Tree-structure in separation systems and infinitary combinatorics

Habilitationsschrift

vorgelegt im

Fachbereich Mathematik
Universität Hamburg
von

Joshua Erde
Hamburg, im Juni 2019

Overview

This habilitation thesis consists of twelve papers which I have written together with a number of co-authors during my time as a postdoc at the University of Hamburg, which started in October 2014.

The papers can be naturally split into two topics, tree-structure in separation systems and infinitary combinatorics, which constitute the two parts of the thesis.

Part I

Tree structure in separation systems

The concept of a tree-decomposition of a graph was first introduced by Halin [74], and later developed by Roberston and Seymour [109] as part of their work on graph minors. Graphs with tree-decompositions of low width have a global tree-like structure, which makes them simpler to work with. For example many difficult algorithmic problems are much simpler on graphs with bounded tree-width. It is therefore useful to know what can be said about the structure of a graph which does not admit a tree-decomposition of low width. In a shift of paradigm, Roberston and Seymour [110] introduced the concept of a tangle to represent a highly connected structure in a graph. Rather than a concrete structure in the graph itself, a tangle is a consistent way to orient the low order separators of a graph. It can be shown that the existence of a high order tangle in a graph is dual to the existence of a tree-decomposition of low width. Furthermore, Roberston and Seymour showed that there is a tree-decomposition of a graph which distinguishes all of its maximal tangles, and so we can view the global structure of the graph as being built of of this collection of tangles in a tree-like manner.

Recently Diestel developed an unified abstract framework (see for example [44]), called *separation systems*, in which these and other results about tree-structure could be expressed. This framework is broad enough to encompass many of the varied types of tree-decompositions that have been considered in the literature, and many others beyond that, allowing for unified proofs of many know results, as well as allowing one to apply the theory of tree-decompositions to other mathematical structures, combinatorial or otherwise.

Refining a tree-decomposition which distinguishes tangles

Roberston and Seymour introduced tangles of order k as objects representing highly connected parts of a graph and showed that every graph admits a tree-decomposition of adhesion $< k$ in which each tangle of order k is contained in a different part. Recently, Carmesin, Diestel, Hamann and Hundertmark [35] showed that such a tree-decomposition can be constructed in a canonical way, which makes it invariant under automorphisms of the graph. These canonical tree-decompositions necessarily have parts which contain no tangle of order k. We call these parts inessential. Diestel asked what could be said about the structure of the inessential parts. In this paper we show that the torsos of the inessential parts in these tree-decompositions have branch width $< k$, allowing us to further refine the canonical tree-decompositions, and also show that a similar result holds for k-blocks.

This paper appears in the SIAM Journal on Discrete Mathematics [58].

Duality theorems for blocks and tangles in graphs

We prove a duality theorem applicable to a a wide range of specialisations, as well as to some generalisations, of tangles in graphs. It generalises the classical tangle duality theorem of Robertson and Seymour, which says that every graph either has a large-order tangle or a certain low-width tree-decomposition witnessing that it cannot have such a tangle. Our result also yields duality theorems for profiles and for k-blocks. This solves a problem studied, but not solved, by Diestel and Oum [50] and answers an earlier question of Carmesin, Diestel, Hamann and Hundertmark [34].

This paper is joint work with Reinhard Diestel and Philipp Eberenz and appears in the SIAM Journal on Discrete Mathematics[46].

A unified treatment of linked and lean tree-decompositions

There are many results asserting the existence of tree-decompositions of minimal width which still represent local connectivity properties of the underlying graph, perhaps the best-known being Thomas' theorem [124] that proves for every graph G the existence of a linked tree-decompositon of width $\mathrm{tw}(G)$. We prove a general theorem on the existence of linked and lean tree-decompositions, providing a unifying proof of many known results in the field, as well as implying some new results. In particular we prove that every matroid M admits a lean tree-decomposition of width $\mathrm{tw}(M)$, generalizing the result of Thomas.

This paper appears in the Journal of Combinatorial Theory, Series B [59].

Structural submodularity and tangles in abstract separation systems

We prove a tangle-tree theorem and a tangle duality theorem for abstract separation systems $\vec{S}$ that are submodular in the structural sense that, for every pair of oriented separations, $\vec{S}$ contains either their meet or their join defined in some universe $\vec{U}$ of separations containing $\vec{S}$. This holds, and is widely used, if $\vec{U}$ comes with a submodular order function and $\vec{S}$ consists of all its separations up to some fixed order. Our result is that for the proofs of these two theorems, which are central to abstract tangle theory, it suffices to assume the above structural consequence for $\vec{S}$, and no order function is needed.

This paper is joint work with Reinhard Diestel and Daniel Weißauer and appears in the Journal of Combinatorial Theory, Series A [47].

Directed path-decompositions

Many of the tools developed for the theory of tree-decompositions of graphs do not work for directed graphs. In this paper we show that some of the most basic tools do work in the case where the model digraph is a directed path. Using these tools we define a notion of a directed blockage in a digraph and prove a min-max theorem for directed path-width analogous to the result of Bienstock, Roberston, Seymour and Thomas [21] for blockages in graphs. Furthermore, we show that every digraph with directed path width $\geqslant k$ contains each arboresence of order $\leqslant k+1$ as a butterfly minor. Finally we also show that every digraph admits a linked directed path-decomposition of minimum width, extending a result of Kim and Seymour [85] on semi-complete digraphs.

A short derivation of the structure theorem for graphs with excluded topological minors

As a major step in their proof of Wagner's conjecture, Robertson and Seymour [111] showed that every graph not containing a fixed graph H as a minor has a tree-decomposition in which each torso is almost embeddable in a surface of bounded genus. Recently, Grohe and Marx [70] proved a similar result for graphs not containing H as a topological minor. They showed that every graph which does not contain H as a topological minor has a tree-decomposition in which every torso is either almost embeddable in a surface of bounded genus, or has a bounded number of vertices of high degree. We give a short proof of the theorem of Grohe and Marx, improving their bounds on a number of the parameters involved.

This paper is joint work with Daniel Weißauer.

Part II

The Reconstruction conjecture

We say that two graphs G and H are *hypomorphic* if there exists a bijection φ between the vertices of G and H such that the induced subgraphs $G - v$ and $H - \varphi(v)$ are isomorphic for each vertex v of G. Any such bijection is called a *hypomorphism*. We say that a graph G is *reconstructible* if $H \cong G$ for every H hypomorphic to G. The following conjecture, attributed to Kelly and Ulam, is perhaps one of the most famous unsolved problems in the theory of graphs.

Conjecture (The Reconstruction Conjecture)**.** *Every finite graph with at least three vertices is reconstructible.*

For an overview of results towards the Reconstruction Conjecture for finite graphs see the survey of Bondy and Hemminger [23]. Harary [75] proposed the Reconstruction Conjecture for infinite graphs, however Fisher [62] found a counterexample, which was improved to a simpler counterexample by Fisher, Graham and Harary [63]. These graphs, however, contain vertices of infinite degree. A graph is *locally finite* if every vertex has finite degree. Locally finite graphs are a particular simple class of infinite graphs, which in many ways posses similar properties to finite graphs.

Harary, Schwenk and Scott [76] showed that there exists a non-reconstructible locally finite forest. However, they conjectured that the Reconstruction Conjecture should hold for locally finite trees.

Conjecture (The Harary-Schwenk-Scott Conjecture)**.** *Every locally finite tree is reconstructible.*

This conjecture has been verified in a number of special cases. Bondy and Hemminger [22] showed that every tree with at least two but a finite number of ends is reconstructible, and Thomassen [125] showed that this also holds for one-ended trees. Andreae [10] proved that also every tree with countably many ends is reconstructible.

A counterexample to the reconstruction conjecture for locally finite trees

It is well known that not all infinite graphs are reconstructible. However, the Harary-Schwenk-Scott Conjecture from 1972 [76] suggests that all locally finite trees are reconstructible. In this paper, we construct a counterexample to the Harary-Schwenk-Scott Conjecture. Our example also answers four other questions of Nash-Williams [103], Halin and Andreae [12] on the reconstruction of infinite graphs.

This paper is joint work with Nathan Bowler, Peter Heinig, Florian Lehner and Max Pitz and appears in the Bulletin of the London Mathematical Society [28].

Non-reconstructible locally finite graphs

Nash-Williams [102, 104] proved that all locally finite graphs with a finite number $\geqslant 2$ of ends are reconstructible, and asked whether locally finite graphs with one end or countably many ends are also reconstructible. In this paper we construct non-reconstructible graphs of bounded maximum degree with one and countably many ends respectively, answering the two questions of Nash-Williams about the reconstruction of locally finite graphs in the negative.

This paper is joint work with Nathan Bowler, Peter Heinig, Florian Lehner and Max Pitz and appears in the Journal of Combinatorial Theory, Series B [29].

The ubiquity conjecture

Let $\triangleleft$ be a relation between graphs, for example the subgraph relation $\subseteq$, the topological minor relation $\leqslant$ or the minor relation $\preccurlyeq$. We say that a graph G is *$\triangleleft$-ubiquitous* if whenever Γ is a graph with $nG \triangleleft \Gamma$ for all $n \in \mathbb{N}$, then one also has $\aleph_0 G \triangleleft \Gamma$, where αG is the disjoint union of α many copies of G.

Two classic results of Halin [71, 72] say that both the ray and the double ray are $\subseteq$-ubiquitous, i.e. any graph which contains arbitrarily large collections of disjoint (double) rays must contain an infinite collection of disjoint (double) rays. However, even quite simple graphs can fail to be $\subseteq$ or $\leqslant$-ubiquitous, see e.g. [7, 130, 91].

However, for the minor relation, no such simple examples of non-ubiquitous graphs are known. Indeed, one of the most important problems in the theory of infinite graphs is the so-called *Ubiquity Conjecture* due to Andreae [13].

Conjecture (The Ubiquity Conjecture)**.** *Every locally finite connected graph is $\preccurlyeq$-ubiquitous.*

In [13], Andreae constructed a graph that is not $\preccurlyeq$-ubiquitous. However, this construction relies on the existence of a counterexample to the well-quasi-ordering of infinite graphs under the minor relation, for which counterexamples are only known with very large cardinality [122]. In particular, it is still an open question whether or not there exists a countable connected graph which is not $\preccurlyeq$-ubiquitous.

Topological ubiquity of trees

The *Ubiquity Conjecture* of Andreae, a well-known open problem in the theory of infinite graphs, asserts that every locally finite connected graph is ubiquitous with respect to the minor relation. In this paper, which is the first of a series of papers making progress towards the Ubiquity Conjecture, we show that all trees are ubiquitous with respect to the topological minor relation, irrespective of their cardinality. This answers a question of Andreae [8] from 1979.

This paper is joint work with Nathan Bowler, Christian Elbracht, Pascal Gollin, Karl Heuer, Max Pitz and Maximilian Teegen.

Ubiquity of graphs with non-linear end structure

In this paper we give a sufficient condition on the structure of the ends of a graph G which implies that G is ubiquitous with respect to the minor relation. In particular this implies that the full grid is ubiquitous with respect to the minor relation.

This paper is joint work with Nathan Bowler, Christian Elbracht, Pascal Gollin, Karl Heuer, Max Pitz and Maximilian Teegen.

Ubiquity of locally finite graphs with extensive tree decompositions

In this paper we show that locally finite graphs admitting a certain type of tree-decomposition, which we call an extensive tree decomposition, are ubiquitous with respect to the minor relation. In particular this includes all locally finite graphs of finite tree-width and locally finite graphs with finitely many ends, all of which are thin.

This paper is joint work with Nathan Bowler, Christian Elbracht, Pascal Gollin, Karl Heuer, Max Pitz and Maximilian Teegen.

Hamilton decompositions of infinite Cayley graphs

A *Hamiltonian cycle* of a finite graph is a cycle which includes every vertex of the graph. A finite graph $G = (V, E)$ is said to have a *Hamilton decomposition* if its edge set can be partitioned into disjoint sets $E = E_1 \dot\cup E_2 \dot\cup \cdots \dot\cup E_r$ such that each E_i is a Hamiltonian cycle in G.

The starting point for the theory of Hamilton decompositions is an old result by Walecki from 1890 according to which every finite complete graph of odd order has a Hamilton decomposition (see [3] for a description of his construction). Since then, this result has been extended in various different ways, and we refer the reader to the survey of Alspach, Bermond and Sotteau [4] for more information.

Hamiltonicity problems have also been considered for infinite graphs, see for example the survey by Gallian and Witte [129]. While it is sometimes not obvious which objects should be considered the correct generalisations of a Hamiltonian cycle in the setting of infinite graphs, for one-ended graphs the undisputed solution is to consider *double-rays*, i.e. infinite, connected, 2-regular subgraphs. Thus, for us a *Hamiltonian double-ray* is then a double-ray which includes every vertex of the graph, and we say that an infinite graph $G = (V, E)$ has a *Hamilton decomposition* if we can partition its edge set into edge-disjoint Hamiltonian double-rays.

One well known conjecture on the existence of Hamilton decompositions for finite graphs concerns Cayley graphs: Given a finitely generated abelian group $(\Gamma, +)$ and a finite generating set S of Γ, the *Cayley graph* $G(\Gamma, S)$ is the multi-graph with vertex set Γ and edge multi-set

$$\{(x, x+g) \,:\, x \in \Gamma, g \in S\}.$$

Conjecture (Alspach [2])**.** *If Γ is an abelian group and S generates G, then the simplification of $G(\Gamma, S)$ has a Hamilton decomposition, provided that it is $2k$-regular for some k.*

Hamilton decompositions of one-ended Cayley graphs

We prove that any one-ended, locally finite Cayley graph with non-torsion generators admits a decomposition into edge-disjoint Hamiltonian (i.e. spanning) double-rays. In particular, the n-dimensional grid $\mathbb{Z}^n$ admits a decomposition into n edge-disjoint Hamiltonian double-rays for all $n \in \mathbb{N}$. We also prove an infinite version of a conjecture of Bermond [18] that whenever G and H are two graphs which admit a decomposition into spanning double-rays, then so does their cartesian product $G \Box H$.

This paper is joint work with Florian Lehner and Max Pitz and has been accepted for publication in the Journal of Combinatorial Theory, Series B [61].

Acknowledgements.

I would like to thank my many co-authors, in particular Max, Florian and Nathan, who have introduced me to many interesting ideas and problems, and without whom many of these papers would not have been possible. I would also like to thank Reinhard Diestel for introducing me to both of the fields of study covered in this thesis, as well as his support and encouragement over the past 5 years. Finally I would like to thank my wife Barbara for coming with me to Germany and supporting me during my time at Hamburg.

Contents

Part I

Tree-structure in separation systems

Chapter 1

Refining a tree-decomposition which distinguishes tangles

1.1 Introduction

A classical notion in graph theory is that of the block-cut vertex tree of a graph. It tells us that if we consider the maximal 2-connected components of a connected graph G then they are arranged in a 'tree-like' manner, separated by the cut vertices of G. A result of Tutte's [126] says that we can decompose any 2-connected graph in a similar way. Broadly, it says that every 2-connected graph can be decomposed in a 'tree-like' manner, so that the parts are separated by vertex sets of size at most 2, and every part, together with the edges in the separators adjacent to it, is either 3-connected or a cycle. We call the union of a part and the edges in the separators adjacent to it the *torso* of the part. In contrast to the first example not every part, or even torso, of this decomposition is 3-connected, and indeed it is easy to show that not every 2-connected graph can be decomposed in this way such that every torso is 3-connected.

It has long been an open problem how best to extend these results for general k, the aim being to decompose a $(k-1)$-connected graph into its 'k-connected components', where the precise meaning of what these 'k-connected components' should be considered to be has varied. Tutte's example shows us that there may be parts of this decomposition which are not highly connected, but rather play a structural role in the graph of linking the highly connected parts together, and further that the highly connected parts of the decomposition may not correspond exactly to k-connected subgraphs.

Whereas initially these 'k-connected components' were considered as concrete structures in the graph itself, Robertson and Seymour [110] radically re-interpreted them as *tangles of order* k, which for brevity we will refer to as *k-tangles*[1]. Instead of being defined in terms of the edges and vertices of a graph, these objects were defined in terms of structures on the set of low-order separations of a graph.

Robertson and Seymour showed that, given any set of distinct k-tangles $T_1, T_2, \dots, T_n$ in a graph G, there is a tree-decomposition $(T, \mathcal{V})$ of G with precisely n parts in which the orientations induced by the tangles T_i on $E(T)$ each have distinct sinks, where we say the tangle is *contained* in this sink. We say that such a tree-decomposition *distinguishes* the tangles $T_1, T_2, \dots, T_n$. They showed further that these tree-decompositions can be chosen so that the separators between the parts are in some way minimal with respect to the tangles considered. We say that such a tree-decomposition distinguishes the k-tangles *efficiently*. If we call the largest size of a separator

[1]Precise definitions of many of the terms in the introduction will be postponed until Section 1.2, where all the necessary background material will be introduced.

in a tree-decomposition the *adhesion* of the tree-decomposition, then in particular their result implies the following:

Theorem 1.1.1 (Robertson and Seymour [110])**.** *For every graph G and $k \geqslant 2$ there exists a tree-decomposition $(T, \mathcal{V})$ of G of adhesion $< k$ which distinguishes the set of k-tangles in G efficiently.*

More recently Carmesin, Diestel, Hamann and Hundertmark [35] described a family of algorithms that can be used to build tree-decompositions which distinguish the set of k-tangles in a graph and are *canonical*, that is, they are invariant under every automorphism of the graph.

Just as in Tutte's theorem, where there were parts of the tree-decomposition whose torsos were not 3-connected, it is easy to show that the tree-decompositions formed in [35] must contain parts which do not contain any k-tangle. Since the general motivation for these tree-decompositions is to decompose the graph into its 'k-connected components' in a way that displays the global structure of the graph, it is natural to ask further questions about the structure of these tree-decompositions. In [33] Carmesin et al. analysed the structure of the trees that the various algorithms given in [35] produced. One particular question that was asked is what can be said about the structure of the parts which do not contain a k-tangle. We will call the parts of a tree-decomposition that contain a k-tangle *essential*, and those that do not *inessential*.

For example, if the whole graph contains no k-tangle, then these canonical tree-decompositions tell us nothing about the graph, as they consist of just one inessential part. However there are theorems which describe the structure of a graph which contains no k-tangle. In the same paper where they introduced the concept of tangles, Roberston and Seymour [110] showed that a graph which contains no k-tangle has branch-width $< k$, and in fact that the converse is also true, a graph with branch-width $\geqslant k$ contains a k-tangle. Having branch-width $< k$ can be rephrased in terms of the existence of a certain type of tree-decomposition (See e.g. [52]). A nice property of these tree-decompositions is that each of the parts is in some sense 'too small' to contain a k-tangle. In this way these tree-decompositions witness that a graph has no k-tangle by splitting the graph into a number of parts, each of which cannot contain a k-tangle and similarly a k-tangle witnesses that a graph does not have such a tree-decomposition.

