.

Let L_1,L_2 and L_3 be three disjoint s-ladders in $\mathcal{P}$. Since none of these is redundant there exists some r such that there is no set of more than r disjoint s-ladders each of which avoids at least one of L_1, L_2, L_3. Let D be a set of $2(s+2)^2R(s(r+1),2)+r$ disjoint s-ladders. Deleting any elements which do not meet at least one of L_1, L_2 and L_3 leaves a set D' of $2(s+2)^2R(s(r+1),2)$ disjoint s-ladders meeting any of these. Each element of D' then has a segment that is an L_1-$L_2 \cup L_3$-path, let $\mathcal{Q}$ contain a choice of one such path for each element. Since any s-ladder has $s+2$ parts, by the pigeonhole principle there exists a part P of L_1 and a part Q of L_2 or L_3, say L_2, such that at least $R(s(r+1),2)$ elements of $\mathcal{Q}$ are P-Q-paths. Let $\mathcal{P}$ be the subset of $\mathcal{Q}$ containing these. By Lemma 4.2.3 there is a set of $r+1$ disjoint s-ladders contained in $P \cup Q \cup \mathcal{P}$. But these s-ladders all avoid L_3, contradicting the fact that at most r disjoint s-ladders miss any of L_1, L_2 and L_3. $\square$

4.3 Networks

Our goal now is to extend the proof idea of Theorem 4.2.4 to show that IG for G finite and planar are ubiquitous. In that proof it was important that our ladders missed exactly one of the ladders fixed before. First we will prove a lemma which will help us reach a similar situation more generally. For this we work with a simple connected path space, which we call *network*. It will be useful to first give an alternate characterization of networks. Thus we call a finite enumeration of paths $P_1 \dots P_n$ of a path space $\mathcal{P}$ *passable* if it covers the ground set and P_i meets $\bigcup_{j<i} P_j$ for every $i > 1$.

Lemma 4.3.1. *A path space is a network if and only if it has a passable enumeration.*

Proof. First assume that $\mathcal{P}$ has a passable enumeration $P_1, \dots, P_n$. Then clearly it is simple. To prove that it is connected, it suffices to prove that for every $k \leq n$ the set $\bigcup_{i \leq k} P_i$ is contained in a single equivalence class. This we do by induction on k. The case $k = 1$ is trivial. For $k > 1$ we already know that $\bigcup_{i<k} P_i$ is contained in a single equivalence class. But P_k meets this set and is contained in a single equivalence class, so we are done.

Now let $\mathcal{P}$ be a network. Since $\mathcal{P}$ is simple there exist some $Q_1, \dots, Q_k$ covering its ground set. Since $\mathcal{P}$ is connected, we may choose for every $(i, j) \in \{1, \dots k\}^2$ a Q_i-Q_j-path $R_{i,j}$ in $\mathcal{P}$. Let X be the set of all the Q_i and $R_{i,j}$ and define a graph G on X by adding an edge between $x_1, x_2 \in X$ if they meet. Since G contains a path through all Q_i and each $R_{i,j}$ is adjacent to one of them, G is connected. Thus there exists an enumeration $P_1 \dots P_n$ of its vertices, where P_i has a neighbor in $\{P_j; j < i\}$ for every $i > 1$. By definition of G and choice of $Q_1, \dots, Q_k$ this enumeration is passable. □

Using this enumeration, we can now simply cut off our network at the correct point.

Lemma 4.3.2. *Let $\mathcal{P}$ be a path space and $\mathcal{Q}$ a network in $\mathcal{P}$. Then for any finite set $\mathcal{A}$ of simple sets of $\mathcal{P}$ meeting $\mathcal{Q}$ there exists some network $\mathcal{Q}'$ in $\mathcal{Q}$ avoiding exactly one element of $\mathcal{A}$.*

Proof. By Lemma 4.3.1 $\mathcal{Q}$ has some passable enumeration $P_1, \dots, P_n$. Let k be minimal such that $\bigcup_{i \in k+1} P_i$ avoids at most one element of $\mathcal{A}$. If it does avoid one, we are done, so we may assume it does not. Let $\mathcal{A}'$ be the set of those $A \in \mathcal{A}$ such that $\bigcup_{i \in k} P_i$ does not meet A. For each $A \in \mathcal{A}'$ there is a $\bigcup_{i \in k} P_i$-A-path

Q_A that is a segment of P_k. Let A^* be such that Q_{A^*} is maximal among these with respect to inclusion. Let $\mathcal{Q}'$ be the subspace of $\mathcal{Q}$ which is the completion of the Q_A for $A \in \mathcal{A}' - A^*$ and the P_i for $i \in k$. Then $\mathcal{Q}'$ clearly avoids exactly one element of $\mathcal{A}$. But since adding the Q_A in any order after $P_1, \ldots, P_{k-1}$ gives a passable enumeration for $\mathcal{Q}'$, it is a network by Lemma 4.3.1. $\square$

Now the main difficulty remaining in extending Theorem 4.2.4 is finding a substitute for Lemma 4.2.3. During our proof we will end up with what we call a k, l*-pregrid*, namely a path space consisting of disjoint paths $(P_i)_{i \in k}$ and disjoint networks $(T_j)_{j \in l}$ such that for all $i \in k$ and $j \in l$ there is a vertex in $P_i \cap T_j$.

We will spend the rest of this section showing that any large enough quadratic pregrid contains an $IR_{n,n}$. Of course, this follows easily from Theorem 2.7.1, but it can also be deduced from the graph grid theorem with relative ease. First let us state the grid theorem which was first proven in [46]. Recall that a *bramble* of a graph is a set of connected subgraphs such that any two of its elements either meet in a vertex or are connected by an edge and its *order* is the smallest size of a vertex set meeting every element of the bramble.

Theorem 4.3.3. *There is a function f such that for any natural number n every graph with a bramble of order $f(n)$ has an $R_{n,n}$ minor.*

The main tool is a way to partition a path which captures its interaction with a set of disjoint paths.

An r-partition of a path P is a tuple of segments $(P_1 \ldots P_r)$ of P such that for every $1 \leq i < r$ the segment P_i either connects to P_{i+1} or connects to a segment of length two which connects to P_{i+1}, but such that all the segments are disjoint except where they connect and they cover P.

Proposition 4.3.4. *Let $\mathcal{P}$ be a path space, $P \in \mathcal{P}$ and $\mathcal{Q}$ a finite set of disjoint simple sets in $\mathcal{P}$. Then P has an r-partition $(P_1, \ldots, P_r)$ for some natural number r such that for each $1 \leq i \leq r$ at most one element of $\mathcal{Q}$ meets P_i.*

Proof. We may assume $V(P) \cap V(\mathcal{Q}) \neq \emptyset$, otherwise we are done. We start by constructing a sequence v in $V(P) \cap V(\mathcal{Q})$ recursively as follows. Let v_0 be the infimum in P of $V(\mathcal{Q})$. It remains to define v_{n+1} for some natural number n. Let $Q \in \mathcal{Q}$ be such that $v_\alpha \in Q$. Then we define v_{n+1} to be the infimum in the segment of P starting in v_n of $V(\mathcal{Q} - Q)$. If for some n we have defined v_n to be the supremum of $V(\mathcal{Q})$ in P, we stop the construction. We claim that the construction stops at some natural number k. Assume not. Since v_n and v_{n+1}

for any natural number n are never contained in the same element of $\mathcal{Q}$, there are different Q and R in $\mathcal{Q}$ which each contain v_k for infinitely many k. But then the supremum of $\{v_n; n \in \omega\}$ is contained in $Q \cap R$, a contradiction.

Now consider the partition into which $v_0, \ldots, v_k$ separate P. Each segment P', say with endpoints v_n and v_{n+1}, of this partition meets $Q \in \mathcal{Q}$ only if $v_n \in Q$ or $v_{n+1} \in Q$. Indeed, if it met some other $Q \in \mathcal{Q}$ then the infimum of $V(Q)$ in P' would have been a candidate for v_{n+1}, a contradiction. If P' does not meet two elements of $\mathcal{Q}$, we are done, so assume that it does meet some Q containing v_n and Q' containing v_{n+1}. Then by construction P' meets Q' only in v_{n+1}. Let x be the supremum in P' of points of $V(Q)$. If there is a point on P' between x and v_{n+1}, subdividing P' along it ensures that each segment only meets at most one $Q \in \mathcal{Q}$. Otherwise, subdividing P' along x gives a segment of cardinality two and a segment meeting at most one $Q \in \mathcal{Q}$. Finally, we go through all the segments R with $|R| = 2$ which meet two different elements of $\mathcal{Q}$, dropping R from the partition and replacing it by singleton segments for any points of R not contained in any other segment (if they exist). □

Proposition 4.3.5. *Any $f(n), f(n)$-pregrid contains an $IR_{n,n}$.*

Proof. Let $\mathcal{P}$ be an $f(n), f(n)$-pregrid with set of paths $\mathcal{Q}$ and set of networks $\mathcal{T}$. Fix an r-partition as in Proposition 4.3.4 for each $Q \in \mathcal{Q}$ with respect to $\mathcal{T}$ and let X be the set of all their segments together with all the components of $T - \bigcup \mathcal{Q}$ for each $T \in \mathcal{T}$. Let $\psi : V(\mathcal{P}) \to X$ be defined by mapping each point contained in some $Q \in \mathcal{Q}$ to its segment in the relevant r-partition, choosing the earlier if there are two possibilities, and mapping each other point to the component containing it. Let G be a the graph on X with an edge whenever there is a path between two elements of X not meeting any other element of X. Then $\mathcal{P}$ contains an IG with branch sets $\psi(x)$ for $x \in X$. Let G' be obtained from G by identifying for every $T \in \mathcal{T}$ all those vertices arising from T which have a common neighborhood and let φ be the corresponding map from G to G'. Then G' is a finite graph and any $IR_{n,n}$ in G' must clearly also be contained in G.

As a consequence it is enough to show that G' contains a bramble of order at least $f'(n)$, since by Theorem 4.3.3 it and then also G will contain an $IR_{n,n}$ and by transitivity so will $\mathcal{P}$. For every $Q \in \mathcal{Q}$ and $T \in \mathcal{T}$ let $X_{Q,T} = \varphi(\psi(Q) \cup \psi(T))$. We claim that this forms the desired bramble.