A natural question to then ask is, do the inessential parts in the tree-decompositions from [35] admit tree-decompositions of the same form, into parts which are too small to contain a k-tangle? If so we might hope to refine these canonical tree-decompositions by decomposing further the inessential parts. By combining these decompositions we would get an overall tree-decomposition of G consisting of some essential parts, each containing a k-tangle in G, and some inessential parts, each of which is 'small' enough to witness the fact that no k-tangle is contained in that part.

We first note that we cannot hope for these refinements to also be canonical. For example consider a graph formed by taking a large cycle C and adjoining to each edge a large complete graph K_n. Then a canonical tree-decomposition which distinguishes the 3-tangles in this graph will contain the cycle C as an inessential part. However there is no canonical tree-decomposition of C with branch-width < 3. Indeed, if such a tree-decomposition contained any of the 2-separations of C as an adhesion set then, since all the rotations of C lie in the automorphism group of G, every rotation of this separation must appear as an adhesion set. However these separations cannot all appear as the adhesion sets in any tree-decomposition, as every pair of vertices in a 2-separation of C are themselves separated by some rotation of that separation.

If we drop the restriction that the refinement be canonical then, at first glance, it might seem like there should clearly be such a refinement. If there is no k-tangle contained in a part V_t in a tree-decomposition, $(T, \mathcal{V})$, then by the theorem of Robertson and Seymour there should be a tree-decomposition of that part with branch-width $< k$. However there is a problem with this naive approach, in that we have no guarantee that we can insert the tree-decomposition of this part into the existing tree-decomposition. In particular it could be the case that this tree-decomposition splits up the separators of the part V_t in $(T, \mathcal{V})$. One way to avoid this problem is to instead consider the torso of the part V_t. If we have a tree-decomposition of the torso we can insert it into the original tree-decomposition, but it is not clear that adding these extra edges can not increase the branch-width of the part. In fact it is easy to find examples where choosing a bad canonical tree-decomposition to distinguish the set of k-tangles in a graph results in inessential parts whose torsos have branch-width $\geqslant k$.

For example consider the following graph: We start with the union of three large complete graphs, K_{N_1}, K_{N_2} and K_{N_3}, for $N_1, N_2, N_3 >> k$. We pick a set of $(k-1)/2$ vertices from each graph, which we denote by X_1, X_2 and X_3 respectively, and join each of these sets completely to a new vertex x. It is a simple check that there are three k-tangles in this graph, corresponding to the three large complete subgraphs. However, consider the following tree-decomposition of the graph into four parts $K_{N_1} \cup X_2$, $K_{N_2} \cup X_3$, $K_{N_3} \cup X_1$ and $X_1 \cup X_2 \cup X_3 \cup \{x\}$. This is a tree-decomposition which distinguishes the k-tangles in the graph, and the part $X_1 \cup X_2 \cup X_3 \cup \{x\}$ is inessential. However the torso of this middle part is a complete graph of order $3(k-1)/2+1$, which can be seen to have branch-width $\geqslant k$.

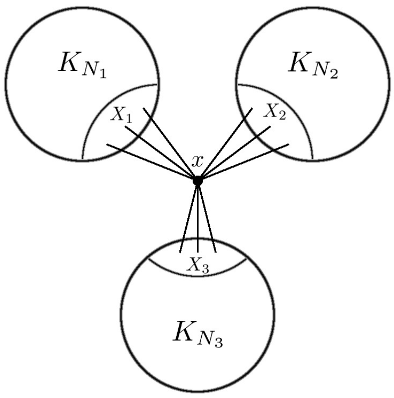

Figure 1.1: A graph with a bad tangle-distinguishing tree-decomposition.

We will show that, for the canonical tree-decompositions of Carmesin et al, the torsos of the inessential parts all have branch-width $< k$ and so it is possible to decompose the torsos of the inessential parts in this way.

Theorem 1.1.2. *For every graph G and $k \geqslant 3$ there exists a canonical tree-decompositon $(T, \mathcal{V})$ of G of adhesion $< k$ such that*

- *$(T, \mathcal{V})$ distinguishes the set of k-tangles in G efficiently;*
- *The torso of every inessential part has branch-width $< k$.*

More recently another potential candidate for these 'k-connected components' has been considered in the literature, called *k-blocks.* We say that a set of at least k vertices in a graph is $(< k)$-*inseparable* if no set of $< k$ vertices can separate any two of the vertices. A *k-block* is a maximal $(< k)$-inseparable set of vertices. These k-blocks differ from subgraphs which are k-connected in the classical sense in that their connectivity is measured in the ambient graph rather than the subgraph itself. For example if we take a large independent set, I, and join each pair of vertices in I by k vertex disjoint paths, then I is a k-block, even though as a subgraph it is independent. Carmesin, Diestel, Hundertmark and Stein [36] showed that, for any graph G, there is a canonical tree-decomposition which distinguishes the set of k-blocks. The work of Carmesin et al [35] extended the results of [36] to more general types of highly connected substructures in graphs, and these results have been extended further by Diestel, Hundertmark and Lemanczyk [48] to more general combinatorial structures, such as matroids.

As before, these tree-decompositions will have some parts which are *essential*, that is they contain a k-block, and some parts which are *inessential*, and it is natural to ask about the structure of these parts. Recently, Diestel, Eberenz and Erde [46] proved a duality theorem for k-blocks, analogous to the tangle/branch-width duality of Robertson and Seymour. The result implies that a graph contains a k-block if and only if it does not admit a tree-decomposition of *block-width* $< k$, where as before, every part in a tree-decomposition of block-width $< k$ is in some sense 'too small' to contain a k-block. We also show a corresponding result for blocks.

Theorem 1.1.3. *For every graph G and $k \geqslant 3$ there exists a canonical tree-decompositon $(T, \mathcal{V})$ of G of adhesion $< k$ such that*

- *$(T, \mathcal{V})$ distinguishes the set of k-blocks in G efficiently;*
- *The torso of every inessential part has block-width $< k$.*

The main result in this paper, of which Theorems 1.1.2 and 1.1.3 are corollaries, is a lemma that gives sufficient conditions on the separators of an inessential part in a distinguishing tree-decomposition for the torso to have small width. These conditions seem quite natural and reasonable, in particular they are satisfied by every part of the canonical tangle/block-distinguishing tree-decompositions constructed by Carmesin et al.

In some sense the canonical tangle-distinguishing tree-decompositions tell us most about the structure of the graph when the essential parts correspond closely to the tangles inside them. For example consider the following two graphs, firstly two K_Ns overlapping in $k-1$ vertices and secondly two $K_{3k/2}$s each with a long path attached, of length $N' = N - 3k/2$, overlapping in a similar way, see Figure 1.2.

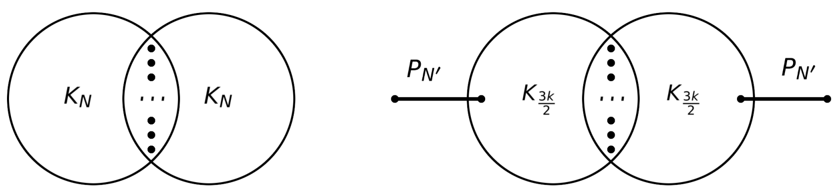

Figure 1.2: Two graphs with the same canonical k-tangle-distinguishing tree-decomposition.

Since the tangle-distinguishing tree-decompositions of Carmesin et al. only use essential separations, that is separations which distinguish some pair of k-tangles, they will construct the same tree-decomposition for both of these graphs, with just two parts of size N. However in the second example a more sensible tree-decomposition would further split up the long paths. This could be done in a way to maintain the property that the inessential parts have small branch-width, and by separating these inessential parts from the essential part we have more precisely exhibited the structure of the graph. We will also apply our methods to the problem of further refining the essential parts of these tree-decompositions.

In Section 1.2 we introduce the background material necessary for our proof and in Section 1.3 we prove our central lemma and deduce the main results in the paper. In Section 1.4 we discuss how our methods can also be used to further refine the essential parts of a tree-decomposition.

1.2 Background material

1.2.1 Separation systems and tree-decompositions

A *separation* of a graph G is a set $\{A, B\}$ of subsets of $V(G)$ such that $A \cup B = V$ and there is no edge of G between $A \setminus B$ and $B \setminus A$. There are two *oriented separations* associated with a separation, (A, B) and (B, A). Informally we think of (A, B) as *pointing towards* B and *away from* A. We can define a partial ordering on the set of oriented separations of G by

$$(A, B) \leqslant (C, D) \text{ if and only if } A \subseteq C \text{ and } B \supseteq D.$$

The *inverse* of an oriented separation (A, B) is the separation (B, A), and we note that mapping every oriented separation to its inverse is an involution which reverses the partial ordering.

In [51] Diestel and Oum generalised these properties of separations of graphs and worked in a more abstract setting. They defined a *separation system* $(\vec{S}, \leqslant, *)$ to be a partially ordered set $\vec{S}$ with an order reversing involution, $*$. The elements of $\vec{S}$ are called *oriented separations.* Often a given element of $\vec{S}$ is denoted by $\vec{s}$, in which case its inverse $\vec{s}^{\,*}$ will be denoted by $\overleftarrow{s}$, and vice versa. Since $*$ is ordering reversing we have that, for all $\vec{r}, \vec{s} \in S$,

$$\vec{r} \leqslant \vec{s} \text{ if and only if } \overleftarrow{r} \geqslant \overleftarrow{s}.$$

A *separation* is a set of the form $\{\vec{s}, \overleftarrow{s}\}$, and will be denoted by simply s. The two elements $\vec{s}$ and $\overleftarrow{s}$ are the *orientations* of s. The set of all such pairs $\{\vec{s}, \overleftarrow{s}\} \subseteq \vec{S}$ will be denoted by S. If $\vec{s} = \overleftarrow{s}$ we say s is *degenerate.* Conversely, given a set $S' \subseteq S$ of separations we write $\vec{S'} := \bigcup S'$ for the set of all orientations of its elements. With the ordering and involution induced from $\vec{S}$, this will form a separation system. When we refer to a oriented separation in a context where the notation explicitly indicates orientation, such as $\vec{s}$ or (A, B), we will usually suppress the prefix “oriented” to improve the flow of the paper.

Given a separation of a graph $\{A, B\}$ we can identify it with the pair $\{(A, B), (B, A)\}$ and in this way any set of separations in a graph which is closed under taking inverses forms a separation system. We will work within the framework developed in [51] since we will need to use directly some results proved in this abstract setting, but also because our results are most easily expressible in this framework. An effort has been made to state the results in the widest generality, so as to be applicable in the broadest sense, however we will always have in mind the motivating example of separation systems which arise as sets of separations in a graph, and so a reader will not lose too much by thinking about these separation systems solely in those terms.

The *separator* of a separation $\vec{s} = (A, B)$ in a graph is the intersection $A \cap B$ and the *order* of a separation, $|\vec{s}| = \mathrm{ord}(A, B)$, is the cardinality of the separator $|A \cap B|$. Note that if $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$ are separations then so are the *corner separations* $\vec{r} \vee \vec{s} := (A \cup C, B \cap D)$ and $\vec{r} \wedge \vec{s} := (A \cap C, , B \cup D)$ and the orders of these separations satisfy the equality

$$|\vec{r} \vee \vec{s}| + |\vec{r} \wedge \vec{s}| = |\vec{r}| + |\vec{s}|.$$

Hence the order function is a submodular function on the set of separations of a graph, and we note also that it is clearly symmetric.

For abstract separations systems, if there exists binary operations $\vee$ and $\wedge$ on $\vec{S}$ such that $\vec{r} \vee \vec{s}$ is the supremum and $\vec{r} \wedge \vec{s}$ is the infimum of $\vec{r}$ and $\vec{s}$ then we call $(\vec{S}, \leqslant, *, \vee, \wedge)$ a

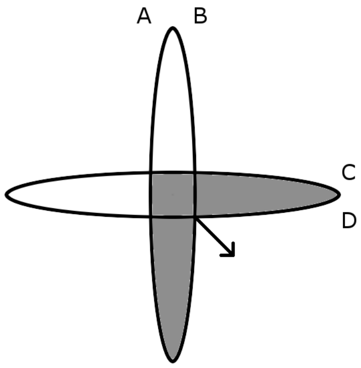

Figure 1.3: Two separations (A, B) and (C, D) with the corner separation $(A \cup C, B \cap D)$ marked.

universe of (oriented) separations, and we call any real, non-negative, symmetric and submodular function on a universe an *order function.*

Two separations r and s are *nested* if they have $\leqslant$-comparable orientations. Two oriented separations $\vec{r}$ and $\vec{s}$ are *nested* if r and s are nested [2]. If $\vec{r}$ and $\vec{s}$ are not nested we say that the two separations *cross.* A set of separations S is *nested* if every pair of separations in S is nested, and a separation s is *nested* with a set of separations S if $S \cup \{s\}$ is nested.

A separation $\vec{r} \in \vec{S}$ is *trivial in* $\vec{S}$, and $\overleftarrow{r}$ is *co-trivial*, if there exist an $s \in S$ such that $\vec{r} < \vec{s}$ and $\vec{r} < \overleftarrow{s}$. Note that if $\vec{r}$ is trivial, witnessed by some s, then, since the involution is order reversing, we have that $\vec{r} < \vec{s} < \overleftarrow{r}$. So, in particular, $\overleftarrow{r}$ cannot also be trivial. Separations $\vec{s}$ such that $\vec{s} \leqslant \overleftarrow{s}$, trivial or not, will be called *small* and their inverses *co-small.*

In the case of separations of a graph, it is a simple check that the small separations are precisely those of the form (A, V). Furthermore the trivial separations can be characterised as those of the form (A, V) such that $A \subseteq C \cap D$ for some separation (C, D) such that $\{C, D\} \neq \{A, B\}$. Finally we note that there is only one degenerate separation in a graph, (V, V).

A *tree-decomposition* of a graph G is a pair $(T, \mathcal{V})$ consisting of a tree T and family $\mathcal{V} = (V_t)_{t \in T}$ of vertex sets $V_t \subseteq V(G)$, one for each vertex $t \in T$ such that:

- $V(G) = \bigcup_{t \in T} V_t$;
- for every edge $e \in G$ there exists some $t \in T$ such that $e \in G[V_t]$;
- $V_{t_1} \cap V_{t_2} \subseteq V_{t_3}$ whenever t_3 lies on the $t_1 - t_2$ path in T.

The sets V_t in a tree-decomposition are its *parts* and the sets $V_t \cap V_{t'}$ such that (t, t') is an edge of T are the *adhesion sets.* The *torso* of a part $\overline{V_t}$ is the union of that part together with the completion of the adhesion sets adjacent to that part, that is

$$\overline{V_t} = G|_{V_t} \cup \bigcup_{(t,t') \in T} K_{V_t \cap V_{t'}}.$$

The *width* of a tree-decomposition is $\max\{|V_t| - 1 : \text{ such that } t \in T\}$, and the *adhesion* is the size of the largest adhesion set. Deleting an oriented edge $e = (t_1, t_2) \in \vec{E}(T)$ divides $T - e$ into two components $T_1 \ni t_1$ and $T_2 \ni t_2$. Then $(\bigcup_{t \in T_1} V_t, \bigcup_{t \in T_2} V_t)$ can be seen to be a separation

[2]In general we will use terms defined for separations informally for oriented separations when the meaning is clear, and vice versa

of G with separator $V_{t_1} \cap V_{t_2}$. We say that the edge e *induces* this separation. Given a tree-decomposition $(T, \mathcal{V})$ it is easy to check that the set of separations induced by the edges of T form a nested separation system. Conversely it was shown in [36] that every nested separation system is induced by some tree-decomposition, and so in a sense these two concepts can be thought of as equivalent.

We say that a nested set of separations $\mathcal{N}'$ *refines* a nested set of separations $\mathcal{N}$ if $\mathcal{N}' \supseteq \mathcal{N}$, and similarly a tree-decomposition $(T', \mathcal{V}')$ *refines* a tree-decomposition $(T, \mathcal{V})$ if the set of separations induced by the edges of T' refines the corresponding set of separations for T.

1.2.2 Duality of tree-decompositions

There are a number of theorems that assert a duality between certain structurally 'large' objects in a graph and an overall tree structure. For example a graph has small tree-width if and only if it contains no large order bramble [118]. In [51] a general theory of duality, in terms of separation systems, was developed which implied many of the existing theorems. Following on from the notion of tangles in graph minor theory [110] these large objects were described as orientations of separations systems avoiding certain forbidden subsets.

An *orientation* of a set of separations S is a subset $O \subseteq \vec{S}$ which for each $s \in S$ contains exactly one of its orientations $\vec{s}$ or $\overleftarrow{s}$. A *partial orientation* of S is an orientation of some subset of S, and we say that an orientation O *extends* a partial orientation P if $P \subseteq O$.

In our context we will think of an orientation O on some set of graph separations as choosing a side of each separation $s = \{A, B\}$ to designate as large. For example given a graph G and the set S of all separations of the graph G, we denote by

$$\vec{S}_k = \{\vec{s} \in \vec{S} \,:\, |\vec{s}| < k\},$$

the set of all orientations of order less than k. If there is a large clique (of size $\geqslant k$) in G then for every $s = \{A, B\} \in S_k$ we have that the clique is contained entirely in A or B. So this clique defines an orientation of S_k by picking, for each $\{A, B\} \in S_k$ the orientated separation such that the clique is contained in second set in the pair.

We call an orientation O of a set of separations S *consistent* if whenever we have distinct r and s such that $\vec{r} < \vec{s}$, O does not contain both $\overleftarrow{r}$ and $\vec{s}$. Note that a consistent orientation must contain all trivial separations $\vec{r}$, since if $\vec{r} < \vec{s}$ and $\vec{r} < \overleftarrow{s}$ then, whichever orientation of s is contained in O would be inconsistent with $\overleftarrow{r}$.