The set $X_{Q,T}$ is connected because T is connected. Furthermore, for $Q_1, Q_2 \in \mathcal{Q}$ and $T_1, T_2 \in \mathcal{T}$ the set X_{Q_1,T_1} meets X_{Q_2,T_2} because T_1 meets Q_2. It remains to be proven that this bramble has order $f(n)$. Note that all the $\psi(Q)$ for $Q \in \mathcal{Q}$ are disjoint because the $Q \in \mathcal{Q}$ are. Similarly all the $\psi(T)$ for $T \in \mathcal{T}$ are disjoint because every segment of the partitions meets only one $T \in \mathcal{Q}$. In both cases this remains true after applying φ because the only vertices arising from a common $Q \in \mathcal{Q}$ were identified. Thus any set $Y \subseteq V(G')$ of order less than $f'(n)$ avoids at least one $\varphi(\psi(Q))$ for $Q \in \mathcal{Q}$ and at least one $\varphi(\psi(T))$ for $T \in \mathcal{T}$, and thus does not meet $X_{Q,T}$, completing the proof. □

4.4 Grid minors

At this point one might think that slotting Proposition 4.3.5 into the argument from Theorem 4.2.4 will be enough to show the ubiquity of IG for finite planar G. However, we run into a difficulty during the initial choices. Recall that that early in the proof of Theorem 4.2.4 we chose three disjoint ladders and these ladders gaves us some maximum number r of disjoint ladders avoiding them, which we then contradicted by constructing more. We can of course start with some larger number k of IG instead of three ladders, but only afterwards will we know the number r of disjoint IG avoiding them that we will have to construct. We will then be able to construct an a, b-pregrid, where a depends only on k and b also depends on r. Applying Proposition 4.3.5 to this will only give us some $IR_{n,n}$ where n depends on a and thus not on r, which will generally not contain r disjoint IG.

However, it would be enough to find an $IR_{n,m}$ where m is much bigger then r and thus we want to prove the following statement.

Proposition 4.4.1. *There are functions $g : \mathbb{N} \to \mathbb{N}$ and $h : \mathbb{N}^2 \to \mathbb{N}$ such that for $n, m \in \mathbb{N}$ any $g(n), h(n,m)$-pregrid contains an $IR_{n,m}$.*

Note that even for graphs this statement is new and not a direct consequence of the theory of tree width. The idea for the proof is to try to split the $g(n)$ paths into disjoint parts repeatedly to build an l, l-pregrid for $l \gg m$ and then apply Proposition 4.3.5 and if we get stuck and cannot find any more paths to split we can use the extra structure this gives to construct the $IR_{n,m}$ by hand.

Let $\mathcal{P}$ be a path space containing some path P and some set of networks $\mathcal{T}$. Then we say that $(P, \mathcal{T})$ *r-splits into* $((Q, R), \mathcal{T}')$ if Q and R are disjoint segments of P, $\mathcal{T}'$ is a subset of $\mathcal{T}$ of size at least r and Q and R both meet every element of $\mathcal{T}'$.

We say that $(P, \mathcal{T})$ *r-segments into* $((P_1, \ldots, P_r), \mathcal{T}')$ if $(P_1, \ldots, P_r)$ is an r-partition of P, $\mathcal{T}' \subseteq \mathcal{T}$ and there is a bijection f from $\{1, \ldots, r\}$ to $\mathcal{T}'$ such that for any $i \in \{1, \ldots, r\}$ the segment P_i meets $f(i)$, but no other element of $\mathcal{T}'$. When we just write that $(P, \mathcal{T})$ r-splits this means that there is some pair it r-splits into and similarly for r-segmenting.

Lemma 4.4.2. *Let $r \in \mathbb{N}$ and let $\mathcal{P}$ be a path space containing some path P and some set of networks $\mathcal{T}$ which all meet P. If $|\mathcal{T}| \geq r^2$, then $(P, \mathcal{T})$ r-splits or r-segments.*

Proof. We first enumerate the $T \in \mathcal{T}$ as $T_1, \ldots, T_{r^2}$ ordered by their first point in the order of P, which we call x_i for the network T_i. Let $(P_1, \ldots, P_r)$ be the r-partition with P_1 ending at the predecessor of x_r if it exists or x_r itself otherwise, P_r starting at $x_{r(r-1)}$ and every other P_i starting at $x_{(i-1)r}$ and ending at the predecessor of x_{ir} if it exists or x_{ir} itself otherwise.

If for every $1 \leq i \leq r-1$ there is some $ir \leq j < (i+1)r$ such that T_j only meets P_i, then $(P, \mathcal{T})$ r-segments.

Otherwise there is some $1 \leq i \leq r-1$ such that for every $ir \leq j < (i+1)r$ the network T_j has some point y_j after $x_{(i+1)r}$. Let y be a point of P strictly larger than $x_{(i+1)r}$ and at most y_j for all the relevant j. Let Q be the segment of P up to $x_{(i+1)r}$, let R be the segment of P starting at y and let $\mathcal{T}' = \{T_j; ir \leq j < (i+1)r\}$. Then $|\mathcal{T}'| = (i+1)r - ir = r$. Thus $(P, \mathcal{T})$ r-splits. $\square$

Let $g(n) = n^n$, $s(n, m) = \max(R(mng(n)^{g(n)}, g(n)), f(m))$ and $h(n, m) = s(n, m)^{2(f(m)+g(n))}$

Proposition 4.4.3. *For $n, m \in \mathbb{N}$ any $g(n), h(n, m)$-pregrid contains an $IR_{n,m}$.*

Proof. Without loss of generality $m \geq n$. Let $\mathcal{P}$ be a $g(n), h(n, m)$-pregrid. Let $r(i) = s(n, m)^{2(f(m)+g(n)-i)}$ for all $i \in \mathbb{N}$.

For $0 \leq i \leq f(m) + g(n)$ we will inductively construct pregrids $\mathcal{P}_i$ with set of networks $\mathcal{T}_i$ of size at least $r(i)$ starting with $\mathcal{P}$. During this construction we will always either segment or split a path and we keep track of which paths have been segmented.

To construct $\mathcal{P}_{i+1}$ from $\mathcal{P}_i$ let P be a path of $\mathcal{P}_i$ which has not been segmented. By Lemma 4.4.2 $(P, \mathcal{T}_i)$ either $r(i+1)$-splits or $r(i+1)$-segments. If it $r(i+1)$-segments, we simply obtain $\mathcal{P}_{i+1}$ by dropping all networks which do not occur in the pair it segments into. If it $r(i+1)$-splits into $((Q, R), \mathcal{T}')$, we obtain $\mathcal{P}_{i+1}$ by again dropping all networks outside $\mathcal{T}'$ and also dropping P and replacing it by Q and R.

If we split at least $f(m)$ times during this construction, then $\mathcal{P}_{f(m)+g(n)}$ contains an $f(m), f(m)$-pregrid and so we are done by Proposition 4.3.5. So we may assume that we have segmented at least $g(n)$ times. Now let $\mathcal{T}' = \mathcal{T}_{f(m)+g(n)}$ and let $P_1, \ldots, P_{g(n)}$ be paths of $\mathcal{P}_{f(n)+g(n)}$ which $s(n,m)$-segment into $((Q_{i,1}, \ldots, Q_{i,s(n,m)}), \mathcal{T}')$.

For each $1 \leq i \leq n^n$ let $\leq_i$ be the total order on $\mathcal{T}'$ defined by the unique $1 \leq j \leq s(n,m)$ for which they meet $Q_{i,j}$. Since $|\mathcal{T}'| \geq R(mn(n^n)^{n^n}, n^n)$ there is a subset $\mathcal{S}$ of $\mathcal{T}'$ of size $mn(n^n)^{n^n}$ sorted in $\{\leq_1, \ldots, \leq_{n^n}\}$. For each $S \in \mathcal{S}$ let G_S be the graph on $\{1, \ldots, n^n\}$ with an edge (i,j) if S contains a P_i-P_j-path not meeting any other P_l for $1 \leq l \leq n^n$. Since S is a network, G_S is connected. Because $|\mathcal{S}| \geq mn(n^n)^{n^n}$, there is some $\mathcal{S}' \subseteq \mathcal{S}$ of size mn and a tree T on $\{1, \ldots, n^n\}$ such that T is a subgraph of G_S for any $S \in \mathcal{S}'$.

Now T contains either a path of length $n-1$ or a vertex of degree n. First assume that T has a path with vertices $i_1, \ldots, i_n$. Select some $\mathcal{S}'' \subseteq \mathcal{S}'$ of size m arbitrarily. Then combining $P_{i_1} \ldots P_{i_n}$ with a P_{i_j}-$P_{i_{j+1}}$-path $S \in \mathcal{S}''$ for every $1 \leq j \leq n-1$ gives an $IR_{n,m}$.

Now assume that T has a vertex i with neighbors $i_1 \ldots, i_n$. List $\mathcal{S}'$ as $S_1 \ldots S_{mn}$ according to $\leq_1$. Then combining $P_{i_1} \ldots P_{i_n}$ with a P_i-P_{i_j}-path and a P_i-$P_{i_{j+1}}$-path from S_{on+m} together with a segment of P_i connecting them for every $1 \leq j \leq n$ and $0 \leq o < m$ gives an $IR_{n,m}$. $\square$

A class $\mathcal{C}$ of networks is called *constructible* if there exist $m, n \in \mathbb{N}$ such that any $IR_{m,n}$ contains an element of $\mathcal{C}$. In particular the class of IG for G finite and planar is constructible.

Theorem 4.4.4. *Any constructible class of networks is ubiquitous.*

Proof. Let $\mathcal{C}$ be a constructible class of networks. Assume that $\mathcal{P}$ contains arbitrarily many disjoint elements of $\mathcal{C}$. Recursively delete any redundant element of $\mathcal{C}$ as long as there is one. It is enough to show that this never terminates, since then we will have found an infinite set of disjoint elements of $\mathcal{C}$ in $\mathcal{P}$. So we assume for a contradiction that the process terminates after only finitely many steps, without loss of generality we may then assume that $\mathcal{P}$ contains no redundant elements of $\mathcal{C}$.

Since $\mathcal{C}$ is constructible, there are $m, n \in \mathbb{N}$ such that every $IR_{m,n}$ contains some element of $\mathcal{C}$. Let $\mathcal{A}$ be a set of $g(m)+1$ disjoint elements of $\mathcal{C}$ contained in $\mathcal{P}$. Since none of these is redundant there exists some $r \in \mathbb{N}$ such that at most r disjoint elements of $\mathcal{C}$ in $\mathcal{P}$ may avoid any of the these. Since every element of

$\mathcal{C}$ is simple, there is some $b \in \mathbb{N}$ such that every element of $\mathcal{A}$ is covered by b of its paths.