Given a set of subsets $\mathcal{F} \subseteq 2^{\vec{S}}$ we say that an orientation O is $\mathcal{F}$*-avoiding* if there is no $F \in \mathcal{F}$ such that $F \subseteq O$. So for example an orientation is consistent if it avoids $\mathcal{F} = \{\{\overleftarrow{r}, \vec{s}\} \,:\, r \neq s, \vec{r} < \vec{s}\}$. In general we will define the 'large' objects we consider by the collection $\mathcal{F}$ of subsets they avoid. For example a k-tangle in a graph G can easily be seen to be equivalent to an orientation of S_k which avoids the set of triples

$$\mathcal{T}_k = \{\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} \subseteq \vec{S}_k \,:\, \bigcup_{i=1}^{3} G[A_i] = G\}.$$

(Where the three separations need not be distinct). That is, a tangle is an orientation such that no three small sides cover the entire graph, it is a simple check that any such orientation must in fact also be consistent. We say that a consistent orientation which avoids a set $\mathcal{F}$ is an $\mathcal{F}$*-tangle.*

Given a set $\mathcal{F} \subseteq 2^{\vec{S}}$, an *$S$-tree over* $\mathcal{F}$ is a pair (T, α), of a tree T with at least one edge and a function $\alpha : \vec{E}(T) \to \vec{S}$ from the set

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

of orientations of it's edges to $\vec{S}$ such that:

- For each edge $(t_1, t_2) \in \vec{E}(T)$, if $\alpha(t_1, t_2) = \vec{s}$ then $\alpha(t_2, t_1) = \overleftarrow{s}$;
- For each vertex $t \in T$, the set $\{\alpha(t', t) : (t', t) \in \vec{E}(T)\}$ is in $\mathcal{F}$;

For any leaf vertex $w \in T$ which is adjacent to some vertex $u \in T$ we call the separation $\vec{s} = \alpha(w, u)$ a *leaf separation* of (T, α). A particularly interesting class of such trees is when the set $\mathcal{F}$ is chosen to consist of stars. A set of non-degenerate oriented separations σ is called a *star* if $\vec{r} \leqslant \overleftarrow{s}$ for all distinct $\vec{r}, \vec{s} \in \sigma$. In what follows, if we refer to an S-tree without reference to a specific family $\mathcal{F}$ of stars, it can be assumed to be over the set of *all* stars in $2^{\vec{S}}$. We say that an S-tree over $\mathcal{F}$ is *irredundant* if there is no $t \in T$ with two neighbours, t' and t'' such that $\alpha(t, t') = \alpha(t, t'')$. If (T, α) is an irredundant S-tree over a set of stars $\mathcal{F}$, then it is easy to verify that the map α preserves the natural ordering on $\vec{E}(T)$, defined by letting $(s, t) \leqslant (u, v)$ if the unique path in T between those edges starts at s and ends at v (see [[51], Lemma 2.2]).

Given an irredundant S-tree (T, α) over a set of stars and an orientation O of S, O induces an orientation of the edges of T, which will necessarily contain a sink vertex. If the orientation O is consistent then this sink vertex, which we will denote by t, will be unique. We say that O is *contained in* t. If $S = S_k$ for some graph G, we have that (T, α) defines some tree-decomposition $(T, \mathcal{V})$ of G, and we say that O is *contained in the part* V_t. So, each $\mathcal{F}$-tangle of S must live in some vertex of every such S-tree, and by definition this vertex give rise to a star of separations in $\mathcal{F}$. In this way, each of the vertices in an S-tree over $\mathcal{F}$ (and each of the parts in the corresponding tree-decomposition when one exists) is 'too small' to contain an $\mathcal{F}$-tangle.

Suppose we have a separation $\vec{r}$ which is neither trivial nor degenerate. In applications $\vec{r}$ will be a leaf separation in some irredundant S-tree over a set $\mathcal{F}$ of stars. Given some $\vec{s} \geqslant \vec{r}$, it will be useful to have a procedure to 'shift' the S-tree (T, α) in which $\vec{r}$ is a leaf separation to a new S-tree (T, α') such that $\vec{s}$ is a leaf separation. Let $S_{\geqslant \vec{r}}$ be the set of separations $x \in S$ that have an orientation $\vec{x} \geqslant \vec{r}$. Since $\vec{r}$ is a leaf separation in an irredundant S-tree over a set of stars we have by the previous comments that the image of α is contained in $\vec{S}_{\geqslant \vec{r}}$.

Given $x \in S_{\geqslant \vec{r}} \setminus \{r\}$ we have, since $\vec{r}$ is non-trivial, that only one of the two orientations of x, say $\vec{x}$ is such that $\vec{x} \geqslant \vec{r}$. So, we can define a function $f \downarrow^{\vec{r}}_{\vec{s}}$ on $\vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$ by[3]

$$f \downarrow^{\vec{r}}_{\vec{s}} (\vec{x}) := \vec{x} \vee \vec{s} \text{ and } f \downarrow^{\vec{r}}_{\vec{s}} (\overleftarrow{x}) := (\vec{x} \vee \vec{s})^*.$$

Given an S-tree (T, α) and $\vec{s} \geqslant \vec{r}$ as above let $\alpha' := f \downarrow^{\vec{r}}_{\vec{s}} \circ \alpha$. The *shift of* (T, α) *onto* $\vec{s}$ is the S-tree (T, α').

We say that $\vec{s}$ *emulates* $\vec{r}$ *in* $\vec{S}$ if $\vec{r} \leqslant \vec{s}$ and for every $\vec{t} \in \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$, $\vec{s} \vee \vec{t} \in \vec{S}$. Given a particular set of stars $\mathcal{F} \subseteq 2^{\vec{S}}$ we say further that $\vec{s}$ *emulates* $\vec{r}$ *in* $\vec{S}$ *for* $\mathcal{F}$ if $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$ and for any star $\sigma \subset \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$ in $\mathcal{F}$ that contains an element $\vec{t} \geqslant \vec{r}$ we also have $f \downarrow^{\vec{r}}_{\vec{s}} (\sigma) \in \mathcal{F}$. The usefulness of this property is exhibited by the following lemma, which is key both in the proof of Theorem 1.2.2 from [51], and will be essential for the proof of our central lemma.

Lemma 1.2.1. *[[51], Lemma 4.2] Let $(\vec{S}, \leqslant, ^*)$ be a separation system, $\mathcal{F} \subseteq 2^{\vec{S}}$ a set of stars, and let (T, α) be an irredundant S-tree over $\mathcal{F}$. Let $\vec{r}$ be a nontrivial and nondegenerate separation which is a leaf separation of (T, α), and is not the image of any other edge in T, and let $\vec{s}$ emulate $\vec{r}$ in $\vec{S}$. Then the shift of (T, α) onto $\vec{s}$ is an S-tree over $\mathcal{F} \cup \{\{\overleftarrow{s}\}\}$ in which $\vec{s}$ is a leaf separation, associated with a unique leaf.*

[3] The exclusion of $\overleftarrow{r}$ here is for a technical reason, since it could be the case that $\vec{r} < \overleftarrow{r}$, however we want to insist that $f \downarrow^{\vec{r}}_{\vec{s}} (\overleftarrow{r})$ is the inverse of $f \downarrow^{\vec{r}}_{\vec{s}} (\vec{r})$

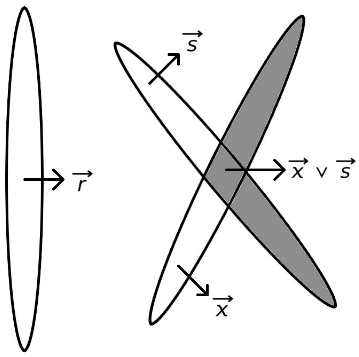

Figure 1.4: Shifting a separation $\vec{x} \geqslant \vec{r}$ under $f \downarrow^{\vec{r}}_{\vec{s}}$.

It is shown in [[51], Lemma 2.4] that if we have an S-tree over $\mathcal{F}$, (T, α), and a set of non-trivial and non-degenerate leaf separations, $\vec{r}_i$, of (T, α) then there also exists an irredundant S-tree over $\mathcal{F}$, (T', α'), such that each $\vec{r}_i$ is a leaf separation of (T', α) and is not the image of any other edge in T'.

We say a set $\mathcal{F} \subseteq 2^{\vec{S}}$ *forces* a separation $\vec{r}$ if $\{\overleftarrow{r}\} \in \mathcal{F}$ or r is degenerate. Note that the non-degenerate forced separations in $\mathcal{F}$ are precisely those separations which can appear as leaf separations in an S-tree over $\mathcal{F}$. We say $\mathcal{F}$ is *standard* if it forces every trivial separation in $\vec{S}$.

We say that a separation system $\vec{S}$ is *separable* if for any two non-trivial and non-degenerate separations $\vec{r}, \overleftarrow{r}' \in \vec{S}$ such that $\vec{r} \leqslant \vec{r}'$ there exists a separation $s \in S$ such that $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$ and $\overleftarrow{s}$ emulates $\overleftarrow{r}'$ in $\vec{S}$. We say that $\vec{S}$ is $\mathcal{F}$*-separable* if for all non-trivial and non-degenerate $\vec{r}, \overleftarrow{r}' \in \vec{S}$ that are not forced by $\mathcal{F}$ such that $\vec{r} \leqslant \vec{r}'$ there exists a separation $s \in S$ such that $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$ and $\overleftarrow{s}$ emulates $\overleftarrow{r}'$ in $\vec{S}$ for $\mathcal{F}$. Often one proves that $\vec{S}$ is $\mathcal{F}$-separable in two steps, first by showing it is separable, and then by showing that $\mathcal{F}$ is *closed under shifting*: that whenever $\vec{s}$ emulates some $\vec{r}$ in $\vec{S}$, it also emulates that $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$.

We are now in a position to state the Strong Duality Theorem from [51].

Theorem 1.2.2. *[[51], Theorem 4.3] Let $(\vec{U}, \leqslant, *, \vee, \wedge)$ be a universe of separations containing a separation system $(\vec{S}, \leqslant, *)$. Let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a standard set of stars. If $\vec{S}$ is $\mathcal{F}$-separable, exactly one of the following assertions holds:*

- *There exists an S-tree over $\mathcal{F}$.*
- *There exists an $\mathcal{F}$-tangle of S.*

The property of being $\mathcal{F}$-separable may seem a rather strong condition to hold, however in [52] it is shown that for all the sets $\mathcal{F}$ describing classical 'large' objects (such as tangles or brambles) the separation systems $\vec{S}_k$ are $\mathcal{F}$-separable. More specifically, by definition a k-tangle is a consistent orientation which avoids the set $\mathcal{T}_k$ as defined earlier. In fact it is shown in [52] that a consistent orientation avoids $\mathcal{T}_k$ if and only if it avoids the set of stars in $\mathcal{T}_k$

$$\mathcal{T}_k^* = \{\{(A_i, B_i)\}_1^3 : \{(A_i, B_i)\}_1^3 \subseteq S_k \text{ is a star and } \bigcup_i G[A_i] = G\}.$$

Note that $\mathcal{T}_k^*$ is standard. Indeed it forces all the small separations (A, V), and so it forces the trivial separations. It can also be checked that $\vec{S}_k$ is $\mathcal{T}_k^*$-separable.

The dual structure to a k-tangle is therefore an S_k-tree over $\mathcal{T}_k^*$. It is shown in [52] that the existence of such an S_k-tree is equivalent to the existence of a branch-decomposition of width $<k$ for all $k \geqslant 3$. We note that the condition that $k \geqslant 3$ is due to a quirk in how branch-width is traditionally defined, which results in, for example, stars having branch-width 1 but all other trees having branch-width 2, whilst both contain 2-tangles.

If a tree-decomposition $(T, \mathcal{V})$ of a graph G is such that the set of separations induced by the edges of T is an S_k-tree over $\mathcal{T}_k^*$ for some k, then there is some smallest such k', and we say the *branch-width* of the tree-decomposition is $k' - 1$. If no such k exists then we will let the branch-width be infinite. By the preceding discussion we have that the branch-width (in the traditional sense) of a graph is the smallest k such that G has a tree-decomposition of branch-width k (except when the branch-width of G is 1), and so this should not cause too much confusion.

1.2.3 Canonical tree-Decompositions distinguishing tangles

Given two orientations O_1 and O_2 of a set of separations S we say that a separation s *distinguishes* O_1 and O_2 if $\vec{s} \in O_1$ and $\overleftarrow{s} \in O_2$. As in the previous section, every tree-decomposition, $(T, \mathcal{V})$, corresponds to some nested set of separations, $\mathcal{N}$. We say that a tree-decomposition *distinguishes* O_1 and O_2 if there is some separation in $\mathcal{N}$ which distinguishes O_1 and O_2. If O_1 and O_2 are consistent, then the tree-decomposition will distinguish them if and only if they are contained in different parts of the tree.

As in Section 1.2.2 a k-block b can be viewed as an orientation of S_k. Indeed given any separation (A, B) with $\mathrm{ord}(A, B) < k$, since b is $(<k)$-inseparable, $b \subseteq A$ or $b \subseteq B$, so we can think of b as orienting each $s \in S_k$ towards the side of the separations that b lies in. In [36] Carmesin, Diestel, Hundertmark and Stein showed how to algorithmically construct a nested set of separations in a graph G (and so a tree-decomposition) in a canonical way, that is, invariant with respect to the automorphism group of G, which distinguishes all of its k-blocks, for a given k.

These ideas were extended in [35] to construct canonical tree-decompositions which distinguish all the k-profiles in a graph, a common generalization of k-tangles and k-blocks. A *k-profile* can be defined as a $\mathcal{P}_k$-tangle of S_k, where

$$\mathcal{P}_k = \{\sigma = \{(A, B), (C, D), (B \cap D, A \cup C)\} \,:\, \sigma \subseteq \vec{S}_k\}.$$

More generally, given a universe of separations $(\vec{U}, \leqslant, *, \vee, \wedge)$ with an order function containing a separation system $(\vec{S}, \leqslant, *)$, we can define as before an *S-profile* to be a $\mathcal{P}_S$-tangle of S where

$$\mathcal{P}_S = \{\sigma = \{\vec{r}, \vec{s}, \overleftarrow{r} \wedge \overleftarrow{s}\} \,:\, \sigma \subseteq \vec{S}\}.$$

Given two distinct S-profiles P_1 and P_2 there is some $s \in S$ which distinguishes them. Furthermore, there is some minimal l such that there is a separation of order l which distinguishes P_1 and P_2, and we define $\kappa(P_1, P_2) := l$. We say that a separation s distinguishes P_1 and P_2 *efficiently* if s distinguishes P_1 and P_2 and $|s| = \kappa(P_1, P_2)$. Given a set of profiles ϕ we say that a separation s is *ϕ-essential* if it efficiently distinguishes some pair of profiles in ϕ. We will often consider in particular, as in the case of graphs, the separation system arising from those separations in a universe of order $<k$, that is we define

$$\vec{S}_k = \{\vec{u} \in \vec{U} \,:\, |\vec{u}| < k\},$$

where in general it should be clear from the context which universe S_k lives in.

In [35] a number of different algorithms, which they call *k-strategies*, are described for constructing a nested set of separations distinguishing a set of profiles. These algorithms build the set of separations in a series of steps, and at each step there is a number of options for how to pick the next set of separations. A k-strategy is then a description of which choice to make at each step. The authors showed that, regardless of which choices are made at each step, this algorithm will produce a nested set of separations distinguishing all the profiles in G. We say a set of profiles is *canonical* if it is fixed under every automorphism of G. In particular the following is shown.

Theorem 1.2.3. *[[35] Theorem 4.4] Every k-strategy Σ determines for every canonical set ϕ of k-profiles of a graph G a canonical nested set $\mathcal{N}_\Sigma(G,\phi)$ of ϕ-essential separations of order $< k$ that distinguishes all the profiles in ϕ efficiently.*

Note that any k-tangle, O, is also a k-profile. Indeed, it is a simple check that O is consistent. Also for any pair of separations $(A,B),(C,D) \in \vec{S}_k$ we have that $G[A] \cup G[C] \cup G[B \cap D] = G$, since any edge not contained in A or C is contained in both B and D. Hence, $\{(A,B),(C,D),(B\cap D, A \cup C)\} \in \mathcal{T}_k$, and so $\mathcal{P}_k \subseteq \mathcal{T}_k$. Therefore any k-tangle, which by definition avoids $\mathcal{T}_k$, must also avoid $\mathcal{P}_k$, and so must be a k-profile. Similarly one can show that the orientations defined by k-blocks are consistent and $\mathcal{P}_k$ avoiding, and so k-profiles. Even more, there is some family $\mathcal{B}_k \supseteq \mathcal{P}_k$ such that the orientations defined by k-blocks are $\mathcal{B}_k$-tangles, and if there is a $\mathcal{B}_k$-tangle of S_k then the graph G contains a unique k-block corresponding to this orientation.

One of the aims of [51, 52] had been to develop a duality theorem which would be applicable to k-profiles and k-blocks. The same authors showed in [50] that there is a more general duality theorem of a similar kind which applies in these cases, however the dual objects in this theorem correspond to a more general object than the classical notion of tree-decompositions.

Nevertheless, it was posed as an open question whether or not there was a duality theorem for k-profiles or k-blocks expressible within the framework of [51]. By Theorem 1.2.2 it would be sufficient to show that there is a standard set of stars $\mathcal{F}$ such that the set of k-profiles or k-blocks coincides with the set of $\mathcal{F}$-tangles. Recently Diestel, Eberenz and Erde [46] showed that, if we insist the orientations satisfy a slightly stronger consistency condition, this will be the case. We say that an orientation O of a separation system S is *regular* if whenever we have r and s such that $\vec{r} \leqslant \vec{s}$, O does not contain both $\overleftarrow{r}$ and $\vec{s}$. We note that a consistent orientation is regular if and only if it contains every small separation. A *regular $\mathcal{F}$-tangle of S* is then a regular $\mathcal{F}$-avoiding orientation of S, and a *regular S-profile* is a regular $\mathcal{P}_S$-tangle. For most natural examples of separation systems there will not be a difference between regular and irregular profiles. Indeed, in [46] it is shown that for $k \geqslant 3$ every k-profile of a graph is in fact a regular k-profile[4].

We say a separation system is *submodular* if whenever $\vec{r}, \vec{s} \in \vec{S}$ either $\vec{r} \wedge \vec{s}$ or $\vec{r} \vee \vec{s} \in \vec{S}$. Note that, if a universe U has an order function, then the separation systems S_k are submodular.

Theorem 1.2.4. *[Diestel, Eberenz and Erde [46]] Let S be a separable submodular separation system contained in some universe of separations $(\vec{U}, \leqslant, *, \vee, \wedge)$, and let $\mathcal{F} \supseteq \mathcal{P}_S$. Then there exists a standard set of stars $\mathcal{F}^*$ (which is closed under shifting, and contains $\{\vec{r}\}$ for every co-small $\vec{r}$) such that every regular $\mathcal{F}$-tangle of S is an $\mathcal{F}^*$-tangle of S, and vice versa, and such that the following are equivalent:*

- *There is no regular $\mathcal{F}$-tangle of S;*
- *There is no $\mathcal{F}^*$-tangle of S;*
- *There is an S-tree over $\mathcal{F}^*$.*

[4] There do exist pathological examples of 2-profiles in graphs which are not regular, however they can be easily characterized.

In the case where $S = S_k$ is the set of separations of a graph with $k \geqslant 3$, we have that $\mathcal{F} \supseteq \mathcal{P}_k$, and so every $\mathcal{F}$-tangle is a k-profile, and so regular. Hence, in this case, we can omit the word regular from the statement of the theorem. We note that $\mathcal{T}_k \supseteq \mathcal{P}_k$, (and in fact the $\mathcal{T}_k^*$ of the theorem can be taken to be the $\mathcal{T}_k^*$ defined earlier) and so Theorem 1.2.4 also implies the tangle/branch-width duality theorem.