Let D be a set of $h(m, (r+1)n)(g(m)+1)b^{g(m)} + r$ disjoint elements of $\mathcal{C}$. Deleting any elements which do not meet an element of $\mathcal{A}$ leaves a set $D' \subseteq D$ of size $h(m, (r+1)n)(g(m)+1)b^{g(m)}$. By Lemma 4.3.2 for each $d \in D'$ we can find some network in d such that it avoids exactly one element of $\mathcal{A}$, so there is some $D'' \subseteq D'$ of size $h(m, (r+1)n)b^{g(m)}$ for which this is the same $A^* \in \mathcal{A}$. Let E contain for each element of D'' a network witnessing this. Any $A \in \mathcal{A} - A^*$ is covered by at most b of its paths, so we can recursively find a set $\mathcal{A}'$ of one path from each $A \in \mathcal{A} - A^*$ and a set $E' \subseteq E$ of size $h(m, (r+1)n)$ such that every element of $\mathcal{A}'$ meets every element of E'. Now $\mathcal{A}'$ and E' form a $g(m), h(m, (r+1)n)$-pregrid, which contains an $IR_{m,(r+1)n}$ $\mathcal{H}$ by Proposition 4.4.3. Clearly $\mathcal{H}$ contains disjoint $IR_{m,n}$ $\mathcal{H}_i$ for $i \in r+1$. For all $i \in r+1$ $\mathcal{H}_i$ contains some $C_i \in \mathcal{C}$ by choice of m and n. Since these are all disjoint and do not meet A^*, this contradicts the choice of r. $\square$

The fact that planarity is pivotal both for our ubiquity proof and for our discussion on the Erdős-Pósa property in Section 2.8 suggests a possible connection. This connection is provided by an infinite analogue to the Erdős-Pósa property. A path space $\mathcal{Q}$ has the *ω-EP* if every path space $\mathcal{P}$ either contains infinitely many disjoint $I\mathcal{Q}$ or finitely many points meeting every $I\mathcal{Q}$.

Proposition 4.4.5. *Let $\mathcal{Q}$ be a path space. If $\mathcal{Q}$ has the ω-EP, it is ubiquitous. If $\mathcal{Q}$ has the EP and is ubiquitous, it has the ω-EP.*

Proof. First assume that $\mathcal{Q}$ has the ω-EP. If $\mathcal{P}$ does not contain infinitely many disjoint $I\mathcal{Q}$, then it contains a finite set of points X meeting every $I\mathcal{Q}$. In particular, it cannot contain $|X|+1$ disjoint $I\mathcal{Q}$.

Now assume that $\mathcal{Q}$ has the EP and is ubiquitous. If $\mathcal{P}$ does not contain infinitely many disjoint $I\mathcal{Q}$, then it does not contain k disjoint $I\mathcal{Q}$ for some finite k. Thus it contains some set X of size $l \gg k$ meeting every $I\mathcal{Q}$. $\square$

Corollary 4.4.6. *Let G be a finite, planar graph. Then IG has the ω-EP.*

Proof. By Proposition 4.4.5 this follows from Theorems 2.8.6 and 4.4.4. $\square$

4.5 Edge ubiquity

In graphs ubiquity is considered not only with respect to total disjointness, but also edge-disjointness. In path spaces edges are not necessary for connectivity, but we can still define a notion analogous to edge-disjointness. Given a path space $\mathcal{P}$ we call two subspaces of $\mathcal{P}$ *edge-disjoint* if none of their paths share a nontrivial segment. Note that this definition agrees which the usual one for graph-like spaces.

Helpfully, in showing that cycles are edge-ubiquitous we may assume that any two paths meet in only finitely many segments, since otherwise there are already infinitely many disjoint cycles between them, as we will prove next.

Let P and Q be paths. We call a pair of points x and y in their intersection *segment-equivalent* if the segment of P between x and y is the same as that segment on Q. This is clearly an equivalence relation. In particular, if P and Q are edge-disjoint, no different points in $P \cap Q$ are segment-equivalent. Furthermore P and Q induce a linear order on the segment-equivalence classes, since these are closed intervals.

Lemma 4.5.1. *Let $\mathcal{P}$ be a path space containing paths P and Q such that $P \cap Q$ has infinitely many segment-equivalence classes. Then $P \cup Q$ contains infinitely many disjoint circuits.*

Proof. Let X be the set of segment-equivalence classes of $P \cap Q$ and let l_P and l_Q be the orders P and Q respectively induce on X. Since $R(\aleph_0, 2) = \aleph_0$, there exists an infinite subset Y of X sorted in $\{l_P, l_Q\}$. Find in Y a set of infinitely many pairs whose intervals do not intersect. For each of these pairs $p = \{x, y\}$ let C_p be the network induced by the segments between x and y in P and Q. Since x and y are not segment-equivalent, C_p contains two different paths between x and y, so it includes a circuit. By construction, the C_p are all disjoint. $\square$

This will allow us to keep the intersection of sets of edge-disjoint cycles finite. The following lemma will then be useful to extract paths with no inner points in such cycles.

Lemma 4.5.2. *Let $\mathcal{P}$ be a network and C and D simple subsets of its ground set with finite intersection and $C \setminus D$ nonempty. Then $\mathcal{P}$ contains a $C \setminus D$-D-path or a nontrivial path from C to $C \cap D$ not meeting D in an inner point.*

Proof. Let $c \in C \setminus D$ and $d \in D$ be arbitrary. Since $\mathcal{P}$ is connected, it contains a path P from c to d. If P does not meet $C \cap D$, we are done. Otherwise let x be the first point of $C \cap D$ in P after c. Let P' be the segment of P from c to x. If P' does not meet D in an inner point, we are again done. If it does, in y say, then the part of P' between c and y is a network which meets C and D in nonempty disjoint sets, so it contains a $C \setminus D$-D-path. $\square$

The finite intersection between the cycles is also useful for our equivalent of Lemma 4.2.3, together with the classic matching theorem of Kőnig first proven in [38].

Lemma 4.5.3. *Let $\mathcal{P}$ be a path space and let P and Q be edge-disjoint paths in it meeting in at most n points. Let $\mathcal{Q}$ be a set of $2kR(2k,2)+4kn$ edge-disjoint $P \setminus Q$-Q-paths. Then there exists a set of k edge-disjoint circuits contained in $\mathcal{P}$.*

Proof. If there are $2k$ elements of $\mathcal{Q}$ from some point v to P or Q, say P, combining them in pairs along P together with the segment of P between their endpoints gives k edge-disjoint networks, each of which contains a circuit. So we may assume that this never occurs and so deleting any elements of $\mathcal{Q}$ with an endpoint in $P \cap Q$ still leaves a set $\mathcal{Q}'$ with at least $2kR(2k,2)$ elements. Now the graph H on the vertices of P and Q where each element of $\mathcal{Q}'$ defines an edge between its endpoints is bipartite. Since H has maximum degree less than $2k$, every vertex cover of H has at least $R(2k,2)$ elements and by König's theorem there exists a matching of that size. Let $\mathcal{Q}''$ be the corresponding subset of $\mathcal{Q}'$.

Let l_P and l_Q be the linear orders on $\mathcal{Q}''$ given by the order of their endpoints on P and Q respectively. Since $|\mathcal{Q}''| \geq R(2k,2)$ there is a set X of size $2k$ sorted in $\{l_P, l_Q\}$. Let P consist of k adjacent pairs from X. For each $p \in P$ let C_p consist of the two paths in p together with those segments of P and Q between their endpoints. These circuits are then edge-disjoint by construction. $\square$

Call a circuit C *edge-redundant* if there are arbitrarily many edge-disjoint circuits in the space obtained from $\mathcal{P}$ by deleting the paths of C. Now we can once again follow the basic structure of Theorem 4.2.4.

Theorem 4.5.4. *Cycles are edge-ubiquitous in path-spaces.*

Proof. Assume that $\mathcal{P}$ contains arbitrarily many edge-disjoint circuits. Recursively delete any edge-redundant circuit in $\mathcal{P}$ as long as there is one.

It is enough to show that this never terminates, since then we will have found an infinite set of edge-disjoint circuits. So we assume for a contradiction that the process terminates after only finitely many steps, without loss of generality we may then assume that $\mathcal{P}$ has no edge-redundant circuits. Let C_1, C_2 and C_3 be three edge-disjoint circuits in $\mathcal{P}$. By Lemma 4.5.1 and edge-disjointness of the C_i we may assume that the set X of all points meeting more than one of C_1, C_2 and C_3 is finite. Since none of these is edge-redundant there exists some r such that at most r edge-disjoint circuits are edge-disjoint from any of them.

Let $\mathcal{C}$ be a set of $4(r+1)R(2r+2,2)(2)+10(r+1)|X|+r+12|X|$ edge-disjoint circuits. Deleting any circuits which are edge-disjoint from C_1, C_2 or C_3 leaves a set $\mathcal{C}'$ of $4(r+1)R(2r+2,2)+10(r+1)|X|+12|X|$ edge-disjoint circuits not edge-disjoint from all three. Since there are at most two directions at any vertex in a path, deleting all those elements of D' which have a direction at a vertex of X which is the same as one in C_1, C_2 or C_3 still leaves a set $\mathcal{C}''$ of $4(r+1)R(2r+2,2)+10(r+1)|X|$ edge-disjoint circuits. Let $D = C_2 \cup C_3$. By Lemma 4.5.2 each element of $\mathcal{C}''$ contains either a $C_1 \setminus D$-D-path or a nontrivial path from C_1 to $C_1 \cap D$ whose inner points avoid D. Let Z consist of one of these for every element of $\mathcal{C}''$.

Let us first assume that Z contains a set $\mathcal{Q}$ of edge-disjoint $C_1 \setminus D$-D-paths of size $4(r+1)R(2r+2,2)+8(r+1)|X|$. By the pigeonhole principle we may assume that $2(r+1)R(2r+2,2)+4(r+1)|X|$ of them are $C_1 \setminus D$-C_2-paths, otherwise exchange the roles of C_2 and C_3. Let $\mathcal{P}$ be the subset of $\mathcal{Q}$ containing all such paths. By Lemma 4.5.3 there is a set of $r+1$ edge-disjoint circuits contained in $C_1 \cup C_2 \cup \mathcal{P}$. But these circuits all avoid C_3, contradicting the fact that at most r edge-disjoint circuits miss any of C_1, C_2 and C_3.