Applying the result to $\mathcal{P}_k$ or $\mathcal{B}_k$ also gives a duality theorem for k-blocks and k-profiles. As in the case of tangles, if a tree-decomposition $(T, \mathcal{V})$ of a graph G is such that the set of separations induced by the edges of T is an S_k-tree over $\mathcal{P}_k^*$ for some k, then there is some smallest such k', and we say the *profile-width* of the tree-decomposition is $k' - 1$. If no such k exists then we will let the profile-width be infinite. The *profile-width* of a graph is then the smallest k such that G has a tree-decompositions of profile-width k. Then, as was the case with tangles, Theorem 1.2.4 tells us that the profile-width of a graph is the largest k such that G contains a k-profile. We define the *block-width* of a tree-decomposition and graph in the same way.

In a similar way as before, we can think of any part in a tree-decomposition of block-width at most $k-1$ as being 'too small' to contain a k-block, as the corresponding star of separations must lie in $\mathcal{B}_k^*$, and by Theorem 1.2.4 every k-block defines an orientation of S_k which avoids $\mathcal{B}_k^*$.

1.3 Refining a tree-decomposition

Given a set of profiles of a graph, ϕ, we say a part V_t of a tree-decomposition is *ϕ-essential* if some profile from ϕ is contained in this part. We will keep in mind as a motivating example the case $\phi = \tau_k$, the set of k-tangles and, when the set of profiles considered is clear, we will refer to such parts simply as *essential.* Conversely if no such profile is contained in the part we call it *inessential.* The main result of the paper can now be stated formally.

Lemma 1.3.1. *Let $(\vec{U}, \leqslant, *, \vee, \wedge)$ be a universe of separations with an order function. Let ϕ be a set of S_k-profiles and let $\mathcal{F}$ be a standard set of stars which contains $\{\vec{r}\}$ for every co-small $\vec{r}$, and which is closed under shifting, such that ϕ is the set of $\mathcal{F}$-tangles. Let $\sigma = \{\vec{s}_i : i \in [n]\} \subseteq \vec{S}_k$ be a non-empty star of separations such that each s_i is ϕ-essential, and let $\mathcal{F}' = \mathcal{F} \cup \bigcup_1^n \{\overleftarrow{s_i}\}$*

Then either there is an $\mathcal{F}'$-tangle of S_k, or there is an S_k-tree over $\mathcal{F}'$ in which each $\vec{s}_i$ appears as a leaf separation.

If we compare Lemma 1.3.1 to Theorem 1.2.2, we see that Lemma 1.3.1 can be viewed in some way as a method of building a new duality theorem from an old one, by adding some singleton separations to our set $\mathcal{F}$. The restriction to considering only S_k-profiles rather than those of arbitrary separation systems S contained in U comes from the proof, where we need to use the submodularity of the order function to show that certain separations emulate others. It would be interesting to know if the result would still be true for any S which is separable, or even any pair $\mathcal{F}$ and S such that S is $\mathcal{F}$-separable. The condition that $\mathcal{F}$ contains every co-small separation as a singleton is to ensure that the $\mathcal{F}$-tangles are regular $\mathcal{F}$-tangles, as we will need to use the slightly stronger consistency condition in the proof.

What does Lemma 1.3.1 say in the case of k-tangles arising from graphs? Recall that τ_k is the set of $\mathcal{T}_k^*$-tangles, and that $\mathcal{T}_k^*$ is closed under shifting, and contains $\{\vec{r}\}$ for every co-small $\vec{r}$. Given a star $\sigma = \{\vec{s}_i : i \in [n]\} \subseteq \vec{S}_k$ we note that a $\mathcal{T}_k^* \cup \bigcup_1^n \{\overleftarrow{s_i}\}$-tangle is just a $\mathcal{T}_k^*$-tangle which contains $\vec{s}_i$ for each i, and so it is a k-tangle which orients the star inwards. Conversely, an S_k-tree over $\mathcal{T}_k^* \cup \bigcup_1^n \{\overleftarrow{s_i}\}$ in which each $\vec{s}_i$ appears as a leaf separation will give a tree-decomposition of the part of the graph at σ. In particular, since each of the separations in the tree will be nested with σ, the separators $A_i \cap B_i$ of the separations $\vec{s}_i$ will lie entirely on one side of every separation in the tree, and so this will in fact be a decomposition of the torso of the part (since any extra edges in the torso lie inside the separators).

Therefore, in practice this tells us that if we have a part in a tree-decomposition whose separators are τ_k-essential then either there is a k-tangle in the graph which is contained in that part, or there is a tree-decomposition of the torso of that part with branch-width $< k$. In the second case we can then refine the original tree-decomposition by combining it with this new tree-decomposition. By applying this to each inessential part of one of the canonical tree-decompositions formed in [35] we get the following result, which easily implies Theorems 1.1.2 and 1.1.3 by taking $\mathcal{F} = \mathcal{T}_k^*$ and $\mathcal{B}_k^*$ respectively.

Corollary 1.3.2. *Let $k \geqslant 3$ and let $\mathcal{F} \supseteq \mathcal{P}_k$ be such that the set ϕ of regular $\mathcal{F}$-tangles is canonical. If $\mathcal{F}^*$ is defined as in Theorem 1.2.4 then there exists a nested set of separations $\mathcal{N} \subseteq S_k$ corresponding to an S_k-tree (T, α) of G such that:*

- *there is a subset $\mathcal{N}' \subseteq \mathcal{N}$ that is fixed under every automorphism of G and distinguishes all the regular $\mathcal{F}$-tangles in ϕ efficiently;*
- *every vertex $t \in T$ either contains a regular $\mathcal{F}$-tangle or $\{\alpha(t', t) : (t', t) \in \vec{E}T)\} \in \mathcal{F}^*$.*

Proof. By Theorem 1.2.3 there exists a canonical nested set $\mathcal{N}'$ of ϕ-essential separations of order $< k$ that distinguishes all the regular $\mathcal{F}$-tangles in ϕ efficiently, and by Theorem 1.2.4 ϕ is also the set of $\mathcal{F}^*$-tangles. Given an inessential part V_t in the corresponding tree-decomposition $(T, \mathcal{V})$, this part corresponds to some star of separations $\sigma = \{\vec{s}_i : i \in [n]\} \subseteq \mathcal{N}'$. Each $\vec{s}_i \in \mathcal{N}'$ is ϕ-essential, and, by Theorem 1.2.4, $\mathcal{F}^*$ is a standard set of stars which is closed under shifting, and contains $\{\vec{r}\}$ for every co-small $\vec{r}$. Hence, by Lemma 1.3.1, if we let $\mathcal{F}' = \mathcal{F}^* \cup \bigcup_1^n \{\vec{s}_i\}$, there is either an $\mathcal{F}'$-tangle of S_k, or an S_k-tree over $\mathcal{F}'$ in which each $\vec{s}_i$ appears as a leaf separation.

Suppose that there exists an $\mathcal{F}'$-tangle O. Since O avoids $\mathcal{F}' \supseteq \mathcal{F}^*$, it is also an $\mathcal{F}^*$-tangle, and so $O \in \phi$. By assumption $\mathcal{N}'$ distinguishes all the regular $\mathcal{F}$-tangles in ϕ, so O is contained in some part of the tree-decomposition, and since O avoids $\{\{\overleftarrow{s}_i\} : i \in [n]\}$, it must extend σ, and so this part must be V_t. However, this contradicts the assumption that V_t is inessential.

Therefore, by Lemma 1.3.1, there exists an S_k-tree over $\mathcal{F}^* \cup \bigcup_1^n \{\overleftarrow{s}_i\}$. This gives a nested set of separations $\mathcal{N}_t$ which contains the set σ. If we take such a set for each inessential V_t then the set

$$\mathcal{N} = \mathcal{N}' \cup \bigcup_{V_t \text{ inessential}} \mathcal{N}_t$$

satisfies the conditions of the corollary. □

We note that, whilst the existence of such a tree-decomposition is interesting in its own right, perhaps a more useful application of Lemma 1.3.1 is that we can conclude the same for *every* tree-decomposition constructed by the algorithms in [35]. So, we are able to choose whichever algorithm we want to construct our initial tree-decomposition, perhaps in order to have some control over the structure of the essential parts, and we can still conclude that the inessential parts have small branch-width.

Apart from the set τ_k of k-tangles there is another natural set of tangles for which tangle-distinguishing tree-decompositions have been considered. Since a k-tangle, as a $\mathcal{T}_k$-avoiding orientation of S_k, induces an orientation on S_i for all $i \leqslant k$, it induces an i-tangle for all $i \leqslant k$. If an i-tangle for some i is not induced by any k-tangle with $k > i$ we say it is a *maximal tangle.*

Robertson and Seymour [110] showed that there is a decomposition of the graph which distinguishes its maximal tangles, but the theorem does not tell us much about the structure of this tree-decomposition. The approach of Carmesin et al was extended by Diestel, Hundertmark and Lemanczyk [48] to show how an iterative approach to Theorem 1.2.3 could be used to build canonical tree-decompositions distinguishing the maximal tangles in a graph (in fact they showed a stronger result for a broader class of profiles which implies the result for tangles). In particular, the results of [48] imply the following.

Theorem 1.3.3. *If ϕ is a canonical set of tangles in a graph G, then there exists a canonical nested set $\mathcal{N}(G,\phi)$ of ϕ-essential separations that distinguishes all the tangles in ϕ efficiently.*

In particular we can apply this to the set of maximal tangles. By looking directly at the proof in [48] one can see the structure of the tree-decomposition formed. The proof proceeds iteratively, by choosing for each i in a turn a nested set of $(i-1)$-separations (that is, separations of order $(i-1)$), which distinguishes efficiently the pairs of i-tangles which are distinguished efficiently by an $(i-1)$-separation, such that this set is also nested with the previously constructed sets.

At each stage in the construction we have a tree-decomposition which distinguishes all the tangles of order $\leqslant i$ in the graph. Some of these i-tangles however will extend to $(i+1)$-tangles in different ways (induced by distinct maximal tangles in the graph). The next stage constructs a nested set of separations distinguishing such tangles, which gives a tree-decomposition of the torsos of the relevant parts. In these tree-decompositions some parts will be 'essential', and containing $(i+1)$-tangles, but some will be inessential.

It is natural to expect that the inessential parts constructed at stage i should have branch-width $< i$, by a similar argument as Corollary 1.3.2. However it is not always the case that the separators of the inessential part satisfy the conditions of Lemma 1.3.1, since it can be the case that these inessential parts have separators which are separations constructed in an earlier stage of the process, and as such might not efficiently distinguish a pair of tangles of order i.

Question 1.3.4. *Can we bound the branch-width of the inessential parts in such a tree-decomposition in a similar way?*

A positive answer to the previous question in the strongest form would give the following analogue of Theorem 1.1.2.

Conjecture 1.3.5. *For every graph G there exists a canonical sequence of tree-decompositions $(T_i,\mathcal{V}_i)$ for $1 \leqslant i \leqslant n$ of G such that*

- *$(T_i,\mathcal{V}_i)$ distinguishes every i-tangle in G for each i;*
- *$(T_n,\mathcal{V}_n)$ distinguishes the set of maximal tangles in G.*
- *$(T_{i+1},\mathcal{V}_{i+1})$ refines $(T_i,\mathcal{V}_i)$ for each i;*
- *The torso of every inessential part in $(T_i,\mathcal{V}_i)$ has branch-width $< i$.*

1.3.1 Proof of Lemma 1.3.1

Proof of Lemma 1.3.1. Let us write

$$\overline{\mathcal{F}} = \mathcal{F} \cup \{\{\overleftarrow{x}\} \,:\, \overleftarrow{s}_i \leqslant \overleftarrow{x} \text{ for some } i \in [n]\}.$$

We first claim that $\vec{S}_k$ is $\overline{\mathcal{F}}$-separable. We note that by [[52], Lemma 3.4] for every universe $\vec{U}$ and any $k \in \mathbb{N}$, the separation system $\vec{S}_k$ is separable. Therefore it is sufficient to show that $\overline{\mathcal{F}}$ is closed under shifting. By assumption $\mathcal{F}$ is closed under shifting, and the image of any singleton star $\{\overleftarrow{x}\} \in \overline{\mathcal{F}}$ under some relevant $f \downarrow^{\vec{r}}_{\vec{s}}$ is $\{\overleftarrow{y}\}$ for some separation $\overleftarrow{x} \leqslant \overleftarrow{y}$, and hence $\{\overleftarrow{y}\} \in \overline{\mathcal{F}}$. Therefore, $\overline{\mathcal{F}}$ is closed under shifting. Furthermore, since $\mathcal{F}$ was standard, so is $\overline{\mathcal{F}}$. Hence, we can apply Theorem 1.2.2 to $\overline{\mathcal{F}}$.

By Theorem 1.2.2, either there exists an S_k-tree over $\overline{\mathcal{F}}$, or there exists an $\overline{\mathcal{F}}$-tangle. Since $\overline{\mathcal{F}} \supset \mathcal{F}'$, every $\overline{\mathcal{F}}$-tangle is also an $\mathcal{F}'$-tangle, and so in the second case we are done. Therefore we may assume that there exists an S_k-tree over $\overline{\mathcal{F}}$, (T, α). We will use (T, α) to form an S_k-tree over $\mathcal{F}'$.

Since there is no $\overline{\mathcal{F}}$-tangle, each $\mathcal{F}$-tangle O must contain some $\overleftarrow{s}_i$. We note that, since by assumption $\mathcal{F}$ contains every co-small separation, O is regular. Hence, since σ is a star, this $\overleftarrow{s}_i$ is unique. We claim that, for every $\mathcal{F}$-tangle O such that $\overleftarrow{s}_i \in O$ there is some leaf separation $\vec{x} \in \alpha(\overrightarrow{E}(T))$ such that $\overleftarrow{x} \leqslant \overleftarrow{s}_i$.

Indeed, since O is a consistent orientation of $\overrightarrow{S}_k$, it is contained in some vertex of $(T\alpha)$. However, the star of separations at that vertex, by definition of an $\mathcal{F}$-tangle, cannot lie in $\mathcal{F}$, and so must lie in $\overline{\mathcal{F}} \setminus \mathcal{F}$. Since each of these stars are singletons, the vertex must be a leaf. Therefore, there is some leaf separation $\vec{x}$ such that $\overleftarrow{x} \in O$. Since $\{\overleftarrow{x}\} \in \overline{\mathcal{F}} \setminus \mathcal{F}$, it follows that $\overleftarrow{s}_r \leqslant \overleftarrow{x}$ for some $r \in [n]$. However, since $\overleftarrow{s}_i \in O$, and it was the unique separation in σ with that property, it follows that $r = i$, and so $\overleftarrow{s}_i \leqslant \overleftarrow{x}_i$ as claimed.

If the only leaf separations in $\overline{\mathcal{F}} \setminus \mathcal{F}$ were the separations $\{\vec{s}_i : i \in [n]\}$ then (T, α) would be the required S-tree over $\mathcal{F}'$. In general however the tree will have a more arbitrary set $\{\vec{x}_{i,j}\}$ of leaf separations (along with some leaf separations arising as separations forced by $\mathcal{F}$) where $\overleftarrow{s}_i \leqslant \overleftarrow{x}_{i,j}$, see Figure 1.5. Note that there may not necessarily be any edges in this tree corresponding to the separations s_i.

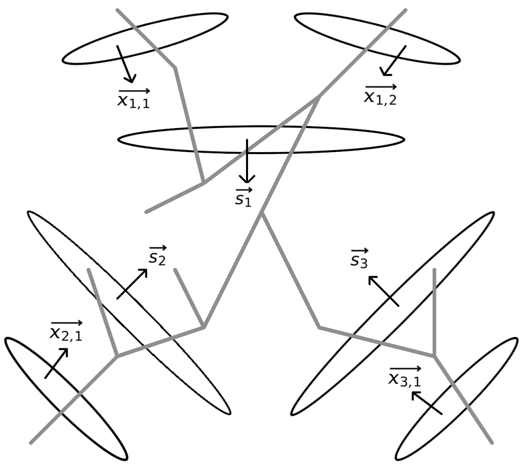

Figure 1.5: The S_k-tree over $\overline{\mathcal{F}}$ with unlabelled leafs corresponding to separations forced by $\mathcal{F}$.

We claim that each $\vec{s}_i$ emulates some $\vec{x}_{i,j}$ in $\overrightarrow{S}_k$ for $\overline{\mathcal{F}}$. By assumption, every $s_i \in \sigma$ distinguishes efficiently some pair O_1 and O_2 of $\mathcal{F}$-tangles. Suppose that $\vec{s}_i \in O_1$ and $\overleftarrow{s}_i \in O_2$. By our previous claim, there is some leaf separation $\vec{x}_{i,j}$ such that $\overleftarrow{x}_{i,j} \in O_2$. We claim that $\vec{s}_i$ emulates this $\vec{x}_{i,j}$ in $\overrightarrow{S}_k$. Note that, since $\overline{\mathcal{F}}$ is closed under shifting, it would follow that $\vec{s}_i$ emulates $\vec{x}_{i,j}$ in $\overrightarrow{S}_k$ for $\overline{\mathcal{F}}$. Note that, since $\vec{s}_i$ and $\vec{x}_{i,j}$ both distinguish two $\mathcal{F}$-tangles, they are non-trivial and non-degenerate.

Indeed, given any separation $\vec{r} \geqslant \vec{x}_{i,j}$ we have that $\vec{s}_i \geqslant \vec{s}_i \wedge \vec{r} \geqslant \vec{x}_{i,j}$ and so $\vec{s}_i \wedge \vec{r}$ distinguishes O_1 and O_2. Therefore, since s_i distinguises O_1 and O_2 efficiently, $|\vec{s}_i \wedge \vec{r}| \geqslant |\vec{s}_i|$. Hence, by submodularity, $|\vec{s}_i \vee \vec{r}| \leqslant |\vec{r}| < k$ and so $\vec{s}_i \vee \vec{r} \in S_k$. Therefore the image of $f \downarrow^{\vec{x}_{i,j}}_{\vec{s}_i}$ is contained in S_k and so $\vec{s}_i$ emulates $\vec{x}_{i,j}$ in $\overrightarrow{S}_k$. Furthermore, since $\vec{x}_{i,j}$ is non-trivial and non-degenerate, by the comment after Lemma 1.2.1 we can assume that T is irredundant, and that $\overrightarrow{x}_{i,j}$ is not the image of any other edge in T.