So we may assume that Z contains a set $\mathcal{R}$ of $2(r+1)|X|$ edge-disjoint nontrivial paths from C_1 to $C_1 \cap D$ whose inner points avoid D. If there is some $R \in \mathcal{R}$ such that the common points of R and one of the paths of C_1 have infinitely many segment-equivalence classes, then we have found infinitely many disjoint cycles by Lemma 4.5.1. Otherwise each $R \in \mathcal{R}$ can be shortened to a C_1-$C_1 \cap D$-path. Let $\mathcal{R}'$ be the set of all these shortened elements of $\mathcal{R}$. By the pigeonhole principle $\mathcal{R}'$ has a subset $\mathcal{R}''$ of size $2(r+1)$ all of whose elements have a common endpoint v in $C_1 \cap D$. Then combining them in pairs along the order of C together with the segment of C between their endpoints gives $r+1$ edge-disjoint networks, each of which contains a circuit. Since these are all edge-disjoint from D, this again contradicts the fact that at most r edge-disjoint circuits miss any of C_1, C_2 and C_3. □

4.6 Counterexamples

To construct our counterexamples, we make use of a counterexample to Menger's theorem for cardinality $\aleph_0$. We will however need certain additional properties which the example given in Section 1.5 does not have, such as having maximal degree three. While it is possible to modify that example to remedy this, we will instead give a different construction which will naturally have the required properties.

For the construction we will work with graph-like spaces to show that these counterexamples are not an artifact of our definition of path spaces. Of course any such example for graph-like spaces defines a path space counterexample. Furthermore, since we only work with the usual arcs instead of pseudo-lines, the counterexample also covers the other main example, being the set of arcs of a Hausdorff space.

A *turbulence* is an equivalence relation on $\mathbb{N}^2$ satisfying the following conditions:

1. (a, b) is equivalent to (a, c) only if $b = c$.
2. Each equivalence class has size at most 2.
3. For any $a \neq b$ there are at most finitely many x and y for which $(a, x) \sim (b, y)$.

For some turbulence $\sim$ let $\Gamma(\sim)$ be the graph obtained from the union of ω rays (which we number as R_i and their edges as $e_{i,j}$) by identifying $e_{a,3b}$ with $e_{c,3d}$ along the order of the rays whenever $(a, b) \sim (c, d)$. We write a_i for the initial vertex of R_i Let $G(\sim)$ be the space obtained from the geometric realization of $\Gamma(\sim)$ by adding an additional vertex v_i for each ray R_i and additionally declaring as open all those sets whose intersection with $G(\sim)$ is open and which contain precisely one v_i together with a tail of the corresponding R_i.

Lemma 4.6.1. *Let $\sim$ be a turbulence. Then $G(\sim)$ is a graph-like space.*

Proof. Since the geometric realization is a graph-like space, all that needs to be shown is that for any vertices v and w of $G(\sim)$ there are disjoint open sets containing v and w respectively and covering $V(G(\sim))$. When constructing these open sets, we always take less than half of any edge with only one endpoint in the set to ensure disjointness.

If both v and w were already vertices of $\Gamma(\sim)$, constructing such open sets is easy, so we may assume without loss of generality that $v \notin \Gamma(\sim)$. Let R be the R_i ending in v. Note that R has a tail not containing w since every equivalence class of $\sim$ is finite. Let V be an open set arising from that tail by adding half-edges for any edge incident with the tail.

By the third condition for a turbulence every R_i has a tail avoiding V. Taking the union of open sets arising from all these tails together with open sets around every vertex not in V gives an open set W disjoint from V such that V and W cover the vertex set.

$\square$

A turbulence $\sim$ is called *disastrous* if sets of arbitrarily many R_i are still disjoint in $\Gamma(\sim)$, but $G(\sim)$ does not contain infinitely many edge-disjoint arcs each from some a_m to some v_n.

Proposition 4.6.2. *There is a disastrous turbulence.*

Proof. Let $g : \mathbb{N} \to \mathbb{N}$ be an increasing function such that $|g^{-1}(n)| = n$ for $n \in \mathbb{N}$. Clearly such a function exists. Let $\sim$ be the equivalence relation induced by $(a, b) \sim (b, a)$ if $g(a) \neq g(b)$. Since there are arbitrarily long intervals on which g is constant, we can find arbitrarily many of the initial rays in $G(\sim)$ which are disjoint. On the other hand, any arc f ending in some v_n contains a tail of R_n, since the image of f is homeomorphic to $[0, 1]$ and there is an open set containing v_n and otherwise just consisting of such a tail and half-edges to it. Then there is some $m > n$ such that f meets all R_k for $k \geq m$ in $e_{k,3n}$. Clearly, there are at most finitely many edge-disjoint arcs edge-disjoint from f starting with $e_{l,1}$ or ending at v_l for $l < m$. Let f' be some other arc from some a_i to some v_j. For some $X \subseteq \mathbb{N}$ and $z \in \mathbb{N}$ variable let $B(X, z)$ be the set of all vertices on any R_x with $x \in X$ before $e_{x,3z}$ and $A(X, z)$ the set of those vertices after $e_{x,3z}$. Then f' contains a $B(\mathbb{N}_{\geq m}, n)$-$A(\mathbb{N}_{\geq m}, n)$-path P. Since only pairs which are identical except for the order are $\sim$-equivalent, edges in $B(\mathbb{N}_{\geq m}, n)$ are identified only with edges in $A([0, n-1], n)$ (because $n < m$). But the first edge of P must leave $A(\mathbb{N}_{\geq m}, n)$ and thus P must already start in a vertex of $A([0, n-1], n)$. Similarly, P must end in $A([n+1, m], n)$. So P is a path starting in $A([0, n-1], n)$ and ending in $A([n+1, m], n)$, but not meeting any R_i for $i \geq m$ outside its first and last vertices and thus P must include $e_{l,n}$ for some $l < m$. Since there are only finitely many such edges, we are done. $\square$

For the rest of this section we fix some disastrous turbulence $\sim$ and define $X = G(\sim)$. Now let H be some finite graph and F some subset of its edge set.

We can construct a graph-like space $X(H, F)$ starting with one copy X_e of X for each edge e of F by adding ω copies of the graph with vertex set $V(H)$ and edge set $E(H) \setminus F$ and then for each edge e of F with endpoints a and b[1] joining the version of a in copy n with v_n in X_e and the version of b in copy n with the

[1] These are chosen consistently, that is we fix an orientation of e and call its first vertex a and its last vertex b.

starting vertex of R_n in X_e with an edge and then deleting the corresponding copy of e.

Example 4.6.3. Let $k \geq 5$ be odd. Then $G = X(K^k, E(K^k))$ obviously contains arbitrarily many disjoint TK^k. Let us first show that every TK^k in G' contains an arc in some copy of X from some v_n to some a_m.

Since the vertices of the added K^k are the only ones at which there are at least four directions, the branch vertices of any TK^k must be such vertices. If two branch vertices are copies of different vertices, then our TK^k contains the required arc. Thus we may assume that they are all copies of the same vertex v and all its edges are embedded as arcs starting and ending on the same side of some copy of X which corresponds to some e incident with the vertex v. In each copy of X each of the k branch vertices can only be adjacent with one arc of our copy, so at most $\frac{k-1}{2}$ edges can be embedded in each. But since there are only $k-1$ copies which we can use and K^k has $\frac{k(k-1)}{2}$ edges, this is a contradiction.

Thus the assertion is proven. But by choice of X there is no infinite edge-disjoint set of such arcs. In particular, there is no infinite edge-disjoint set of TK^k in G.

For a finite graph G and a natural number k, we write $|G|_{\geq k}$ for its number of vertices of degree at least k. We call finite graphs G_1 and G_2 *degree-incompatible* if there are $i, j > 3$ with $|G_1|_{\geq i} > 2|G_2|_{\geq i}$ and $|G_2|_{\geq j} > 2|G_1|_{\geq j}$. An bridge e of a connected graph G with no vertex of degree three is called *troublesome* if it separates G into two components C_1 and C_2 such that $C_1 + e$ and $C_2 + e$ are degree-incompatible.

We will now show that if G has a troublesome bridge, the class of TG is not ubiquitous. Note that there are trees with a troublesome bridge, like the one in Fig. 15.

Example 4.6.4. Let G be a graph with a troublesome bridge e. Then $H = X(G, \{e\})$ obviously contains arbitrarily many disjoint TG. As in Example 4.6.3 it is enough to show that that every TG in H contains an arc in the sole copy of X connecting its two sides, since then there cannot be even infinitely many edge-disjoint TG in H by choice of X.

Let C_1 and C_2 be the two components of $G - e$. Then $C_1 + e$ and $C_2 + e$ are degree-incompatible witnessed by some natural numbers i and j. Then the vertices of $C_1 + e$ of degree at least i must be embedded to points of degree i in H, so clearly into points of copies of G.

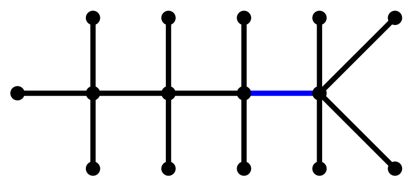

Figure 15: A tree with the blue edge as a troublesome bridge.

We claim that at least one of them is embedded into a copy of C_1. Assume the contrary. Then by the pigeonhole principle they need to be embedded into at least three different copies of C_2. But since each of these copies is separated from its complement by a single vertex, there must be a point of G of degree at least three embedded into the sole copy of X. But G has no vertex of degree exactly three and X has maximal degree three, so this is impossible.

Similarly, at least one vertex of degree at least j in $C_2 + e$ must be embedded into a copy of C_2. But since G is connected and e separates these two vertices, we are done.

4.7 Outlook

Our results still leave many questions unanswered, chief among them perhaps if either of the following two mutually exclusive conjectures is true.

Conjecture 4.7.1. *For all finite graphs G the class IG is ubiquitous.*

Conjecture 4.7.2. *For a finite graph G, IG is ubiquitous if and only if G is planar.*

Solving the case of $K_{3,3}$ could be a significant step towards answering this. The structure of our proof also suggests a question about finite graphs which could be helpful. For a graph G we ask if there is a natural number k such that for all natural numbers l, there is some natural number m such that for every graph H containing m disjoint IH and for every set $\mathcal{A}$ of k disjoint IG in H there is some $A \in \mathcal{A}$ such that $H - A$ contains l disjoint IH. Our proof of Theorem 4.4.4 builds on the grid theorem answering this question in the affirmative for planar G, but for other G finding an answer appears to be much harder. One could also ask whether this property is equivalent to ubiquity of IG.

Examples 4.6.3 and 4.6.4 seem to suggest that we should not expect many positive results for TG beyond those which follow directly from the results for minors. For edge-ubiquity, the situation is even less clear. Our methods do not seem to extend to anything beyond cycles at all, yet we also know of no counterexamples beyond the ones for (standard) ubiquity.

One could also consider more infinitary objects, like embedding of infinite graphs. For rays (either with or without end) the problem can be reduced to graphs, since any embedding of a ray which does not have a ray as a subspace already contains infinitely many disjoint embeddings of rays. We do not have any more interesting results, however.

One final interesting open question is suggested by Proposition 4.4.5, namely whether the ω-EP implies the EP. This would make Proposition 4.4.5 into a true equivalence.