Since $\overrightarrow{s}_n$ and $\overrightarrow{x}_{n,j}$ satisfy the conditions of Lemma 1.2.1, we conclude that the shift of (T, α) onto $\overrightarrow{s}_n$ is an S_k-tree over $\overline{\mathcal{F}}$ which contains $\overrightarrow{s}_n$ as a leaf separation, and not as the image of any other edge. Let us write (T_n, α_n) for this S_k-tree.

If there is some leaf separation $\overleftarrow{r}$ of (T_n, α_n) such that $\overleftarrow{s}_n < \overleftarrow{r}$ then, since the leaf separations form a star and $\overleftarrow{s}_i$ is the image of a unique leaf, we also have that $\overrightarrow{s}_i < \overleftarrow{r}$. Hence, $\overrightarrow{r}$ is trivial, and so $\{\overleftarrow{r}\} \in \mathcal{F}$. Therefore (T_n, α_n) is also an S_k-tree over

$$\overline{\mathcal{F}}_n = \mathcal{F} \cup \{\overleftarrow{s}_n\} \cup \{\{\overleftarrow{x}\} : \overleftarrow{s}_i \leqslant \overleftarrow{x} \text{ for some } i \in [n-1]\}.$$

If we repeat this argument for each $1 \leqslant i \leqslant n$, we end up with a sequence of S_k-trees (T_n, α_n),
$(T_{n-1}, \alpha_{n-1}), \ldots (T_1, \alpha_1)$ over $\overline{\mathcal{F}}$ such that (T_j, α_j) is also an S_k-tree over

$$\overline{\mathcal{F}}_j = \mathcal{F} \cup \{\overleftarrow{s}_i : i \geqslant j\} \cup \{\{\overleftarrow{x}\} : \overleftarrow{s}_i \leqslant \overleftarrow{x} \text{ for some } i \in [j-1]\}.$$

We note that $\overline{\mathcal{F}}_1 = \mathcal{F}'$, and so (T_1, α_1) is an S_k-tree over $\mathcal{F}'$, completing the proof. □

1.4 Further refining essential parts of tangle-distinguishing tree-decompositions

In some sense the tree-decompositions of Corollary 1.3.2 tell us most about the structure of the graph when the essential parts correspond closely to the profiles inside them. However, as the example in Figure 1.2 shows, sometimes there can be essential parts which could be further refined, in order to more precisely exhibit the structure of the graph.

In this section we will discuss how the tools from the paper can be used to achieve this goal. Given a graph G we call a separation $\overleftarrow{x} \in \overrightarrow{S}_k$ *inessential* if $\overleftarrow{x} \in O$ for every k-tangle O of G. Given a k-tangle O let $\mathcal{M}(O)$ be the set of maximal separations in O, and let $\mathcal{M}_I(O)$ be the set of maximal inessential separations. Our main tool will be the following lemma.

Lemma 1.4.1. *Let G be a graph, O be a k-tangle of G and let $\overleftarrow{x} \in \mathcal{M}_I(O)$ be non-trivial. Then there is an S_k-tree over $\mathcal{T}_k^* \cup \{\overleftarrow{x}\}$.*

Proof. As in the proof of Lemma 1.3.1 let us consider the family of stars

$$\mathcal{F} = \mathcal{T}_k^* \cup \{\{\overleftarrow{r}\} : \overleftarrow{x} \leqslant \overleftarrow{r}\}.$$

A similar argument show that this family is standard and closed under shifting, and so Theorem 1.2.2 asserts the existence of a $\mathcal{F}$-tangle, or an S_k-tree over $\mathcal{F}$. As before, an $\mathcal{F}$-tangle would be a k-tangle of G which contains $\overrightarrow{x}$, contradicting the fact that $\overleftarrow{x}$ is inessential. Therefore, there is an S_k-tree over $\mathcal{F}$. However, O must live in some part of this tree-decomposition, and since O is $\mathcal{T}_k^*$-avoiding it must live in some leaf vertex, corresponding to a singleton star $\{\overleftarrow{r}\}$ for some $\overleftarrow{x} \leqslant \overleftarrow{r}$. However, $\overleftarrow{x}$ was a maximal separation in O and hence $\overleftarrow{r} \notin O$ unless $\overleftarrow{r} = \overleftarrow{x}$. Therefore the S_k-tree is in fact over $\mathcal{T}_k^* \cup \{\overleftarrow{x}\}$. □

Lemma 1.4.1 tell us that for every $\overleftarrow{x} \in \mathcal{M}_I(O)$ there is a tree-decomposition of the part of the graph behind $\overleftarrow{x}$ with branch-width $< k$. So, we could perhaps hope to refine our canonical k-tangle-distinguishing tree-decompositions further using these tree-decompositions. However, there is no guarantee that $\mathcal{M}_I(O)$ will be nested with the τ_k-essential separations used in a k-tangle-distinguishing tree-decomposition, and so we cannot in general refine such tree-decompositions naively in this way. Moreso, in order to decompose as much of the inessential parts of the graph as possible we would like to take such a tree for *each* such maximal separation, however again in general, $\mathcal{M}_I(O)$ itself may not be nested.

Our plan will be to find, for each $\overleftarrow{x} \in \mathcal{M}_I(O)$, some inessential separation $\overleftarrow{u}$ such that $\overrightarrow{u}$ emulates $\overrightarrow{x}$ in $\overrightarrow{S}_k$, that is also nested with the separations from the k-tangle-distinguishing tree-decomposition. Furthermore we would like to able to do this in such a way that the separations $\overrightarrow{u}$ emulating different maximal separations form a star. Then, for each maximal separation, we could shift the S_k-tree given by Lemma 1.4.1 to an S_k-tree over $\mathcal{T}_k^* \cup \{\overleftarrow{u}\}$. These tree-decompositions could then be used to refine our k-tangle-distinguishing tree-decomposition further. We will in fact show a more general result that may be of interest in its own right.

1.4.1 Uncrossing sets of separations

Given two separations $\overrightarrow{r} \leqslant \overrightarrow{s}$ in an arbitrary universe with an order function, we say that $\overrightarrow{s}$ is *linked* to $\overrightarrow{r}$ if for every $\overrightarrow{x} \geqslant \overrightarrow{r}$ we have that

$$|\overrightarrow{x} \vee \overrightarrow{s}| \leqslant |\overrightarrow{x}|.$$

In particular we note that if $\overrightarrow{r}, \overrightarrow{s} \in \overrightarrow{S}_k$, then $\overrightarrow{s}$ being linked to $\overrightarrow{r}$ implies that $\overrightarrow{s}$ emulates $\overrightarrow{r}$ in $\overrightarrow{S}_k$. We first note explicitly a fact used in the proof of Lemma 1.3.1.

Lemma 1.4.2. *Let $(\overrightarrow{U}, \leqslant, *, \vee, \wedge)$ be a universe of separations with an order function, and let $\overleftarrow{s} \leqslant \overleftarrow{r}$ be two separations in $\overrightarrow{U}$. If $\overleftarrow{x}$ is a separation of minimal order such that $\overleftarrow{s} \leqslant \overleftarrow{x} \leqslant \overleftarrow{r}$, then $\overrightarrow{x}$ is linked to $\overrightarrow{r}$.*

Proof. Given any separation $\overrightarrow{y} > \overrightarrow{r}$ we note that

$$\overleftarrow{s} \leqslant \overleftarrow{x} \leqslant \overleftarrow{x} \vee \overleftarrow{y} \leqslant \overleftarrow{r},$$

and so by minimality of $\overleftarrow{x}$ we have that $|\overleftarrow{y} \vee \overleftarrow{x}| \geqslant |\overleftarrow{x}|$. Hence, by submodularity $|\overrightarrow{y} \vee \overrightarrow{x}| = |\overleftarrow{y} \wedge \overleftarrow{x}| \leqslant |\overrightarrow{y}|$, and so $\overrightarrow{x}$ is linked to $\overrightarrow{r}$. Note that, by symmetry, $\overleftarrow{x}$ is linked to $\overleftarrow{s}$ also. □

In what follows we will need to use two facts about a universe of separations. The first is true for any universe of separations, that for any two separations $\overleftarrow{x}$ and $\overleftarrow{y}$

$$(\overleftarrow{x} \wedge \overleftarrow{y})^* = \overrightarrow{x} \vee \overrightarrow{y} \text{ and } (\overleftarrow{x} \vee \overleftarrow{y})^* = \overrightarrow{x} \wedge \overrightarrow{y}.$$

The second will not be true in general, and so we say a universe of separations is *distributive* if for every three separations $\overleftarrow{x}, \overleftarrow{y}$ and $\overleftarrow{z}$ it is true that

$$(\overleftarrow{x} \wedge \overleftarrow{y}) \vee \overleftarrow{z} = (\overleftarrow{x} \vee \overleftarrow{z}) \wedge (\overleftarrow{y} \vee \overleftarrow{z}) \text{ and } (\overleftarrow{x} \vee \overleftarrow{y}) \wedge \overleftarrow{z} = (\overleftarrow{x} \wedge \overleftarrow{z}) \vee (\overleftarrow{y} \wedge \overleftarrow{z}).$$

It is a simple check that the universe of separations of a graph is distributive.

Lemma 1.4.3. *Let $(\overrightarrow{U}, \leqslant, *, \vee, \wedge)$ be a distributive universe of separations with an order function, and let $\overleftarrow{x}_1$ and $\overleftarrow{x}_2$ be two separations in $\overrightarrow{U}$. Let $\overleftarrow{u}_1$ be any separation of minimal order such that $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leqslant \overleftarrow{u}_1 \leqslant \overleftarrow{x}_1$ and let $\overleftarrow{u}_2 = \overleftarrow{x}_2 \wedge \overrightarrow{u}_1$. Then the following statements hold:*

- *$\overrightarrow{u}_1$ is linked to $\overrightarrow{x}_1$ and $\overrightarrow{u}_2$ is linked to $\overrightarrow{x}_2$;*
- *$|\overleftarrow{u}_1| \leqslant |\overleftarrow{x}_1|$ and $|\overleftarrow{u}_2| \leqslant |\overleftarrow{x}_2|$;*
- *$\overleftarrow{u}_1 = \overleftarrow{x}_1 \wedge \overrightarrow{u}_2$.*

Proof. We note that $\vec{u}_1$ is linked to $\vec{x}_1$ by Lemma 1.4.2. We want to show that $\vec{u}_2$ is linked to $\vec{x}_2$, that is, given any $\vec{r} > \vec{x}_2$ we need that $|\vec{r} \vee \vec{u}_2| \leqslant |\vec{r}|$. We first claim that $\vec{r} \vee \vec{u}_2 = \overleftarrow{u}_1 \vee \vec{r}$. Indeed,

$$\vec{r} \vee \vec{u}_2 = \vec{r} \vee (\vec{x}_2 \vee \overleftarrow{u}_1) = (\vec{r} \vee \vec{x}_2) \vee \overleftarrow{u}_1 = \vec{r} \vee \overleftarrow{u}_1.$$

We also claim that $\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{u}_1 \wedge \vec{r} \leqslant \overleftarrow{x}_1$. Indeed, $\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{u}_1$ and $\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \vec{x}_2 \leqslant \vec{r}$ and so

$$\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{u}_1 \wedge \vec{r} \leqslant \overleftarrow{u}_1 \leqslant \overleftarrow{x}_1.$$

Therefore, by minimality of $\overleftarrow{u}_1$ we have that $|\overleftarrow{u}_1 \wedge \vec{r}| \geqslant |\overleftarrow{u}_1|$ and so, by submodularity, it follows that

$$|\vec{r} \vee \vec{u}_2| = |\overleftarrow{u}_1 \vee \vec{r}| \leqslant |\vec{r}|,$$

as claimed.

By minimality of $\overleftarrow{u}_1$ we have that $|\overleftarrow{u}_1| \leqslant |\overleftarrow{x}_1|$. Also we note that, since $\overleftarrow{u}_2 = \overleftarrow{x}_2 \wedge \vec{u}_1$ we have that

$$|\overleftarrow{u}_2| + |\overleftarrow{x}_2 \vee \vec{u}_1| \leqslant |\overleftarrow{x}_2| + |\overleftarrow{u}_1|.$$

However, $|\overleftarrow{x}_2 \vee \vec{u}_1| = |\overleftarrow{u}_1 \wedge \vec{x}_2|$, and we claim that

$$\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{u}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{x}_1.$$

Indeed, that second inequality is clear since, $\overleftarrow{u}_1 \leqslant \overleftarrow{x}_1$. For the first we note that $\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \vec{x}_2$, and also $\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{x}_1 \wedge \vec{u}_2 = \overleftarrow{u}_1$, and so $\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{u}_1 \wedge \vec{x}_2$. Hence, by the minimality of $\overleftarrow{u}_1$, we have $|\overleftarrow{x}_2 \vee \vec{u}_1| \geqslant |\overleftarrow{u}_1|$. Hence it follows that $|\overleftarrow{u}_2| \leqslant |\overleftarrow{x}_2|$, as claimed.

For the last condition, we have that $\overleftarrow{u}_1 \leqslant \overleftarrow{x}_1$ and $\overleftarrow{u}_1 \leqslant \overleftarrow{u}_1 \vee \vec{x}_2 = \vec{u}_2$, and so $\overleftarrow{u}_1 \leqslant \overleftarrow{x}_1 \wedge \vec{u}_2$. However,

$$\overleftarrow{x}_1 \wedge \vec{u}_2 = \overleftarrow{x}_1 \wedge (\vec{x}_2 \vee \overleftarrow{u}_1) = (\overleftarrow{x}_1 \wedge \vec{x}_2) \vee (\overleftarrow{x}_1 \wedge \overleftarrow{u}_1) = (\overleftarrow{x}_1 \wedge \vec{x}_2) \vee \overleftarrow{u}_1 \leqslant \overleftarrow{u}_1,$$

and so $\overleftarrow{u}_1 = \overleftarrow{x}_1 \wedge \vec{u}_2$. □

We note that if we apply the above lemma to a pair of separations $\overleftarrow{x}_1$ and $\overleftarrow{x}_2$ such that x_1 distinguishes efficiently a pair of regular k-profiles, which x_2 does not distinguish, say $\overleftarrow{x}_1 \in P_1$ and $\vec{x}_1 \in P_2$ and $\overleftarrow{x}_2 \in P_1 \cap P_2$, then $\overleftarrow{x}_1$ is of minimal order over all separations $\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{u}_1 \leqslant \overleftarrow{x}_1$. Hence, in Lemma 1.4.3, we can take $\overleftarrow{u}_1 = \overleftarrow{x}_1$ and $\overleftarrow{u}_2 = \overleftarrow{x}_2 \wedge \vec{x}_1$.

Indeed, suppose $\overleftarrow{x}_1 \wedge \vec{x}_2 \leqslant \overleftarrow{u}_1 \leqslant \overleftarrow{x}_1$ is of minimal order. We note that $\overleftarrow{u}_1 \in P_1$ by regularity. Similarly, let $\overleftarrow{u}_2 = \overleftarrow{x}_2 \wedge \vec{u}_1$ then $\overleftarrow{u}_2 \in P_2$ by regularity. Recall that, by Lemma 1.4.3 $\overleftarrow{u}_1 = \overleftarrow{x}_1 \wedge \vec{u}_2$. Hence, $\vec{u}_1 = \vec{x}_1 \vee \overleftarrow{u}_2 \in P_2$. Therefore, u_1 distinguishes P_1 and P_2 and so, by the efficiency of x_1, $|\overleftarrow{x}_1| \leqslant |\overleftarrow{u}_1|$ as claimed.

The question remains as to what happens for a larger set of separations. It would be tempting to conjecture that the following extension of Lemma 1.4.3 holds, where we note that, in general, $(\vec{x} \wedge \vec{y}) \wedge \vec{z} = \vec{x} \wedge (\vec{y} \wedge \vec{z})$ and so, when writing such an expression we can, without confusion, omit the brackets.

Conjecture 1.4.4. *Let $(\vec{U}, \leqslant, *, \vee, \wedge)$ be a distributive universe of separations with an order function, and let $\{\overleftarrow{x}_i : i \in [n]\}$ be a set of separations in $\vec{U}$. Then there exists a set of separations $\{\overleftarrow{u}_i : i \in [n]\}$ such that the following conditions hold:*

- *$\{\overleftarrow{u}_i : i \in [n]\}$ is a star;*
- *$\vec{u}_i$ is linked to $\vec{x}_i$ for all $i \in [n]$;*
- *$|\overleftarrow{u}_i| \leqslant |\overleftarrow{x}_i|$ for all $i \in [n]$;*

- $\overleftarrow{u}_i = \overleftarrow{x}_i \bigwedge_{j \neq i} \overrightarrow{u}_j$ *for all* $i \in [n]$.

However, it seems difficult to ensure that the fourth condition holds with an inductive argument. We were able to show the following in the case of graph separations, by repeatedly applying Lemma 1.4.3. The extra sets $\{\overleftarrow{r}_i\}$ and ϕ appearing in the statement will be useful for the specific application we have in mind, the conclusion when these are empty is the weakened form of the above conjecture.

Lemma 1.4.5. *Let G be a graph, $k \geqslant 3$, and let ϕ be the set of k-profiles in G. Suppose that $\{\overleftarrow{r}_i = (A_i, B_i) : i \in [n]\}$ is a star composed of ϕ-essential separations, which distinguish efficiently some set ϕ' of regular k-profiles and let $\{\overleftarrow{x}_j = (X_j, Y_j) : j \in [m]\} \subseteq \overrightarrow{S}_k$ be such that $\overleftarrow{x}_j \in P$ for all $j \in [m]$ and $P \in \phi'$. Then there exists a set $\{\overleftarrow{u}_j : j \in [m]\}$ such that the following conditions hold:*

- $\{\overleftarrow{r}_i : i \in [n]\} \cup \{\overleftarrow{u}_j : j \in [m]\}$ *is a star;*
- $|\overleftarrow{u}_j| \leqslant |\overleftarrow{x}_j|$ *for all* $j \in [m]$*;*
- $\overrightarrow{u}_j$ *is linked to* $\overrightarrow{x}_j$ *for all* $j \in [m]$*;*
- $\overleftarrow{x}_j \bigwedge_i \overrightarrow{r}_i \bigwedge_{k \neq j} \overrightarrow{x}_k \leqslant \overleftarrow{u}_j \leqslant \overleftarrow{x}_j$ *for all* $j \in [m]$*;*
- $\bigcup_{j=1}^m X_j \cup \bigcup_{i=1}^n A_i = \bigcup_{j=1}^m U_j \cup \bigcup_{i=1}^n A_i$.

Proof. Let us start with a set of separations

$$Y = \{\overleftarrow{y}_i : i \in [n+m]\},$$

and some arbitrary order on the set of pairs $Y^{(2)}$. Initially we set $\overleftarrow{y}_i = \overleftarrow{x}_j$ for $j \in [m]$ and $\overleftarrow{y}_{m+i} = \overleftarrow{r}_i$ for $i \in [n]$. For each pair $\{\overleftarrow{y}_i, \overleftarrow{y}_j\}$ in order we apply Lemma 1.4.3 to this pair of separations and replace $\{\overleftarrow{y}_i, \overleftarrow{y}_j\}$ with the nested pair given by Lemma 1.4.3. After we have done this for each pair, we let $\overleftarrow{u}_j := \overleftarrow{y}_j$ for each $j \in [m]$.

Note that, since each $\overleftarrow{x}_j$ is ϕ-inessential, and with each application of Lemma 1.4.3 we only ever replace a separation by one less than or equal to it, $\overleftarrow{y}_j$ is also ϕ-inessential at each stage of this process for $j \in [m]$. Also, $\{\overleftarrow{r}_i : i \in [n]\}$ is a star, and so if we apply Lemma 1.4.3 to a pair $\overleftarrow{r}_i$ and $\overleftarrow{r}_k$, neither is changed. Therefore, by the comment after Lemma 1.4.3, we may assume that at every stage in the process $\overleftarrow{y}_{m+i} = \overleftarrow{r}_i$ for each $i \in [n]$. In particular at the end of the process we have that

$$Y = \{\overleftarrow{r}_i : i \in [n]\} \cup \{\overleftarrow{u}_j : j \in [m]\}.$$

To see that the first condition is satisfied we note that, given any pair of separations $\overleftarrow{y}_i$ and $\overleftarrow{y}_j \in Y$, at some stage in the process we applied Lemma 1.4.3 to this pair, and immediately after this step we have that $\overleftarrow{y}_i \leqslant \overrightarrow{y}_j$. Since Lemma 1.4.3 only ever replaces a separation with one less than or equal to it, it follows that at the end of the process Y is a star. Therefore the family $\{\overleftarrow{r}_i : i \in [n]\} \cup \{\overleftarrow{u}_j : j \in [m]\}$ forms a star.

To see that the second condition is satisfied we note that, whenever we apply Lemma 1.4.3 we only ever replace a separation with one whose order is less than or equal to the order of the original separation.

To see that the third condition is satisfied we note that whenever we apply Lemma 1.4.3 we only ever replace a separation with one whose inverse is linked to the inverse of the original separation. Therefore it would be sufficient to show that the property of being linked to is transitive. Indeed, suppose that $\overrightarrow{r} > \overrightarrow{s} > \overrightarrow{t}$, $\overrightarrow{r}$ is linked to $\overrightarrow{s}$ and $\overrightarrow{s}$ is linked to $\overrightarrow{t}$, all in some separation system S. Let $\overrightarrow{x} > \overrightarrow{t}$ also be in $\overrightarrow{S}_k$.

However, since $\vec{s}$ is linked to $\vec{t}$, it follows that $|\vec{x} \vee \vec{s}| \leqslant |\vec{x}|$. Then, since $\vec{x} \vee \vec{s} \geqslant \vec{s}$ and $\vec{r}$ is linked to $\vec{s}$, it follows that $|(\vec{x} \vee \vec{s}) \vee \vec{r}| \leqslant |\vec{x} \vee \vec{s}| \leqslant |\vec{x}|$. However, since $\vec{s} \leqslant \vec{r}$, $(\vec{x} \vee \vec{s}) \vee \vec{r} = \vec{x} \vee \vec{r}$, and so $\vec{r}$ is linked to $\vec{t}$.

To see that the fourth condition is satisfied let us consider $\overleftarrow{y}_j$ for some $j \in [m]$. There is some sequence of separations $\overleftarrow{x}_j = \overleftarrow{v}_0 \geqslant \overleftarrow{v}_1 \geqslant \ldots \geqslant \overleftarrow{v}_t = \overleftarrow{u}_j$ that are the values $\overleftarrow{y}_j$ takes during this process, corresponding to the t times we applied Lemma 1.4.3 to a pair containing the separation $\overleftarrow{y}_j$. Suppose that the other separations in those pairs were $\overleftarrow{y}_{i_1}, \overleftarrow{y}_{i_2}, \ldots, \overleftarrow{y}_{i_t}$, and let us denote by $\overleftarrow{w}_k$ the value of the separations $\overleftarrow{y}_{i_k}$ at the time which we applied Lemma 1.4.3 to the pair $\{\overleftarrow{y}_j, \overleftarrow{y}_{i_k}\}$.

We claim inductively that for all $0 \leqslant r \leqslant t$

$$\overleftarrow{v}_0 \bigwedge_{k=1}^{r} \vec{w}_k \leqslant \overleftarrow{v}_r \leqslant \overleftarrow{v}_0.$$

The statement clearly holds for $r = 0$. Suppose it holds for $r-1$. We obtain $\overleftarrow{v}_r$ by applying Lemma 1.4.3 to the pair $\{\overleftarrow{v}_{r-1}, \overleftarrow{w}_r\}$, giving us the pair $\{\overleftarrow{v}_r, \overleftarrow{z}\}$. We have that $\overleftarrow{v}_{r-1} \wedge \vec{z} = \overleftarrow{v}_r$ and so, since $\vec{w}_r \leqslant \vec{z}$ it follows that

$$\overleftarrow{v}_{r-1} \wedge \vec{w}_r \leqslant \overleftarrow{v}_r \leqslant \overleftarrow{v}_{r-1}.$$

By the induction hypothesis we know that

$$\overleftarrow{v}_0 \bigwedge_{k=1}^{r-1} \vec{w}_k \leqslant \overleftarrow{v}_{r-1} \leqslant \overleftarrow{v}_0,$$

and so

$$\overleftarrow{v}_0 \bigwedge_{k=1}^{r} \vec{w}_k \leqslant \overleftarrow{v}_{r-1} \wedge \vec{w}_r \leqslant \overleftarrow{v}_r \leqslant \overleftarrow{v}_{r-1} \leqslant \overleftarrow{v}_0$$

as claimed.

For each of the $\overleftarrow{w}_k$ there is some separation $\overleftarrow{s}_k$ from our original set (that is some $\overleftarrow{r}_i$ or $\overleftarrow{x}_j$) such that $\overleftarrow{w}_k \leqslant \overleftarrow{s}_k$ and so, since $\vec{s}_k \leqslant \vec{w}_k$, and since we apply Lemma 1.4.3 to each pair of separations in our original set, we have that

$$\overleftarrow{v}_0 \bigwedge_{i} \vec{r}_i \bigwedge_{k \neq j} \vec{x}_j \leqslant \overleftarrow{v}_0 \bigwedge_{k=1}^{t} \vec{w}_k.$$

So, recalling that $\overleftarrow{v}_0 = \overleftarrow{x}_j$ and $\overleftarrow{v}_t = \overleftarrow{u}_j$, we see that

$$\overleftarrow{x}_j \bigwedge_{i} \vec{r}_i \bigwedge_{k \neq j} \vec{x}_k \leqslant \overleftarrow{u}_j \leqslant \overleftarrow{x}_j$$

as claimed.

Finally we note that, if we apply Lemma 1.4.3 to a pair of separations (C, D) and (E, F), resulting in the nested pair $\{(C', D'), (E', F')\}$, then

$$C \cup E = C' \cup E'.$$

Indeed, we have that $(C \cap F', D \cup E') = (C', D')$ and $(E \cap D', F \cup C') = (E', F')$ and so we have that $C' \cup E' = (C \cap F') \cup E' \supseteq C$ and similarly $C' \cup E' = C' \cup (E \cap D') \supseteq E$ and so $C' \cup E' \supseteq C \cup E$. However, since $C' \subseteq C$ and $E' \subseteq E$ we also have $C' \cup E' \subseteq C \cup E$. $\square$

1.4.2 Refining the essential parts

The content of Lemma 1.4.5 can be thought of as a procedure for turning an arbitrary set of separations into a star which is in some way 'close' to the original set, and is linked pairwise to the original set. We note that the second property guarantees us that this star lies in the same S_k as the original set.

Let us say a few words about the other properties of the star which represent this closeness. It will be useful to think about these properties in terms of how we can use this lemma to refine further an essential part in a k-tangle-distinguishing tree-decomposition.

Suppose $\{\overleftarrow{x}_j : j \in [m]\} = \mathcal{M}_I(O)$ for some k-tangle O, and $\{\overleftarrow{r}_i : i \in [n]\}$ is the star of separations at the vertex where O is contained in a tree-decomposition, specifically one where each r_i distinguishes efficiently some pair of k-tangles. By applying Lemma 1.4.5 we get a star $\{\overleftarrow{u}_j : j \in [m]\}$ satisfying the conclusions of the lemma. For each non-trivial $\overleftarrow{x}_j$, by Lemma 1.4.1, there exists an irredundant S_k-tree over $\mathcal{T}_k^* \cup \{\overleftarrow{x}_j\}$ containing $\overleftarrow{x}_j$ as a leaf separation, such that $\overleftarrow{x}_j$ is not the image of any other edge. We can then use Lemma 1.2.1 to shift each of these S_k-trees onto $\overrightarrow{u}_j$, giving us an S_k-tree over $\mathcal{T}_k^* \cup \{\overleftarrow{u}_j\}$. If $\overleftarrow{x}_j$ is trivial then so is $\overleftarrow{u}_j$, and so there is an obvious S_k-tree over $\mathcal{T}_k^* \cup \{\overleftarrow{u}_j\}$ containing $\overrightarrow{u}_j$ as a leaf separation, that with a single edge corresponding to u_j.

Doing the same for each k-tangle in the graph and taking the union all of these S_k-trees, together with the tree-decomposition from Corollary 1.3.2, will give us a refinement of this tree-decomposition which maintains the property of each inessential part being too small to contain a k-tangle, but also further refines the essential parts. The properties of the star given by Lemma 1.4.5 give us some measurement of how effective this process is in refining the essential parts of the graph.

We first note that, given a k-tangle O, which is contained in some part V_t of a k-tangle-distinguishing tree-decomposition, by the fifth property in Lemma 1.4.5 every vertex in the part V_t which lies on the small side of some maximal inessential separation in O will be in some inessential part of this refinement.

However this property is also satisfied by the rather naive refinement formed by just taking the union of some small separations (A_i, V) with the A_i covering the same vertex set. The problem with this naive decomposition is it does not really refine the part V_t, since there is a still a part with vertex set V_t in the new decomposition. Ideally we would like our refinement to make this essential part as small as possible, to more precisely exhibit how the k-tangle O lies in the graph.

Our refinement comes some way towards this, as evidenced by the fourth condition . For example if we have some separation $\overleftarrow{s} = (A, B)$ which lies 'behind' some maximal inessential separation in O, that is $\overleftarrow{s} \leqslant \overleftarrow{x}_j$ for some j, and is nested $\mathcal{M}_I(O) \cup \{\overleftarrow{r}_i : i \in [n]\}$, then it is easy to check that the fourth property guarantees it will also lie behind some $\overleftarrow{u}_k$ given by Lemma 1.4.5. So, in the refined tree-decomposition, the part containing O will not contain any vertices that lie strictly in the small side of such a separation, $A \setminus B$.

Suppose $\{\overleftarrow{r}_i : i \in [n]\}$ is an essential part in a tree-decomposition $(T, \mathcal{V})$ containing a tangle O. We say a vertex $v \in V$ is *inessentially separated from O relative to* $(T, \mathcal{V})$ if there is a separation (A, B) which is nested with $\mathcal{M}_I(O) \cup \{\overleftarrow{r}_i : i \in [n]\}$ such that $v \in A \setminus B$, and there exists some $(X, Y) \in \mathcal{M}_I(O)$ such that $(A, B) \leqslant (X, Y)$. For example, in Figure 1.2, the vertices in the long paths are inessentially separated from the tangles corresponding to the complete subgraphs relative to the canonical tangle-distinguishing tree-decomposition.

Theorem 1.4.6. *For every graph G and $k \geqslant 3$ there exists a tree-decomposition $(T, \mathcal{V})$ of G of adhesion $<k$ with the following properties*

- *The tree-decomposition $(T', \mathcal{V}')$ induced by the essential separations is canonical and distinguishes every k-tangle in G;*

- *The torso of every inessential part has branch-width $<k$.*

- *For every essential part V_t which contains a tangle O, there are no vertices $v \in V_t$ which are inessentially separated from O relative to $(T', \mathcal{V}')$.*

Given a vertex $v \in V$ we say that that x is *well separated* from O if there is a separation (A, B) which is nested with $\mathcal{M}(O)$ such that $v \in A \setminus B$, and there exists some $(X, Y) \in \mathcal{M}(O)$ such that $(A, B) \leqslant (X, Y)$.

We can think of the vertices which are well separated from O as being 'far away' from O in the graph. Indeed, if $\mathcal{M}(O)$ is a star, then O is a k-block, and the set of vertices well separated from O are just the vertices not in the k-block. In general a tangle will not correspond as closely to a concrete set of vertices as a k-block, and crossing separations in $\mathcal{M}(O)$ somehow demonstrate the uncertainty of whether a vertex 'lives in' O or not. However, if a separation $(A, B) \in O$ is nested with $\mathcal{M}(O)$, then O should be in some way fully contained in B, and so the vertices in $A \setminus B$ are 'far away' from O.

Question 1.4.7. *For every graph G, does there exist a tree-decomposition which distinguishes the k-tangles in a graph, whose essential parts are small in the sense that for each k-tangle O, there is no vertex x which can be well separated from O in the part of the tree-decomposition which contains O? Does there exist such a tree-decomposition with the further property that the inessential parts have branch-width $<k$?*

Chapter 2

Duality theorems for blocks and tangles in graphs

2.1 Introduction

There are a number of theorems about the structure of sparse graphs that assert a duality between the existence of a highly connected substructure and a tree-like overall structure. For example, if a connected graph G has no 2-connected subgraph or minor, it is a tree. Less trivially, the graph minor structure theorem of Robertson and Seymour says that if G has no K_n-minor then it has a tree-decomposition into parts that are 'almost' embeddable in a surface of bounded genus; see [43].

Another example of a highly connected substructure is that of a *k-block*, introduced by Mader [99] in 1978 and studied more recently in [34, 36, 37]. This is a maximal set of at least k vertices in a graph G such that no two of them can be separated in G by fewer than k vertices.

One of our main results is that the non-existence of a k-block, too, is always witnessed by a tree structure (Theorem 2.1.2). This problem was raised in [34, Sec. 7]. In [51], Diestel and Oum used a new theory of 'abstract separation systems' [44] in pursuit of this problem, but were unable to find the tree structures needed: the simplest witnesses to the nonexistence of k-blocks they could find are described in [50], but they are more complicated than trees. Our proof of Theorem 2.1.2, and the rest of this paper, are still based on the theory developed in [44] and [51], and we show that there are tree-like obstructions to the existence of k-blocks after all.

Tangles, introduced by Robertson and Seymour in [110], are substructures of graphs that also signify high local connectivity, but of a less tangible kind than subgraphs, minors, or blocks. Basically, a tangle does not tell us 'what' that substructure is, but only 'where' it is: by orienting all the low-order separations of the graph in some consistent way, which we then think of as pointing 'towards the tangle'. See Section 2.2.2 below, or [43], for a formal introduction to tangles.

Tangles come with dual tree structures called *branch decompositions*. Although defined differently, they can be thought of as tree-decompositions of a particular kind. In [51, 52], Diestel and Oum generalised the notion of tangles to ways of consistently orienting the low-order separations of a graph so as to describe other known types of highly connected substructures too, such as those dual to tree-decompositions of low width. We shall build on [51] to find dual tree structures for various types of tangles, for blocks, and for 'profiles': a common generalisation of blocks and tangles introduced in [48] and defined formally in Section 2.2.3.

Let us describe our results more precisely. A classical k-tangle, as in [110], is an orientation of all the separations $\{A, B\}$ of order $< k$ in a graph G, say as (A, B) rather than as (B, A), so that no three of these cover G by the subgraphs that G induces on their 'small sides' A. Let

$\mathcal{T}$ denote the set of all such forbidden triples of oriented separations of G, irrespective of their order. The tangle duality theorem of Robertson and Seymour then asserts that, given k, either the set S_k of all the separations of G of order $< k$ can be oriented in such a way as to induce no triple from $\mathcal{T}$ – an orientation of S_k we shall call a $\mathcal{T}$-tangle – or G has a tree-decomposition of a particular type: one from which it is clear that G cannot have a k-tangle, i.e., a $\mathcal{T}$-tangle of S_k.

Now consider any superset $\mathcal{F}$ of $\mathcal{T}$. Given k, the orientations of S_k with no subset in $\mathcal{F}$ will then be particular types of tangles. Our $\mathcal{F}$-tangle duality theorem yields duality theorems for all these: if G contains no such 'special' tangle, it will have a tree-decomposition that witnesses this. Formally, to every such $\mathcal{F}$ and k there will correspond a class $\mathcal{T}_{\mathcal{F}}(k)$ of tree-decompositions that witness the non-existence of an $\mathcal{F}$-tangle of S_k, and which are shown to exist whenever a graph has no $\mathcal{F}$-tangle of S_k:

Theorem 2.1.1. *For every finite graph G, every set $\mathcal{F} \supseteq \mathcal{T}$ of sets of separations of G, and every integer $k > 2$, exactly one of the following statements holds:*

- *G admits an $\mathcal{F}$-tangle of S_k;*
- *G has a tree-decomposition in $\mathcal{T}_{\mathcal{F}}(k)$.*

Every k-block also defines an orientation of S_k: as no separation $\{A, B\} \in S_k$ separates it, it lies entirely in A or entirely in B. These orientations of S_k need not be k-tangles, so we cannot apply Theorem 2.1.1 to obtain a duality theorem for k-blocks. But still, we shall be able to define classes $\mathcal{T}_{\mathcal{B}}(k)$ of tree-decompositions that witness the non-existence of a k-block, in the sense that graphs with such a tree-decomposition cannot contain one, and which always exist for graphs without a k-block:

Theorem 2.1.2. *For every finite graph G and every integer $k > 0$ exactly one of the following statements holds:*

- *G contains a k-block;*
- *G has a tree-decomposition in $\mathcal{T}_{\mathcal{B}}(k)$.*

Finally, we define classes $\mathcal{T}_{\mathcal{P}}(k)$ of tree-decompositions which graphs with a k-profile cannot have, and prove the following duality theorem for profiles:

Theorem 2.1.3. *For every finite graph G and every integer $k > 2$ exactly one of the following statements holds:*

- *G has a k-profile;*
- *G has a tree-decomposition in $\mathcal{T}_{\mathcal{P}}(k)$.*

For readers already familiar with profiles [48] we remark that, in fact, we shall obtain a more general result than Theorem 2.1.3: our Theorem 2.3.9 is a duality theorem for all regular profiles in abitrary submodular abstract separation systems, including the standard ones in graphs and matroids but many others too [54].