Appendices

A.1 Cyclic orders and cycle completions

The goal of this appendix is to show that cycle completions exist, are unique and have the properties we outlined in Section 3.4.1. While there is a lot of prior work in this area, some of which is referenced here, none of it matches our requirements exactly, so we include these proofs for completeness.

This section starts with a short collection of basics about cyclic orders and their connection to linear orders. In this section we often consider distinct linear or cyclic orders on the same ground set. Because of this, and in contrast to the rest of the paper, in this section cyclic orders and linear orders are not implicit but introduced more formally as relations on the ground set. So a cyclic order of a set S is a set $Z \subseteq S \times S \times S$ that is cyclic, antisymmetric, linear and transitive.

The notation of $\leq$ is kept for a linear order, but the linear order in question is added as an index where necessary, for example in $s \leq_L t$ for a linear order L. Similarly, for intervals of cyclic orders, the cyclic order in question may be indicated by an index.

Definition A.1.1. [41, Definition 1.1] Given two linear orders A and B on disjoint ground sets, the linear order $A \oplus B$ is the linear order defined on the union of the ground sets of A and B by letting $x \leq y$ if $x \leq_A y$ or $x \leq_B y$ or $x \in A$ and $y \in B$.

Definition A.1.2. [41, Lemma 1.11, Definition 2.1, Theorem 2.3] Given a linear order L on a set S, the cyclic order Z *induced by* L consists of those triples (s, s', t) of elements of S such that in L one of the equations $s < s' < t$, $s' < t < s$ and $t < s < s'$ holds. Given a cyclic order Z on a set S, a *cut of* Z is a linear order L on S such that Z is the cyclic order induced by L.

Lemma A.1.3. *[40, Theorem 3.1] For every cyclic order Z on set S and every s in S there is a cut of Z whose smallest element is s.*

Note that, by Lemma A.1.3, every non-trivial interval of a cyclic order Z is also an interval of a cut L of Z. Also, such an interval inherits a linear order from every cut of which it is an interval, and that linear order does not depend on the chosen cut.

Lemma A.1.4. *Let Z be a cyclic order on a set S, L a cut of Z and s, s' and t elements of S such that $(s, s', t) \in Z$. If $s \leq t$ in L, then $s < s' < t$ in L.* $\square$

Definition A.1.5. Given a cyclic order Z on ground set S, the cyclic order $\{(t, s, r) \colon (r, s, t) \in Z\}$ on S is the *mirror* of Z.

Monotone maps have the property that preimages of intervals are again intervals. Maps with that property are close to being monotone:

Lemma A.1.6. *Let Z be a cyclic order on a set S and Z' a cyclic order on a set S'. Let $f : S \to S'$ be a map such that for all intervals I of Z' the set $f^{-1}(I)$ is an interval of Z. Then f is a monotone map or a mirror of a monotone map.*

Proof. If for all elements r, s and t of S the implication

$$(f(r), f(s), f(t)) \in Z' \Rightarrow (t, s, r) \in Z$$

holds, then f is a mirror of a monotone map. So assume that there are elements r, s and t of S such that $(f(r), f(s), f(t)) \in Z'$ and $(r, s, t) \in Z$. We will start with an observation that we will refer to later in this proof. For this we consider any $u \in S$ such that $(f(r), f(u), f(t)) \in Z'$. Then $f^{-1}([f(t), f(r)])$ is an interval of S which contains t and r but not s, so $[t, r]$ is a subset of $f^{-1}([f(t), f(r)])$. As also $u \notin f^{-1}([f(t), f(r)])$, $u \notin [t, r]$ and thus $(r, u, t) \in Z$.

Now let r', s' and t' be elements of S such that $(f(r'), f(s'), f(t')) \in Z'$. In order to show that $(r', s', t') \in Z$, first consider the case that the number n of elements in $\{f(r'), f(s'), f(t')\}$ which are not contained in $\{f(r), f(s), f(t)\}$ is zero. Assume, by renaming if necessary, that $f(r') = f(r)$, $f(s') = f(s)$ and $f(t') = f(t)$. Then by three applications of our observation, $(r, s', t) \in Z$ and thus $(s', t', r) \in Z$ and hence $(r', s', t') \in Z$.

Next consider the case that $n = 1$. Assume, again by renaming if necessary, that $(f(r), f(r'), f(s)) \in Z'$ (see also the left cyclic order of Fig. 16). By our observation, $(r', s, t) \in Z$ and $(r', t, r) \in Z$ and hence also $(r', s, r) \in Z$. Then there are three cases: Either $f(s') = f(s)$ and $f(t') = f(t)$ or $f(s') = f(t)$ and $f(t') = f(r)$ or $f(s') = f(s)$ and $f(t') = f(r)$. In all three cases, by the case $n = 0$ also $(r', s', t') \in Z$.

Next consider the case that $n = 2$, and that one of the intervals $(f(r), f(s))$, $(f(s), f(t))$, and $(f(t), f(r))$ contains both elements of $\{f(r'), f(s'), f(t')\}$ which are not contained in $\{f(r), f(s), f(t)\}$. Assume, by renaming if necessary, that both $f(r')$ and $f(s')$ are contained in $(f(r), f(s))$. In that case, the fact that $(f(r'), f(s'), f(t')) \in Z'$ implies that $f(r') \in (f(r), f(s'))$ (see also the middle cyclic order in Fig. 16). Also, by the case $n = 1$, (r', s, t), (s', s, t) and (t, r, r') are all contained in Z. As $f^{-1}([f(t), f(r')])$ contains t and r' but neither s or s', $(r', s, t) \in Z$ implies $[t, r'] \subseteq f^{-1}([f(t), f(r')])$. Thus s' is not contained in $[t, r']$ and hence $(r', s', t) \in Z$. Because (s', s, t) and (t, r, r') are contained in Z, also

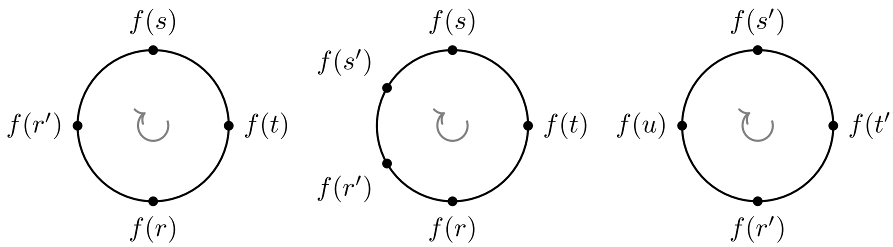

Figure 16: Three of the cases in the proof of Lemma A.1.6

(r', s', s) and (r', s', r) are contained in Z. By the case $n = 0$, also $(r', s', t') \in Z$.

Next consider the case that $n = 2$ and none of the intervals $(f(r), f(s))$, $(f(s), f(t))$ or $(f(t), f(s))$ contains two elements of $\{f(r'), f(s'), f(t')\}$. Assume, by renaming if necessary, that $f(t') \in \{f(r), f(s), f(t)\}$ and that there is $u \in \{r, s, t\}$ such that $(f(r'), f(u), f(s')) \in Z'$ (see also the right cyclic order of Fig. 16). In this case (t', r', u) and (u, s', t') are both contained in Z' by the case $n = 1$ and thus (r', s', t') is also contained in Z'.

The only case left is the case $n = 3$. Assume, by renaming if necessary, that $(f(r'), f(s), f(s')) \in Z'$. Then (s, s', t') and (s, t', r') are contained in Z by the case $n = 2$ and thus $(r', s', t') \in Z$. $\square$

Definition A.1.7 (Remark 2.3)**.** [40] Given a cyclic order Z of a set S and a subset S' of S, the set of triples in Z which only contain elements of S' is a cyclic order on S', the *induced cyclic order* on S'.

Theorem A.1.8. *[41, Theorem 3.6] Let Z be a cyclic order on set S and let K and L be distinct cuts of Z. Then there are non-empty disjoint subsets A and B of S such that $A \cup B = S$, $K \restriction A = L \restriction A$, $K \restriction B = L \restriction B$, $K = K \restriction A \oplus K \restriction B$ and $L = K \restriction B \oplus K \restriction A$.*

One example of a construction similar to the cycle completion is the following: The cycle completion of a cyclically ordered set I can be obtained from the Dedekind completion of one of its cuts L by adding as many elements to the ground set as necessary such that every element of the original ground set has a predecessor and a successor not in the ground set and then identifying the new smallest and biggest elements. Also constructing the pseudo-line as in [5] from L, contracting all inner points of an edge to one point and then again identifying the new smallest and biggest elements yields the cycle completion. Third, the restriction of the cycle completion to the set of cuts is already described in [41].

We will now give a precise construction which gives a linear order $D(L)$ starting from a linear order L and is very similar to both the Dedekind completion of L and the pseudo-line $L(L)$ as in [5, Definition 4.1]. Similarly to the Dedekind completion, $D(L)$ consists of initial segments of L and of the elements of L itself. But here an element l of L is not identified with an initial segment of L. The construction of $D(L)$ can be obtained from the pseudo-line $L(L)$ by replacing all the intervals $(0,1) \times \{l\}$ by just l. The topology of the pseudo-line is not needed in the context of this paper.

Example A.1.9. (See also [5]) Let L be a linear order on a set S and let $V(L)$ be the set of *initial segments of* L, i.e. subsets S' of S which satisfy that if s is an element of S' and t is an element of S with $t < s$ then also $t \in S'$. The subset relation is a natural linear order on $V(L)$. Define a linear order on the disjoint union of S and $V(L)$ by letting $x \leq y$ if either both x and y are contained in S and $x \leq y$ in L or both are contained in $V(L)$ and $x \leq y$ in $V(L)$ or $x \in y$ or $y \in S \setminus x$. Denote the resulting linear order on $S \cup V(L)$ by $D(L)$. The smallest element of $D(L)$ is the empty set and the biggest element of $D(L)$ is S. Denote $S \cup (V(L) \setminus \{S\})$ by $V'(L)$, the restriction of $D(L)$ to $V'(L)$ by $D'(L)$ and the cyclic order induced by $D'(L)$ by $Z(L)$. For every element s of S, the set $\{t \in S \colon t < s\}$ is the predecessor and the set $\{t \in S \colon t \leq s\}$ is the successor of s in $D(L)$.