Like Theorem 2.1.1, Theorems 2.1.2 and 2.1.3 are 'structural' duality theorems in that they identify a structure that a graph G cannot have if it contains a k-block or k-profile, and must have if it does not. Alternatively, we can express the same duality more compactly in terms of graph invariants, as follows. Let

$$\begin{aligned} \beta(G) &:= \max\{\, k \mid G \text{ has a } k\text{-block} \,\} \\ \pi(G) &:= \max\{\, k \mid G \text{ has a } k\text{-profile} \,\} \end{aligned}$$

be the *block number* and the *profile number* of G, respectively, and let

$$\begin{aligned}\mathrm{bw}(G) &:= \min\{\, k \mid G \text{ has a tree-decomposition in } \mathcal{T}_{\mathcal{B}}(k+1)\,\}\\ \mathrm{pw}(G) &:= \min\{\, k \mid G \text{ has a tree-decomposition in } \mathcal{T}_{\mathcal{P}}(k+1)\,\}\end{aligned}$$

be its *block-width* and *profile-width*. Theorems 2.1.2 and 2.1.3 can now be rephrased as

Corollary 2.1.4. *As invariants of finite graphs, the block and profile numbers agree with the block- and profile-widths:*

$$\beta = \mathrm{bw} \quad \textit{and} \quad \pi = \mathrm{pw}. \qquad \square$$

In Section 2.2 we introduce just enough about abstract separation systems [44] to state the fundamental duality theorem of [51], on which all our proofs will be based. In Section 2.3 we give a proof of our main result, a duality theorem for regular profiles in submodular abstract separation systems. In Section 2.4 we apply this to obtain structural duality theorems for k-blocks and k-profiles, and deduce Theorems 2.1.1–2.1.3 as corollaries. In Section 2.5 we derive some bounds for the above width-parameters in terms of tree-width and branch-width.

Any terms or notation left undefined in this paper are explained in [43].

2.2 Background Material

2.2.1 Separation systems

A *separation* of a graph G is a set $\{A, B\}$ of subsets of $V(G)$ such that $A \cup B = V$, and there is no edge of G between $A \setminus B$ and $B \setminus A$. There are two *oriented separations* associated with a separation, (A, B) and (B, A). Informally we think of (A, B) as *pointing towards* B and *away from* A. We can define a partial ordering on the set of oriented separations of G by

$$(A, B) \leqslant (C, D) \text{ if and only if } A \subseteq C \text{ and } B \supseteq D.$$

The *inverse* of an oriented separation (A, B) is the separation (B, A), and we note that mapping every oriented separation to its inverse is an involution which reverses the partial ordering.

In [51] Diestel and Oum generalised these properties of separations of graphs and worked in a more abstract setting. They defined a *separation system* $(\vec{S}, \leqslant, {}^*)$ to be a partially ordered set $\vec{S}$ with an order-reversing involution *. The elements of $\vec{S}$ are called *oriented separations.* Often a given element of $\vec{S}$ is denoted by $\vec{s}$, in which case its inverse $\vec{s}^{\,*}$ will be denoted by $\overleftarrow{s}$, and vice versa. Since * is ordering reversing we have that, for all $\vec{r}, \vec{s} \in S$,

$$\vec{r} \leqslant \vec{s} \text{ if and only if } \overleftarrow{r} \geqslant \overleftarrow{s}.$$

A *separation* is a set of the form $\{\vec{s}, \overleftarrow{s}\}$, and will be denoted by simply s. The two elements $\vec{s}$ and $\overleftarrow{s}$ are the *orientations* of s. The set of all such pairs $\{\vec{s}, \overleftarrow{s}\} \subseteq \vec{S}$ will be denoted by S. If $\vec{s} = \overleftarrow{s}$ we say s is *degenerate.* Conversely, given a set $S' \subseteq S$ of separations we write $\vec{S'} := \bigcup S'$ for the set of all orientations of its elements. With the ordering and involution induced from $\vec{S}$, this will form a separation system.

Given a separation of a graph $\{A, B\}$ we can identify it with the pair $\{(A, B), (B, A)\}$ and in this way any set of oriented separations in a graph which is closed under taking inverses forms a separation system. When we refer to an oriented separation in a context where the notation explicitly indicates orientation, such as $\vec{s}$ or (A, B), we will usually suppress the prefix "oriented" to improve the flow of the narrative.

The *separator* of a separation $s = \{A, B\}$ in a graph, and of its orientations $\vec{s}$, is the set $A \cap B$. The *order* of s and $\vec{s}$, denoted as $|s|$ or as $|\vec{s}|$, is the cardinality of the separator,

$|A \cap B|$. Note that if $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$ are separations then so are their *corner separations* $\vec{r} \vee \vec{s} := (A \cup C, B \cap D)$ and $\vec{r} \wedge \vec{s} := (A \cap C, B \cup D)$. Our function $\vec{s} \mapsto |\vec{s}|$ is clearly *symmetric* in that $|\vec{s}| = |\overleftarrow{s}|$, and *submodular* in that

$$|\vec{r} \vee \vec{s}| + |\vec{r} \wedge \vec{s}| \leq |\vec{r}| + |\vec{s}|$$

(in fact, with equality).

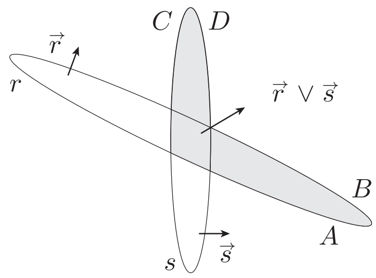

Figure 2.1: The corner separation $\vec{r} \vee \vec{s} = (A \cup C, B \cap D)$

If an abstract separation system $(\vec{S}, \leq, {}^*)$ forms a lattice, i.e., if there exist binary operations $\vee$ and $\wedge$ on $\vec{S}$ such that $\vec{r} \vee \vec{s}$ is the supremum and $\vec{r} \wedge \vec{s}$ is the infimum of $\vec{r}$ and $\vec{s}$, then we call $(\vec{S}, \leqslant, {}^*, \vee, \wedge)$ a *universe* of (oriented) separations. By (2.2.1), it satisfies De Morgan's law:

$$(\vec{r} \vee \vec{s})^* = \overleftarrow{r} \wedge \overleftarrow{s}. \tag{2.2.1}$$

Any real, non-negative, symmetric and submodular function on a universe of separations, usually denoted as $\vec{s} \mapsto |\vec{s}|$, will be called an *order function*.

Two separations r and s are *nested* if they have $\leqslant$-comparable orientations. Two oriented separations $\vec{r}$ and $\vec{s}$ are *nested* if r and s are nested.[1] We say that $\vec{r}$ *points towards* s (and $\overleftarrow{r}$ *points away from* s) if $\vec{r} \leqslant \vec{s}$ or $\vec{r} \leqslant \overleftarrow{s}$. So two nested oriented separations are either $\leqslant$-comparable, or they point towards each other, or they point away from each other. If $\vec{r}$ and $\vec{s}$ are not nested we say that they *cross*. A set of separations S is *nested* if every pair of separations in S is nested, and a separation s is *nested* with a nested set of separations S if $S \cup \{s\}$ is nested.

A separation $\vec{r} \in \vec{S}$ is *trivial in* $\vec{S}$, and $\overleftarrow{r}$ is *co-trivial*, if there exist an $s \in S$ such that $\vec{r} < \vec{s}$ and $\vec{r} < \overleftarrow{s}$. Note that if $\vec{r}$ is trivial, witnessed by some s, then, since the involution on $\vec{S}$ is order-reversing, we have $\vec{r} < \vec{s} < \overleftarrow{r}$. So, in particular, $\overleftarrow{r}$ cannot also be trivial. Separations $\vec{s}$ such that $\vec{s} \leqslant \overleftarrow{s}$, trivial or not, will be called *small*.

In the case of separations of a graph (V, E), the small separations are precisely those of the form (A, V). The trivial separations are those of the form (A, V) with $A \subseteq C \cap D$ for some separation $\{C, D\} \neq \{A, B\}$. Finally we note that there is only one degenerate separation in a graph, (V, V).

2.2.2 Tangle-tree duality in separation systems

Let $\vec{S}$ be a separation system. An *orientation* of S is a subset $O \subseteq \vec{S}$ which for each $s \in S$ contains exactly one of its orientations $\vec{s}$ or $\overleftarrow{s}$. Given a universe $\vec{U}$ of separations with an order function, such as all the oriented separations of a given graph, we denote by

$$\vec{S_k} = \{\vec{s} \in \vec{U} : |\vec{s}| < k\}$$

[1]In general we will use terms defined for separations informally for oriented separations when the meaning is clear, and vice versa.

the set of all its separations of order less than k. Note that $\vec{S_k}$ is again a separation system. But it is not necessarily a universe, since it may fail to be closed under the operations $\vee$ and $\wedge$.

If we have some structure $\mathcal{C}$ in a graph that is 'highly connected' in some sense, we should expect that no low order separation will divide it: that is, for every separation s of sufficiently low order, $\mathcal{C}$ should lie on one side of s but not the other. Then $\mathcal{C}$ will *orient* s as $\vec{s}$ or $\overleftarrow{s}$, choosing the orientation that 'points to where it lies' according to some convention. For graphs, our convention is that the orientated separation (A, B) points towards B. And that if $\mathcal{C}$ is a K_n-minor of G with $n \geqslant k$, say, then $\mathcal{C}$ 'lies on the side B' if it has a branch set in $B \setminus A$. (Note that it cannot have a branch set in $A \setminus B$ then.) Then $\mathcal{C}$ orients $\{A, B\}$ towards B by choosing (A, B) rather than (B, A). In this way, $\mathcal{C}$ induces an orientation of all of S_k.

The idea of [51], now, following the idea of tangles, was to *define* 'highly connected substructures' in this way: as orientations of a given set S of separations.

Any concrete example of 'highly-connected substructures' in a graph, such as a K_n-minor or a k-block, will not induce arbitrary orientations of S_k: these orientations will satisfy some consistency rules. For example, consider two separations $(A, B) < (C, D)$. If our 'highly connected' structure $\mathcal{C}$ orients $\{C, D\}$ towards D then, since $B \supseteq D$ it should not orient $\{A, B\}$ towards A.

We call an orientation O of a set S of separations in some universe $\vec{U}$ *consistent* if whenever we have distinct r and s such that $\vec{r} < \vec{s}$, the set O does not contain both $\overleftarrow{r}$ and $\vec{s}$. Note that a consistent orientation of S must contain all separations $\vec{r}$ that are trivial in S since, if $\vec{r} < \vec{s}$ and $\vec{r} < \overleftarrow{s}$, then $\overleftarrow{r}$ would be inconsistent with whichever orientation of s lies in O.

Given a set $\mathcal{F}$, we say that an orientation O of S *avoids* $\mathcal{F}$ if there is no $F \in \mathcal{F}$ such that $F \subseteq O$. So for example an orientation of S is consistent if it avoids $\mathcal{F} = \{\{\overleftarrow{r}, \vec{s}\} \subseteq \vec{S} : r \neq s, \vec{r} < \vec{s}\}$. In general we will define the highly connected structures we consider by the collection $\mathcal{F}$ of subsets they avoid. For example a *tangle of order k*, or *k-tangle*, in a graph G is an orientation of S_k which avoids the set of triples

$$\mathcal{T} = \{\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} \subseteq \vec{U} : \bigcup_{i=1}^{3} G[A_i] = G\}. \tag{2.2.2}$$

Here, the three separations need not be distinct, so any $\mathcal{T}$-avoiding orientation of S_k will be consistent. More generally, we say that a consistent orientation of a set S of separations which avoids some given set $\mathcal{F}$ is an *$\mathcal{F}$-tangle (of S)*.

Given a set S of separations, an *S-tree* is a pair (T, α), of a tree T and a function $\alpha : \vec{E(T)} \to \vec{S}$ from the set $\vec{E(T)}$ of directed edges of T such that

- For each edge $(t_1, t_2) \in \vec{E(T)}$, if $\alpha(t_1, t_2) = \vec{s}$ then $\alpha(t_2, t_1) = \overleftarrow{s}$.

The S-tree is said to be *over* a set $\mathcal{F}$ if

- For each vertex $t \in T$, the set $\{\alpha(t', t) : (t', t) \in \vec{E(T)}\}$ is in $\mathcal{F}$.

Particularly interesting classes of S-trees are those over sets $\mathcal{F}$ of 'stars'. A set σ of nondegenerate oriented separations is a *star* if $\vec{r} \leqslant \overleftarrow{s}$ for all distinct $\vec{r}, \vec{s} \in \sigma$. We say that a set $\mathcal{F}$ *forces* a separation $\vec{r}$ if $\{\overleftarrow{r}\} \in \mathcal{F}$. And $\mathcal{F}$ is *standard* if it forces every trivial separation in $\vec{S}$.

The main result of [51] asserts a duality between S-trees over $\mathcal{F}$ and $\mathcal{F}$-tangles when $\mathcal{F}$ is a standard set of stars satisfying a certain closure condition. Let us describe this next.

Suppose we have a separation $\vec{r}$ which is neither trivial nor degenerate. Let $S_{\geqslant \vec{r}}$ be the set of separations $x \in S$ that have an orientation $\vec{x} \geqslant \vec{r}$. Given $x \in S_{\geqslant \vec{r}} \setminus \{r\}$ we have, since $\vec{r}$ is nontrivial, that only one of the two orientations of x, say $\vec{x}$, is such that $\vec{x} \geqslant \vec{r}$ and x is not degenerate. For every $\vec{s} \geqslant \vec{r}$ we can define a function $f{\downarrow}^{\vec{r}}_{\vec{s}}$ on $\vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$ by[2]

$$f{\downarrow}^{\vec{r}}_{\vec{s}}\,(\vec{x}) := \vec{x} \vee \vec{s} \text{ and } f{\downarrow}^{\vec{r}}_{\vec{s}}\,(\overleftarrow{x}) := (\vec{x} \vee \vec{s})^*.$$

In general, the image in $\vec{U}$ of this function need not lie in $\vec{S}$.

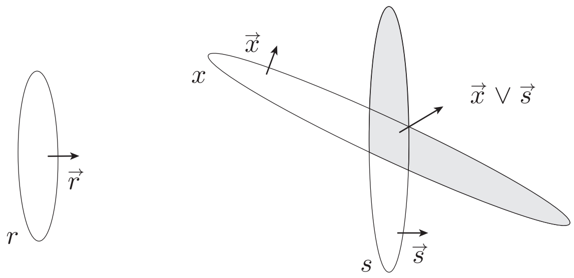

Figure 2.2: Shifting a separation $\vec{x} \geqslant \vec{r}$ to $f{\downarrow}^{\vec{r}}_{\vec{s}}\,(\vec{x}) = \vec{x} \vee \vec{s}$.

We say that $\vec{s} \in \vec{S}$ *emulates* $\vec{r} \in \vec{S}$ *in* $\vec{S}$ if $\vec{r} \leqslant \vec{s}$ and the image of $f{\downarrow}^{\vec{r}}_{\vec{s}}$ is contained in $\vec{S}$. Given a standard set $\mathcal{F}$ of stars, we say further that $\vec{s}$ *emulates* $\vec{r}$ in $\vec{S}$ *for* $\mathcal{F}$ if $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$ and the image under $f{\downarrow}^{\vec{r}}_{\vec{s}}$ of every star $\sigma \subseteq \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$ that contains some separation $\vec{x}$ with $\vec{x} \geqslant \vec{r}$ is again in $\mathcal{F}$.

We say that a separation system $\vec{S}$ is *separable* if for any two nontrivial and nondegenerate separations $\vec{r}, \overleftarrow{r'} \in \vec{S}$ such that $\vec{r} \leqslant \vec{r'}$ there exists a separation $s \in S$ such that $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$ and $\overleftarrow{s}$ emulates $\overleftarrow{r'}$ in $\vec{S}$. We say that $\vec{S}$ is $\mathcal{F}$*-separable* if for all nontrivial and nondegenerate $\vec{r}, \overleftarrow{r'} \in \vec{S}$ that are not forced by $\mathcal{F}$ and such that $\vec{r} \leqslant \vec{r'}$ there exists a separation $s \in S$ with an orientation $\vec{s}$ that emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$ and such that $\overleftarrow{s}$ emulates $\overleftarrow{r'}$ in $\vec{S}$ for $\mathcal{F}$. Often one proves that $\vec{S}$ is $\mathcal{F}$-separable in two steps, by first showing that it is separable, and then showing that $\mathcal{F}$ is *closed under shifting*: that whenever $\vec{s}$ emulates (in $\vec{S}$) some nontrivial and nondegenerate $\vec{r}$ not forced by $\mathcal{F}$, then it does so for $\mathcal{F}$.

We are now in a position to state the Strong Duality Theorem from [51].

Theorem 2.2.1. *Let $\vec{S}$ be a separation system in some universe of separations, and $\mathcal{F}$ a standard set of stars. If $\vec{S}$ is $\mathcal{F}$-separable, exactly one of the following assertions holds:*

- *There exists an S-tree over $\mathcal{F}$;*
- *There exists an $\mathcal{F}$-tangle of S.*

The property of being $\mathcal{F}$-separable may seem a rather strong condition. However in [52] it is shown that for every graph the set $\vec{S_k}$ is separable, and all the sets $\mathcal{F}$ of stars whose exclusion describes classical notions of highly connected substructures are closed under shifting. Hence in all these cases $\vec{S_k}$ is $\mathcal{F}$-separable, and Theorem 2.2.1 applies.

One of our main tasks will be to extend the applicability of Theorem 2.2.1 to sets $\mathcal{F}$ of separations that are not stars, by constructing a related set $\mathcal{F}^*$ of stars whose exclusion is tantamount to excluding $\mathcal{F}$.

2.2.3 Blocks, tangles, and profiles

Suppose we have a graph $G = (V, E)$ and are considering the set U of its separations. As mentioned before, it is easy to see that the tangles of order k in G, as defined by Robertson and

[2]The exclusion of $\overleftarrow{r}$ here is for a technical reason: if $\vec{r} < \overleftarrow{r}$, we do not want to define $f{\downarrow}^{\vec{r}}_{\vec{s}}\,(\overleftarrow{r})$ explicitly, but implicitly as the inverse of $f{\downarrow}^{\vec{r}}_{\vec{s}}\,(\vec{r})$.

Seymour [110], are precisely the $\mathcal{T}$-tangles of $S_k = \{s \in U : |s| < k\}$. In this case, if we just consider the set of stars in $\mathcal{T}$,

$$\mathcal{T}^* := \{ \{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} \subseteq \vec{S} : \\ (A_i, B_i) \leqslant (B_j, A_j) \text{ for all } i, j \text{ and } \textstyle\bigcup_{i=1}^3 G[A_i] = G \},$$

then the $\mathcal{T}^*$-tangles of S_k are precisely its $\mathcal{T}$-tangles [52]. That is, a consistent orientation of S_k avoids $\mathcal{T}$ if and only if it avoids $\mathcal{T}^*$. It is a simple check that $\mathcal{T}^*$ is a standard set of stars which is closed under shifting, and hence Theorem 2.2.1 tells us that every graph either has a tangle of order k or an S_k-tree over $\mathcal{T}^*$, but not both.