Every subset of $D(L)$ has a supremum and an infimum in $D(L)$, which can be seen as follows: Given a subset V' of $V(L)$, the set $\bigcup V'$ is an initial segment of S and is the supremum of V' both in $D(L) \restriction V(L)$ and in $D(L)$. Similarly the set $\bigcap V'$ is the infimum of V' in $D(L) \restriction V(L)$ and $D(L)$. So in order to show that every subset of $S \cup V(L)$ has a supremum and an infimum in $D(L)$, it suffices to consider subsets S' of S, and by symmetry it suffices to show that S' has a supremum in $D(L)$. The set $\{s \in S | \exists\, t \in S' : s \leq t\}$, denoted by S'', is an initial segment of S which is an upper bound of S'. Also, no proper subset of S'' is an upper bound of S'. So if S' has an upper bound in $D(L)$ which is less than S'', then that upper bound is contained in S. In particular, as $D(L)$ contains between any two elements of S at least one element of $V(L)$, there is at most one upper bound of S' which is less than S''. Thus S' has a supremum in $D(L)$.

Lemma A.1.10. *Let L and K be cuts of a cyclic order Z on a set S with at least two elements such that $K = (L \restriction (S \setminus S')) \oplus (L \restriction S')$ for an initial segment*

S' of L. Then the map

$$V'(L) \to V'(K), \quad x \mapsto \begin{cases} x & x \in S \\ x \cup (S \setminus S') & x \in V'(L) \setminus S,\ x \subsetneq S' \\ x \setminus S' & x \in V'(L) \setminus S,\ S' \subseteq x \end{cases}$$

is the unique isomorphism of $Z(L)$ and $Z(K)$ which preserves S.

Proof. The map $F_1 : V'(L) \to V'(L \restriction S') \cup V'(L \restriction (S \setminus S'))$ which maps elements of S to themselves, initial segments which are properly contained in S' to themselves and initial segments I containing S' to $I \setminus S'$ is an isomorphism of the linear orders $D'(L)$ and $D'(L \restriction S') \oplus D'(L \restriction (S \setminus S'))$. Similarly the map $F_2 : V'(K) \to V'(L \restriction S') \cup V'(L \restriction (S \setminus S'))$ which maps every elements of S to themselves, initial segments properly contained in $S \setminus S'$ to themselves and initial segments I containing $S \setminus S'$ to $I \cap S'$ is an isomorphism of the linear orders $D'(K)$ and $D'(L \restriction (S \setminus S')) \oplus (L \restriction S')$. Thus the map given in the lemma, which equals $F_2^{-1} \circ F_1$, is an isomorphism of $Z(L)$ and $Z(K)$.

Let F and G be two isomorphisms of $Z(L)$ and $Z(K)$ which preserve S. Assume for a contradiction that there is $v \in V'(L)$ such that $F(v)$ is less than $G(v)$ in $D'(K)$. As F and G both are bijective and preserve S, $F(v)$ and $G(v)$ are both contained in $V'(K) \setminus S$. Thus there are elements $s \in G(v) \setminus F(v)$ and $t \in S \setminus G(v)$. Then $(F(t), F(v), F(s))$ equals $(t, F(v), s)$ and is thus contained in $Z(K)$. Because F is monotone, this implies that $(t, v, s) \in Z(L)$. But similarly $(s, G(v), t) \in Z(K)$ and thus $(s, v, t) \in Z(L)$, a contradiction. □

Corollary A.1.11. *Let Z be a cyclic order on set S and let L and K be cuts of Z. Then there is a unique isomorphism of $Z(L)$ and $Z(K)$ which preserves S.*

Proof. Let $S' \subseteq S$ such that $K = (L \restriction (S \setminus S')) \oplus (L \restriction S')$ and such that S' is an initial segment of L. Such a set exists by Theorem A.1.8. Then the statement follows from Lemma A.1.10. □

Lemma A.1.12. *Let Z be a cyclic order on a non-empty set S and let $\mathcal{V}$ be the set of cuts of Z. For every cut $L \in \mathcal{V}$ denote the map $V'(L) \to S \cup \mathcal{V}$ which maps every element of S to itself and every initial segment S' to $(L \restriction (S \setminus S')) \oplus (L \restriction S')$ by η_L. Also denote $\{(\eta_L(a), \eta_L(b), \eta_L(c)) \colon (a, b, c) \in Z(L)\}$ by T_L. Then T_L is a cyclic order on $S \cup \mathcal{V}$ which does not depend on the choice of L.*

Proof. By Theorem A.1.8 the maps η_L are surjective, so they are bijections between $V'(L)$ and $S \cup \mathcal{V}$. Thus every T_L arises from $Z(L)$ by renaming the elements of $V'(L)$ and thus is a cyclic order on $S \cup \mathcal{V}$, and η_L is an isomorphism of $Z(L)$ and T_L. Let L and K be elements of $\mathcal{V}$ and let S' be an initial segment of L such that $K = (L \restriction (S \setminus S')) \oplus (L \restriction S')$ (such a segment exists by Theorem A.1.8). Denote the unique isomorphism of $Z(L)$ and $Z(K)$ preserving S, which exists by Lemma A.1.10, by F. Then $\eta_K \circ F(s) = \eta_L(s)$ for all $s \in S$. Also, for all initial segments I of L which are properly contained in S',

$$\begin{aligned}\eta_K \circ F(I) = \eta_K(I \cup (S \setminus S')) &= (K \restriction (S' \setminus I)) \oplus (K \restriction (I \cup (S \setminus S'))) \\ &= (L \restriction (S' \setminus I)) \oplus (L \restriction (S \setminus S')) \oplus (L \restriction I) \\ &= (L \restriction (S \setminus I)) \oplus (L \restriction I) = \eta_L(I),\end{aligned}$$

and similarly for all initial segments I of L which contain S'

$$\begin{aligned}\eta_K \circ F(I) = \eta_K(I \setminus S') &= (K \restriction (S \setminus (I \setminus S'))) \oplus (K \restriction (I \setminus S')) \\ &= (L \restriction (S \setminus I))) \oplus (L \restriction S') \oplus (L \restriction (I \setminus S')) \\ &= (L \restriction (S \setminus I)) \oplus (L \restriction I) = \eta_L(I).\end{aligned}$$

So $\eta_K \circ F = \eta_L$ and thus $\eta_K \circ F \circ \eta_L^{-1}$ is the identity. But $\eta_K \circ F \circ \eta_L^{-1}$ is also a composition of isomorphisms of cyclic orders and thus the identity is an isomorphism of T_L and T_K, so $T_L = T_K$. $\square$

Given a cyclic order Z on set S and a cut L of Z, the previous lemma shows that T_L only depends on Z and not on L. T_L will be a cycle completion of Z. From now on, denote T_L by $\mathcal{Z}(Z)$ and its ground set by $\mathcal{S}(Z)$.

The next lemma shows that $\mathcal{Z}(Z)$ is really a cycle completion of a cyclic order Z. As a result, cycle completions of cyclic orders exists.

Lemma A.1.13. *Let Z be a cyclic order on set S. Then for every non-trivial interval I of S there are unique elements v and w of $\mathcal{Z}(Z) \setminus S$ such that $I = [v, w] \cap S$.*

Proof. Let L be a cut of Z such that I is an interval of L and such that some element of S is bigger than all elements of I in L. By construction of $D(L)$ there are unique elements v and w of $V'(L) \setminus S$ such that $I = [v, w] \cap S$ in $D(L)$. Then v and w are also the unique elements of $V'(L) \setminus S$ such that $I = [v, w] \cap S$ in $Z(L)$, and thus $\eta_L(v)$ and $\eta_L(w)$ are the unique elements of $\mathcal{S}(Z)$ such that $I = [\eta_L(v), \eta_L(w)] \cap S$ in $\mathcal{Z}(Z)$. $\square$

Given a cyclic order Z on set S, $\mathcal{Z}(Z)$ clearly has the property that for distinct elements v and w of $\mathcal{Z}(Z) \setminus S$ the interval $[v, w] \cap S$ is non-trivial. This property holds for cycle completions in general, as the following rephrasing of Lemma 3.4.4 shows:

Lemma A.1.14. *Let Z be a cyclic order on a set S with at least two elements and let T on the set R be a cycle completion of Z. Then for distinct v and w in $R \setminus S$ the interval $[v, w] \cap S$ of S is non-trivial.*

Proof. As S has at least two elements, it has a non-trivial interval and thus there are x and y in $R \setminus S$ such that $[x, y] \cap S$ is a non-trivial interval of S.

If $[v, y] \cap S$ is a trivial interval, then one of $[v, y] \cap S$ and $[y, v] \cap S$ is empty. Suppose $v \neq y$. If $[v, y] \cap S$ is empty, then $x \notin [v, y]$, implying that $v \in [x, y]$ and $[x, v] \cap S = [x, y] \cap S$, a contradiction. Similarly, if $[y, v] \cap S$ is empty, then also $[x, v] \cap S = [x, y] \cap S$, a contradiction to T being a cycle completion. Hence if $[v, y] \cap S$ is a trivial interval of S then $v = y$.

Thus either $[v, y] \cap S$ is a non-trivial interval of S, or $v = y$ which in particular implies that $[v, x] \cap S$ is a non-trivial interval of S. As $v \neq w$, similarly $[v, w] \cap S$ is a non-trivial interval of S. □

In this paper, cycle completions are used as index sets for k-pseudoflowers, and they are related to each other via surjective monotone maps that map only cuts to cuts. The following lemmas establish a few basic facts about such monotone maps. In particular the following lemma is a rephrasing of Lemma 3.4.6.

Lemma A.1.15. *Let Z and Z' be cyclic orders on sets S and S' with at least two elements. Let T and T' be cycle completions of Z and Z' on sets R and R'. Let $F : R \to R'$ be surjective and monotone such that $F(S) \subseteq S'$. Then there is for every $v' \in R' \setminus S'$ exactly one $v \in R$ with $F(v) = v'$, and there is for every $s' \in S'$ some $s \in S$ with $F(s) = s'$.*

Proof. Every interval of T with at least two elements contains at least one element of S. As $F^{-1}(v')$ does not contain elements of S, it has at most one element. Because F is surjective, $F^{-1}(v')$ contains exactly one element v.

Let t' and q' be the predecessor and successor of s' in T'. Then $[t', q']$ taken in T' consists of t', q' and s'. Let t and q be the unique elements of R such that $F(t) = t'$ and $F(q) = q'$, and let s be an element of $[t, q]$ taken in T which is contained in S. As F is monotone and $F(s)$, q' and t' are pairwise disjoint, $F(s) \in [t', q']$ in T'. Thus $F(s) = s'$. □

So F naturally induces two other monotone maps: The restriction of F to S is surjective and monotone, and $g : R' \setminus S' \to R \setminus S$ which maps every element to its unique preimage under F is injective and monotone. In the other direction, all surjective monotone f from Z to Z' are derived from a surjective monotone map $T \to T'$ with $F(S) \subseteq S'$.

Lemma A.1.16. *Let Z and Z' be cyclic orders on sets S and S' respectively, each with at least two elements, and let T and T' by cycle completions of Z and Z' on sets R and R' respectively. Let $F : R \to R'$ be surjective and monotone with $F(S) \subseteq S'$. Then for all elements v and w of R such that $F(v)$ and $F(w)$ are distinct elements of $R' \setminus S'$ the equations $F^{-1}((F(v), F(w))) = (v, w)$ and $F^{-1}([F(v), F(w)]) = [v, w]$ hold.*

Proof. By Lemma A.1.15, v is the only element of R which is mapped to $F(v)$ by F and similarly for w. So for all $x \in R - v - w$, $(v, x, w) \in T$ if and only if $(F(v), F(x), F(w)) \in T'$ and thus the two equations hold. □

Together with Lemma A.1.15, the previous lemma shows Lemma 3.4.7. With its help, it is now possible to show the following phrasing of Lemma 3.4.5.

Lemma A.1.17. *Let Z and Z' be cyclic orders on sets S and S' and let T and T' be cycle completions of Z and Z' respectively on sets R and R'. Let $f : S \to S'$ be surjective and monotone. If S' has at least two elements, then there is a unique surjective monotone map F from T to T' such that the restriction of F to S equals f.*

Proof. For $v \in R \setminus S$, define $F(v)$ as follows: If there are s and t in S such that $v \in [s, t]$ and $[s, t] \cap S \subseteq f^{-1}(s')$ for some $s' \in S'$, then let $F(v) = s'$. This is well defined as there is at most one such s'. Otherwise let s and t be elements of S with $f(s) \neq f(t)$ and $v \in [s, t]$. Then $S_1 := f([s, v] \cap S)$ and $S_2 := f([v, t] \cap S)$ are intervals of S', and they non-trivial because they are disjoint. So they can be written uniquely as $S_1 = [w_1, w_1'] \cap S'$ and $S_2 = [w_2, w_2'] \cap S'$ for elements w_1, w_1', w_2, w_2' of $R' \setminus S'$. Then $w_1' = w_2$, and this element of R' does not depend on the choice of s and t. Let $F(v) = w_1'$. In particular $F(v) \in R' \setminus S'$.

In order to show that F is surjective, consider $v' \in R' \setminus S'$. Let s' and t' be distinct elements of S' with $v' \in [s', t']$ and let s and t be elements of S with $f(s) = s'$ and $f(t) = t'$. Then $f^{-1}([s', v'] \cap S')$ is a non-trivial interval of S and thus can be written uniquely as $[v, w] \cap S$ for elements v and w of $R \setminus S$. As $[s, v] \cap S \subseteq f^{-1}([v', t'] \cap S')$ and $[v, t] \cap S \subseteq f^{-1}([v', t'] \cap S')$, this implies that $F(v) = v'$.

In order to show that F is monotone, consider $s, t, v \in R$ such that the triple $(F(s), F(v), F(t))$ is contained in T'. For an element r of S', $F^{-1}(r) \setminus S$ only contains elements of $R \setminus S$ whose image under F is defined via the first case, which implies that $F^{-1}(r)$ is an interval of R. So in the case where all three elements $F(s)$, $F(t)$ and $F(v)$ are contained in S', $(s, v, t) \in T$ follows from the fact that f is monotone. As between any two elements of $R' \setminus S'$ there is some element of S', in order to show $(s, v, t) \in T$ it suffices to consider the case that $F(s)$ and $F(t)$ are contained in S' but $F(v)$ is not. Again, because $F^{-1}(r)$ is an interval of R for every $r \in S'$, it suffices to consider the case that both s and t are contained in S. The fact that $F(v)$ is not contained in S' implies that v is not contained in S and that $F(v)$ is defined via the second case. In particular if $v \in [t, s]$ then $F(v) \in [f(t), f(s)] = [F(t), F(s)]$, which is impossible. Hence $(s, v, t) \in T$. $\square$

By the previous lemma, given two cycle completions of a cyclic order Z on set S, the identity on S can be extended uniquely to a surjective monotone map between the cycle completions, showing that cycle completions are essentially unique. Together with the existence of cycle completions proved above this proves Lemma 3.4.3.

A.2 Summary

In this dissertation we introduce path spaces, an infinitary generalization of graphs, which allows us to prove general statements for a variety of path-like objects, including topological arcs in Hausdorff spaces. We spend the rest of the dissertation proving results about them. This happens over the course of four chapters.

In Chapter 1 we consider general connectivity theory. By using different types of alternating paths, we prove two main results. The first is a version of Menger's theorem, showing that the maximum number disjoint of paths between two sets is equal to the minimum size of a separator assuming the first number is finite. The second result shows that we can find a similar duality, though with a more complex witness, for the maximum number of disjoint paths starting and ending in some set. This is a theorem of Gallai for graphs and corresponds to the base case of Mader's theorem.

In Chapter 2 we consider tree-like decompositions of path spaces, making use of the recent theory of separation systems. We begin with a decomposition of connected path spaces into blocks and a decomposition of 2-connected path spaces into 2-blocks, which are parts whose torsos are 3-connected or cycles. Then we move to general widths and use the concept of necklaces to show that every path space of high tree width contains a large grid minor.

In Chapter 3 we note that separations systems of path spaces are limit-closed and analyze in general limit-closed vertex separation systems. We start with a short proof that limit-closed profiles can distinguished by a tree set. Afterward we investigate the flowers of limit-closed separation systems and prove in particular the existence of maximal such objects.

In Chapter 4 we consider questions of ubiquity, more specifically we ask whether path spaces contain infinitely many disjoint copies of a certain substructure whenever they contain arbitrarily many disjoint copies. We first prove that cycles and more generally ladder-like structures are ubiquitous and then apply the grid theorem to show that this applies to minor embeddings of finite planar graphs more generally. We also talk about a notion of edge-ubiquity and show that cycles have this property.

A.3 Zusamenfassung

Diese Dissertation führt Wegräume ein, eine Verallgemeinerung von Graphen, die unendliche Wege zulässt. Mithilfe dieser können wir Aussagen gleichzeitig für verschiedene wegartige Objekte beweisen, darunter topologische Bögen in Hausdorffräumen. Im weiteren Verlauf der Dissertation zeigen wir im Laufe von vier Kapiteln ebensolche Resultate.

In Kapitel 1 betrachten wir Zusammenhang im Allgemeinen. Wir verwenden zwei verschiedene Arten von alternierenden Wegen, um zwei Hauptresultate zu beweisen. Das erste ist eine Version des Satzes von Menger; wir beweisen genauer, dass die maximale Anzahl von disjunkten Wegen zwischen zwei Mengen gleich der minimalen Größe eines Trenners ist, wenn die erste Zahl endlich ist. Das zweite Resultat zeigt eine ähnliche Dualität, wenn auch mit einem komplizierteren Gegenpart, für die maximale Anzahl von disjunkten Wegen die in einer bestimmten Menge beginnen und enden. Für Graphen ist das ein Satz von Gallai, der dem Induktionsanfang vom Satz von Mader entspricht.

In Kapitel 2 betrachten wir baumartige Zerlegungen von Wegräumen, wobei wir die noch junge Theorie von Separationssytemen verwenden. Zunächst konstruieren wir Zerlegungen von zusammenhängenden Wegräumen in ihre Blöcke und von 2-zusammenhängenden Wegräumen in 2-Blöcke, das heißt Teile der Zerlegung deren Torsos 3-zusammenhängend oder Kreise sind. Danach verwenden wir Halsketten, um das zu zeigen, dass jeder Wegraum mit hoher Baumweite einen großen Gitterminor enthält.

In Kapitel 3 bemerken wir, dass Separationssysteme von Wegräumen unter Limites abgeschlossen sind und betrachten solche Separationssystem im Allgemeinen. Zunächst zeigen wir, dass Profile, die unter Limites abgeschlossen sind, durch eine geschachtelte Menge von Teilungen unterschieden werden können. Danach betrachten wir Blumen von Separationssystemen und zeigen unter anderem die Existenz von maximalen Blumen unter Annahme dieser Abschlußeigenschaft.

In Kapitel 4 untersuchen wir Ubiquität, das heißt wir fragen ob Wegräume, die beliebig viele disjunkte Kopien von bestimmten Unterstrukturen enthalten auch unendlich viele solche Kopien enthalten. Wir zeigen zunächst, dass Kreise und allgemeiner leiterartige Strukturen diese Eigenschaft haben und wenden dann den Gittersatz an um allgemeiner die Ubiquität von Einbettungen von endlichen, planaren Graphen als Minor zu beweisen. Weiterhin geben wir eine Definition von Kantenubiquität und zeigen, dass Kreise dies erfüllen.

A.4 Related publications

The only published preprint related to this thesis is [35], an unpublished expanded version of which roughly corresponds to Chapter 1.

Chapter 2 consists of two unfinished drafts, [33] making up the first half and [34] the second.

Similarly, Chapter 3 consists of the two drafts [22] and [21].

Finally, Chapter 4 corresponds to the draft [36].

A.5 Contributions

Chapter 1, Chapter 2 and Chapter 4 are based on drafts of which I am sole author. Chapter 3 is joint work with Ann-Kathrin Elm. While the research and writing were done collaboratively, in the following I will give a rough breakdown of our contributions to it.

To start with I wrote Section 3.1 incorporating some suggestions by Ann-Kathrin. For Section 3.2 I first proved a version of Theorem 3.2.6 and wrote a first draft, Ann-Kathrin introduced the idea of the function $f_{\mathcal{P}}$ and rewrote and simplified the proof using it and finally I rewrote the section again. Section 3.3 is mostly Ann-Kathrin's work. While we came up with our definition of pseudoflower together, Section 3.4 is otherwise mostly her work as well (including the accompanying appendix). I first proved an extension theorem for pseudoflowers and wrote the first draft of Section 3.5 and Ann-Kathrin introduced the notion of separations being anchored, improved the proof using it and wrote the current version of the section. For Section 3.6 I had the idea of using inverse limits and Ann-Kathrin worked out the details and wrote the section. Section 3.7 is mostly my work. Finally, I wrote Section 3.8 based in part on a common list of bullet points.

A.6 Acknowledgement

First I thank my supervisor Nathan Bowler for his support and guidance. You were always there with advice when I needed it and helped my research tremendously. I also thank my second assessor Max Pitz and the rest of my examination commission for their willingness to be a part of this project.

Next I would like to thank Ann-Kathrin Elm for collaborating with me on multiple research projects. It was great working with you. I also thank you for thoroughly proofreading some of my drafts. In addition, I thank Florian Reich for taking up some of the questions left open here and discussing his progress with me.

I further thank Reinhard Diestel, whose graph theory course first excited me for the field. You organized our research group and the results speak for themselves: a great work climate that endured even through the pandemic. Of course everyone in the group, both current and former members, was a part of this and I am grateful to you all.

Finally, I thank my friends for all the fun times and my family for all their love and support. I could not have done this without you.

Bibliography

[1] Ron Aharoni and Eli Berger. "Menger's theorem for infinite graphs". In: *Inventiones mathematicae* 176 (2009), pp. 1–62. DOI: 10.1007/s00222-008-0157-3.

[2] Elad Aigner-Horev, Reinhard Diestel, and Luke Postle. "The structure of 2-separations of infinite matroids". In: *J. Combin. Theory, Series B* 116 (2016), pp. 25–56. DOI: 10.1016/j.jctb.2015.05.011.

[3] Jeremy Aikin and James Oxley. "The structure of crossing separations in matroids". In: *Adv. Appl. Math.* 41.1 (2008), pp. 10–26. DOI: 10.1016/j.aam.2007.05.004.

[4] Jeremy Aikin and James Oxley. "The structure of crossing separations in matroids". English. In: *Adv. Appl. Math.* 41.1 (2008), pp. 10–26. DOI: 10.1016/j.aam.2007.05.004.

[5] Nathan Bowler, Johannes Carmesin, and Robin Christian. "Infinite graphic matroids Part I". In: *Combinatorica* 38 (2013), pp. 305–339. DOI: 10.1007/s00493-016-3178-3.

[6] Nathan Bowler et al. *Ubiquity in graphs I: Topological ubiquity of trees*. 2018. arXiv: 1806.04008 [math.CO].

[7] Henning Bruhn, Reinhard Diestel, and Maya Stein. "Menger's theorem for infinite graphs with ends". In: *J. of Graph Theory* 50 (2005), pp. 199–211. DOI: 10.1002/jgt.20108.

[8] Henning Bruhn and Xingxing Yu. "Hamilton Cycles in Planar Locally Finite Graphs". In: *SIAM J. Discrete Math.* 22 (2008), pp. 1381–1392. DOI: 10.1137/050631458.

[9] W. Bula, J. Nikiel, and E. D. Tymchatyn. "The Königsberg Bridge Problem for Peano Continua". In: *Can. J. Math.* 46 (1994), pp. 1175–1187. DOI: 10.4153/CJM-1994-066-2.

[10] Ben Clark and Geoff Whittle. "Tangles, trees, and flowers". In: *J. Comb. Theory, Ser. B* 103 (2013), pp. 385–407. DOI: 10.1016/j.jctb.2013.03.002.

[11] Ben Clark and Geoff Whittle. "Tangles, trees, and flowers". English. In: *J. Comb. Theory, Ser. B* 103.3 (2013), pp. 385–407. DOI: 10.1016/j.jctb.2013.03.002.

[12] J. van Dalen and E. Wattel. "A topological characterization of ordered spaces". In: *Gen. Topol. Its Appl.* 3.4 (1973), pp. 347–354. DOI: 10.1016/0016-660x(73)90022-6.

[13] Reinhard Diestel. "Abstract Separation Systems". In: *Order* 35 (2018), pp. 157–170. DOI: 10.1007/s11083-017-9424-5.

[14] Reinhard Diestel. *Graph Theory.* 5th ed. Springer, 2017. DOI: 10.1007/978-3-662-53622-3.

[15] Reinhard Diestel, Fabian Hundertmark, and Sahar Lemanczyk. "Profiles of Separations: in Graphs, Matroids, and Beyond". In: *Combinatorica* 39 (2018), pp. 37–75. DOI: 10.1007/s00493-017-3595-y.

[16] Reinhard Diestel and Daniela Kühn. "On Infinite Cycles I". In: *Combinatorica* 24.1 (2004), pp. 69–89. DOI: 10.1007/s00493-004-0005-z.

[17] Reinhard Diestel and Jakob Kneip. "Profinite Separation Systems". In: *Order* 37 (2020), pp. 179–205. DOI: 10.1007/s11083-019-09499-y.

[18] Reinhard Diestel and Daniela Kühn. "Topological Paths, Cycles and Spanning Trees in Infinite Graphs". In: *Eur. J. Comb.* 25.6 (2004), pp. 835–862. DOI: 10.1016/j.ejc.2003.01.002.

[19] Reinhard Diestel et al. "Highly Connected Sets and the Excluded Grid Theorem". In: *J. Combin. Theory, Series B* 75.1 (1999), pp. 61–73. DOI: 10.1006/jctb.1998.1862.

[20] Christian Elbracht, Jakob Kneip, and Maximilian Teegen. "Trees of tangles in infinite separation systems". In: *Math. Proc. Camb. Philos. Soc.* (2021). DOI: 10.1017/S0305004121000530.

[21] Ann-Kathrin Elm and Hendrik Heine. *Flowers in graph-like spaces.* in preparation.

[22] Ann-Kathrin Elm and Hendrik Heine. *Limit-closed profiles*. in preparation.

[23] Benjamin Espinoza, Paul Gartside, and Max Pitz. "Graph-Like Compacta: Characterizations and Eulerian Loops". In: *J. Graph Theory* 95 (2016), pp. 209–239. DOI: 10.1002/jgt.22558.

[24] Tibor Gallai. "Maximum-minimum Sätze und verallgemeinerte Faktoren von Graphen". In: *Acta Math. Hung.* 12.1-2 (1961), pp. 131–173. DOI: 10.1007/BF02066678.

[25] Jim Geelen and Benson Joeris. *A generalization of the Grid Theorem*. 2016. arXiv: 1609.09098 [math.CO].

[26] William J. Gilbert. "Menger's Theorem for topological spaces". In: *Fundam. Math.* 75 (1972), pp. 291–295. DOI: 10.4064/fm-75-3-291-295.

[27] J. Pascal Gollin and Karl Heuer. *Characterising k-connected sets in infinite graphs*. 2020. arXiv: 1811.06411 [math.CO].

[28] Jochen Pascal Gollin and Jakob Kneip. "Representations of Infinite Tree Sets". In: *Order* 38 (2021). DOI: 10.1007/s11083-020-09529-0.

[29] Martin Grohe. "Quasi-4-Connected Components". In: *ICALP 2016*. Vol. 55. 2016, 8:1–8:13. DOI: 10.4230/LIPIcs.ICALP.2016.8.

[30] S.L. Gulden, W.M. Fleischman, and J.Harley Weston. "Linearly ordered topological spaces". In: *Proc. Am. Math. Soc.* 24 (1970), pp. 197–203. DOI: 10.2307/2036727.

[31] R. Halin. "Über die Maximalzahl fremder unendlicher Wege in Graphen". In: *Math. Nachr.* 30 (1965), pp. 63–85. DOI: 10.1002/MANA.19650300106.

[32] Rudolf Halin. "S-functions for graphs". In: *J. Geom.* 8 (1976), pp. 171–186. DOI: 10.1007/BF01917434.

[33] Hendrik Heine. *Connectivity in path spaces II: Tree-like decompositions*. in preparation.

[34] Hendrik Heine. *Connectivity in path spaces III: Grid theorem*. in preparation.

[35] Hendrik Heine. *Path spaces I: A Menger-type result*. 2020. arXiv: 2007.09709 [math.CO].

[36] Hendrik Heine. *Topological ubiquity*. in preparation.

[37] Maya Jakobine Stein. "Arboricity and tree-packing in locally finite graphs". In: *J. Combin. Theory, Series B* 96.2 (2006), pp. 302–312. DOI: 10.1016/j.jctb.2005.08.003.

[38] Dénes Kőnig. “Gráfok és mátrixok”. In: *Matematikai és Fizikai Lapok* 38 (1931), pp. 116–119.

[39] L. Lovász. “Normal hypergraphs and the perfect graph conjecture”. In: *Discrete Mathematics* 2.3 (1972), pp. 253–267. DOI: 10.1016/0012-365X(72)90006-4.

[40] Vitezslav Novak. “Cyclically ordered sets”. English. In: *Czech. Math. J.* 32 (1982), pp. 460–473.

[41] Vítězslav Novák. “Cuts in cyclically ordered sets”. English. In: *Czech. Math. J.* 34 (1984), pp. 322–333.

[42] James Oxley, Charles Semple, and Geoff Whittle. “The structure of the 3-separations of 3-connected matroids”. In: *J. Comb. Theory, Ser. B* 92.2 (2004), pp. 257–293. DOI: 10.1016/j.jctb.2004.03.006.

[43] J.S. Pym. “A proof of the linkage theorem”. In: *J. Math. Anal. Appl.* 27.3 (1969), pp. 636–638. DOI: 10.1016/0022-247X(69)90142-5.

[44] R. Bruce Richter. “Decomposing Infinite 2-Connected Graphs into 3-Connected Components”. In: *Electr. J. Comb.* 11 (2004). DOI: 10.37236/1778.

[45] Neil Robertson and P.D Seymour. “Graph minors. III. Planar tree-width”. In: *J. Combin. Theory, Series B* 36.1 (1984), pp. 49–64. DOI: 10.1016/0095-8956(84)90013-3.

[46] Neil Robertson and P.D Seymour. “Graph minors. V. Excluding a planar graph”. In: *J. Combin. Theory, Series B* 41.1 (1986), pp. 92–114. DOI: 10.1016/0095-8956(86)90030-4.

[47] Neil Robertson and P.D Seymour. “Graph minors. X. Obstructions to tree-decomposition”. In: *J. Comb. Theory, Ser. B* 52.2 (1991), pp. 153–190. DOI: 10.1016/0095-8956(91)90061-N.

[48] Alexander Schrijver. “A Short Proof of Mader’s S-Paths Theorem”. In: *J. Combin. Theory, Series B* 82.2 (2001), pp. 319–321. DOI: 10.1006/jctb.2000.2029.

[49] Carsten Thomassen and Antoine Vella. “Graph-like continua, augmenting arcs, and Menger’s theorem”. In: *Combinatorica* 28 (2008), pp. 595–623. DOI: 10.1007/s00493-008-2342-9.

[50] W. T. Tutte. *Connectivity in Graphs*. University of Toronto Press, 1966. DOI: 10.3138/9781487584863.

[51] G. T. Whyburn. “The Cyclic and Higher Connectivity of Locally Connected Spaces”. In: *Am. J. Math.* 53.2 (1931), pp. 427–442. DOI: 10.2307/2370795.

Eidesstattliche Versicherung

Hiermit versichere ich an Eides statt, dass ich die vorliegende Dissertation selbst verfasst und keine anderen als die angegebenen Hilfsmittel benutzt habe. Darüber hinaus versichere ich, dass diese Dissertation nicht in einem früheren Promotionsverfahren eingereicht wurde.