Another highly connected substructure that has been considered recently in the literature are k-blocks. Given $k \in \mathbb{N}$ we say a set I of at least k vertices in a graph G is *$(<k)$-inseparable* if no set Z of fewer than k vertices separates any two vertices of $I \setminus Z$ in G. A maximal $(<k)$-inseparable set of vertices is called a *k-block*. These objects were first considered by Mader [99], but have been the subject of recent research [34, 36, 37].

As indicated earlier, every k-block b of G defines an orientation $O(b)$ of S_k:

$$O(b) := \{(A, B) \in \vec{S_k} : b \subseteq B\}.$$

Indeed, for each separation $\{A, B\} \in S_k$ exactly one of (A, B) and (B, A) will be in $O(b)$, since $A \cap B$ is too small to contain b and does not separate any two of its vertices. Hence, $O(b)$ is indeed an orientation of S_k. Note also that $O(b) \neq O(b')$ for distinct k-blocks $b \neq b'$: by their maximality as k-indivisible sets of vertices there exists a separation $\{A, B\} \in S_k$ such that $A \setminus B$ contains a vertex of b and $B \setminus A$ contains a vertex of b', which implies that $(A, B) \in O(b')$ and $(B, A) \in O(b)$.

The orientations $O(b)$ of S_k defined by a k-block b clearly avoid

$$\mathcal{B}_k := \big\{\{(A_i, B_i) : i \in I\} \subseteq \vec{U} : |\bigcap_{i \in I} B_i| < k\big\},$$

since $b \subseteq B_i$ for every $(A_i, B_i) \in O(b)$ and $|b| \geqslant k$. Also, it is easily seen that every $O(b)$ is consistent. Thus, every such orientation $O(b)$ is an $\mathcal{F}$-tangle of S_k for $\mathcal{F} = \mathcal{B}_k$. Conversely, if $O \subseteq S_k$ is a $\mathcal{B}_k$-tangle of S_k, then $b := \bigcap \{B \mid (A, B) \in O\}$ is easily seen to be a k-block, and $O = O(b)$. The orientations of S_k that are defined by a k-block, therefore, are precisely its $\mathcal{B}_k$-tangles.

The $\mathcal{B}_k$-tangles of S_k and its $\mathcal{T}$-tangles (i.e., the ordinary k-tangles of G) share the property that if they contain separations (A, B) and (C, D), then they cannot contain the separation $(B \cap D, A \cup C)$. Indeed, clearly this condition is satisfied by $O(b)$ for any k-block b, since if $b \subseteq B$ and $b \subseteq D$ then $b \subseteq B \cap D$ and hence $b \not\subseteq A \cup C$ if $\{B \cap D, A \cup C\} \in S_k$. For tangles, suppose that some tangle contains such a triple $\{(A, B), (C, D), (B \cap D, A \cup C)\}$. Since $\{A, B\}$ and $\{C, D\}$ are separations of G, every edge not contained in $G[A]$ or $G[C]$ must be in $G[B]$ and $G[D]$, and hence in $G[B \cap D]$. Therefore $G[A] \cup G[C] \cup G[B \cap D] = G$, contradicting the fact that the tangle avoids $\mathcal{T}$.

Informally, if we think of the side of an oriented separation to which it points as 'large', then the orientations of S_k that form a tangle or are induced by a k-block have the natural property that if B is the large side of $\{A, B\}$ and D is the large side of $\{C, D\}$ then $B \cap D$ should be the large side of $\{A \cup C, B \cap D\}$ – if this separation is also in S_k, and therefore oriented by O. That is, the largeness of separation sides containing blocks or tangles is preserved by taking intersections.

Consistent orientations with this property are known as 'profiles'. Formally, a *k-profile in G* is a $\mathcal{P}$-tangle of S_k where

$$\mathcal{P} := \big\{\sigma \subseteq \vec{U} \mid \exists\, A, B, C, D \subseteq V : \sigma = \{(A, B), (C, D), (B \cap D, A \cup C)\}\big\}.$$

As we have seen,

Lemma 2.2.2. *All orientations of S_k that are tangles, or of the form $O(b)$ for some k-block b in G, are k-profiles in G.* $\square$

We remark that, unlike in the case of $\mathcal{T}$, the subset $\mathcal{P}'$ consisting of just the stars in $\mathcal{P}$ yields a wider class of tangles: there are $\mathcal{P}'$-tangles of S_k that are not $\mathcal{P}$-tangles, i.e., which are not k-profiles.

More generally, if $\vec{S}$ is any separation system contained in some universe $\vec{U}$, we can define a *profile of* S to be any $\mathcal{P}$-tangle of S where

$$\mathcal{P} := \big\{\sigma \subseteq \vec{U} \mid \exists\, \vec{r}, \vec{s} \in \vec{U} : \sigma = \{\vec{r}, \vec{s}, \overleftarrow{r} \wedge \overleftarrow{s}\}\big\}.$$

In particular, all $\mathcal{F}$-tangles with $\mathcal{F} \supseteq \mathcal{P}$ will be profiles.

The initial aim of Diestel and Oum in developing their duality theory [51] had been to find a duality theorem broad enough to imply duality theorems for k-blocks and k-profiles. Although their theory gave rise to a number of unexpected results [52], a duality theorem for blocks and profiles was not among these; see [50] for a summary of their findings on this problem.

Our next goal is to show that their Strong Duality Theorem does implies duality theorems for blocks and profiles after all.

2.3 A duality theorem for abstract profiles

In this section we will show that Theorem 2.2.1 can be applied to many more types of profiles than originally thought. These will include both k-profiles and k-blocks in graphs.

We say that a separation system in some universe[3] is *submodular* if for every two of its elements $\vec{r}, \vec{s}$ it also contains at least one of $\vec{r} \wedge \vec{s}$ and $\vec{r} \vee \vec{s}$. Given any (submodular) order function on a universe, the separation system

$$\vec{S_k} = \{\vec{r} \in \vec{U} : |\vec{r}| < k\}$$

is submodular for each k. In particular, for any graph G, its universe $\vec{U}$ of separations and, for any integer $k \geqslant 1$, the separation system $\vec{S_k}$, is submodular.

We say that a subset O of U is *strongly consistent* if it does not contain both $\overleftarrow{r}$ and $\vec{s}$ for any $\vec{r}, \vec{s} \in \vec{S}$ with $\vec{r} < \vec{s}$ (but not necessarily $r \neq s$, as in the definition of 'consistent'). An orientation O of S, therefore, is strongly consistent if and only if for every $\vec{s} \in O$ it also contains every $\vec{r} \leq \vec{s}$ with $r \in S$. In particular, then, O cannot contain any $\vec{s}$ such that $\overleftarrow{s} \leq \vec{s}$ (i.e., with $\overleftarrow{s}$ is small).

Let us call an orientation O of S *regular* if it contains all the small separations in $\vec{S}$.

Lemma 2.3.1. *An orientation O of a separation system $\vec{S}$ is strongly consistent if and only if it is consistent and regular.*

Proof. Clearly every strongly consistent orientation O is also consistent. Suppose some small $\vec{s} \in \vec{S}$ is not in O. Then $\overleftarrow{s} \in O$, since O is an orientation of S. Thus, $\vec{s} < \overleftarrow{s} \in O$. But this implies $\vec{s} \in O$, since O is strongly consistent, contradicting the choice of $\vec{s}$. Hence O contains every small separation.

[3] Although submodular separation systems $\vec{S}$ have to lie in some universe $\vec{U}$ in order for $\wedge$ and $\vee$ to be defined on $\vec{S}$ (but with images that may lie in $\vec{U} \setminus \vec{S}$), the choice of $\vec{U}$, given $\vec{S}$, will not matter to us. We shall therefore usually introduce submodular separation systems $\vec{S}$ without formally introducing such a universe $\vec{U} \supseteq \vec{S}$.

Conversely, suppose O is a consistent orientation of S that is not strongly consistent. Then O contains two distinct oriented separations $\overleftarrow{r}$ and $\vec{s}$ such that $\vec{r} < \vec{s}$ and $r = s$.

Thus, $\overleftarrow{s} = \vec{r} < \vec{s}$ is small but not in O, as $\vec{s} \in O$. Hence O does not contain all small separations in $\vec{S}$. □

For example, the $\mathcal{B}_k$-tangles of S_k in a graph, as well as its ordinary tangles of order k, are regular by Lemma 2.3.1: they clearly contain all small separations in $\vec{S_k}$, those of the form (A, V) with $|A| < k$, since they cannot contain their inverses (V, A). More generally:

Lemma 2.3.2. *For $k > 2$, every k-profile in a graph $G = (V, E)$ is regular.*

Proof. We have to show that every k-profile O in G contains every small separation in $\vec{S_k}$. Recall that these are precisely the separations (A, V) of G such that $|A| < k$.

Suppose first that $|A| < k - 1$. Let A' be any set such that $|A'| = k - 1$ and $A \subset A'$. Then $\{A', V\} \in S_k$, and $(A, V) < (A', V)$ as well as $(A, V) < (V, A')$. Since O contains (A', V) or (V, A'), its consistency implies that it also contains (A, V).

If $|A| = k - 1$ then, since $k > 2$, we can pick two non-empty sets $A', A'' \subsetneq A$ such that $A' \cup A'' = A$. Since $|A'|, |A''| < k - 1$, by the preceding discussion both (A', V) and (A'', V) lie in O. As $(V, A) \wedge (V, A'') = (V \cap V, A' \cup A'') = (V, A)$ and $\{(A', V), (A'', V), (V, A)\} \in \mathcal{P}$, the fact that O is a profile implies that $(V, A) \notin O$, so again $(A, V) \in O$ as desired. □

There can be exactly one irregular 1-profile in a graph $G = (V, E)$, and only if G is connected: the set $\{(V, \emptyset)\}$.

Graphs can also have irregular 2-profiles, but they are easy to characterise. Indeed, consider a 2-profile O and small separation $(\{x\}, V)$. Suppose first that x is a cutvertex of G, in the sense that there exists some $\{A, B\} \in S_2$ such that $A \cap B = \{x\}$ and neither A nor B equals V. Then $(\{x\}, V) < (A, B)$ and $(\{x\}, V) < (B, A)$, so the consistency of O implies that $(\{x\}, V) \in O$.

Therefore, if $(V, \{x\}) \in O$ then x is not a cutvertex of G. Then, for every other separation $\{A, B\}$, either $x \in A \setminus B$ or $x \in B \setminus A$, and so either $(B, A) < (V, \{x\})$ or $(A, B) < (V, \{x\})$. The consistency of O then determines that

$$O = O_x := \{(A, B) \in \vec{S_2} \colon x \in B \text{ and } (A, B) \neq (\{x\}, V)\},$$

which is indeed a profile.

We have shown that every graph contains, for each of its vertices x that is not a cutvertex, a unique 2-profile O_x that is not strongly consistent. However, the orientation

$$O'_x := \{(A, B) \in \vec{S_2} \colon x \in B \text{ and } (A, B) \neq (V, \{x\})\}$$

of S_2 is also a 2-profile which does contain every small separation in $\vec{S}_2$. (Indeed, $O'_x = O(b)$ for the unique block b containing x.) Since every graph contains a vertex which is not a cutvertex, it follows that

Lemma 2.3.3. *Every graph G contains a regular 2-profile.* □

Lemma 2.3.3 means that our goal to find a duality theorem for k-profiles in graphs has substance only for $k > 2$, for which Lemma 2.3.2 tells that all k-profiles are regular. In our pursuit of Theorems 2.1.2 and 2.1.3 it will therefore suffice to study regular $\mathcal{F}$-tangles of submodular separation systems $\vec{S}$, such as $\vec{S_k}$ for $\mathcal{F} = \mathcal{P}$.

So, until further notice:

> *Let $\vec{S}$ be any submodular separation system in some universe $\vec{U}$, and let $\mathcal{F}$ be a subset of $2^{\vec{S}}$ containing $\mathcal{P} \cap 2^{\vec{S}}$.*

Our aim will be to prove a duality theorem for the regular $\mathcal{F}$-tangles of S.

It will be instructive to keep in mind, as an example, the case of k-blocks, where $\mathcal{F} = \mathcal{B}_k$. In this case any triple $\{(A, B), (C, D), (D \cap B, A \cup C)\} \in \mathcal{P} \cap 2^{\vec{S_k}}$ is contained in $\mathcal{B}_k$, as $|B \cap D \cap (A \cup C)| < k$ since $(D \cap B, A \cup C) \in S_k$.

For ease of notation, let us write $\mathcal{P}_S := \mathcal{P} \cap 2^{\vec{S}}$, and put $\mathcal{P}_k := \mathcal{P}_{S_k}$ when U is the set of separations of a given graph. Note that an orientation of S avoids $\mathcal{P}$ if and only if it avoids $\mathcal{P}_S$, and an S-tree is over $\mathcal{P}$ if and only if it is over $\mathcal{P}_S$.

Our first problem is that, in order to apply Theorem 2.2.1, we need $\mathcal{F}$ to be a set of stars. Since our assumptions about $\mathcal{F}$ do not require this, our first aim is to turn $\mathcal{F}$ into a set $\mathcal{F}^*$ of stars such that the regular $\mathcal{F}$-tangles of S are precisely its regular $\mathcal{F}^*$-tangles.

Suppose we have some pair of separations $\vec{x_1}$ and $\vec{x_2}$ which are both contained in some set $\sigma \subseteq \vec{S}$. Since S, by assumption, is submodular, at least one of $\vec{x_1} \wedge \overleftarrow{x_2}$ and $\vec{x_2} \wedge \overleftarrow{x_1}$ must also be in $\vec{S}$. To *uncross* $\vec{x_1}$ *and* $\vec{x_2}$ *in* σ we replace $\{\vec{x_1}, \vec{x_2}\}$ with the pair $\{\vec{x_1} \wedge \overleftarrow{x_2}, \vec{x_2}\}$ in the first case and $\{\vec{x_1}, \vec{x_2} \wedge \overleftarrow{x_1}\}$ in the second case. We note that, in both cases the new pair forms a star and is pointwise $\leqslant$ the old pair $\{\vec{x_1}, \vec{x_2}\}$. Uncrossing every pair of separations in σ in turn, we can thus turn σ into a star σ^* of separations in at most $\binom{|\sigma|}{2}$ steps, since any star of two separations remains a star if one of its elements is replaced by a smaller separation, and a set of oriented separations is a star as soon as all its 2-subsets are stars. Note, however, that σ^* will not in general be unique, but will depend on the order in which we uncross the pair of separations in σ. Let us say that $\mathcal{F}^*$ is an *uncrossing* of a set $\mathcal{F}$ of sets $\sigma \subseteq \vec{S}$ if

- Every $\tau \in \mathcal{F}^*$ can be obtained by uncrossing a set $\sigma \in \mathcal{F}$;
- For every $\sigma \in \mathcal{F}$ there is some $\tau \in \mathcal{F}^*$ that can be obtained by uncrossing σ.

Note that $\mathcal{F}^*$, like $\mathcal{F}$, is a subset of $2^{\vec{S}}$. Also, $\mathcal{F}^*$ contains all the stars from $\mathcal{F}$, since these have no uncrossings other than themselves. In particular, if $\mathcal{F}$ is standard, i.e. contains all the singleton stars $\{\overleftarrow{r}\}$ with $\vec{r}$ trivial in $\vec{S}$, then so is $\mathcal{F}^*$.

We have shown the following:

Lemma 2.3.4. *$\mathcal{F}$ has an uncrossing $\mathcal{F}^*$. If $\mathcal{F}$ is standard, then so is $\mathcal{F}^*$.*

The smaller we can take $\mathcal{F}^*$ to be, the smaller will be the class of S-trees over $\mathcal{F}^*$. However, to make $\mathcal{F}^*$ as small as possible we would have to give it exactly one star τ for each $\sigma \in \mathcal{F}$, which would involve making a non-canonical choice with regards to the order in which we uncross σ, and possibly which of the two potential uncrossings of a given pair of separations we select.

If we wish for a more canonical choice of family, we can take $\mathcal{F}^*$ to consist of every star that can be obtained by uncrossing a set in $\mathcal{F}$ in any order. Obviously, this will come at the expense of increasing the class of S-trees over $\mathcal{F}$, i.e. the class of dual objects in our desired duality theorem.

Lemma 2.3.5. *Let $\mathcal{F}^*$ be an uncrossing of $\mathcal{F}$. Then an orientation O of S is a regular $\mathcal{F}$-tangle if and only if it is a regular $\mathcal{F}^*$-tangle.*

Proof. Let us first show that if O is a regular $\mathcal{F}^*$-tangle then it is a regular $\mathcal{F}$-tangle. It is sufficient to show that O avoids $\mathcal{F}$. Suppose for a contradiction that there is some $\sigma = \{\vec{x_1}, \vec{x_2}, \ldots, \vec{x_n}\} \in \mathcal{F}$ such that $\sigma \subseteq O$. Since $\mathcal{F}^*$ is an uncrossing of $\mathcal{F}$ there is some $\tau = \{\vec{u_1}, \vec{u_2}, \ldots, \vec{u_n}\} \in \mathcal{F}^*$ that is an uncrossing of σ. Then, $\vec{u_i} \leqslant \vec{x_i} \in O$ for all i. Since O is strongly consistent, by Lemma 2.3.1, this implies $\vec{u_i} \in O$ for each i. Therefore $\tau \subseteq O$, contradicting the fact that O avoids $\mathcal{F}^*$.

Conversely suppose O is a regular $\mathcal{F}$-tangle. We would like to show that O avoids $\mathcal{F}^*$. To do so, we will show that, if O avoids some set $\sigma = \{\vec{x_1}, \vec{x_2}, \ldots, \vec{x_n}\}$ then it also avoids the set

$\sigma' = \{\vec{x_1} \wedge \overleftarrow{x_2}, \vec{x_2} \dots, \vec{x_n}\}$ obtained by uncrossing the pair $\vec{x_1}, \vec{x_2}$. Then by induction O must also avoid every star obtained by uncrossing a set in $\mathcal{F}$, and thus will avoid $\mathcal{F}^*$.

Suppose then that O avoids σ but $\sigma' \subseteq O$. Since $x_1 \in S$ either $\vec{x_1}$ or $\overleftarrow{x_1}$ lies in O. As $\sigma \setminus \{\vec{x_1}\} = \{\vec{x_2}, \vec{x_3}, \dots, \vec{x_n}\} \subseteq \sigma' \subseteq O$, but O avoids σ, we have $\vec{x_1} \notin O$ and hence $\overleftarrow{x_1} \in O$. But then O contains the triple $\{\overleftarrow{x_1}, \vec{x_2}, \vec{x_1} \wedge \overleftarrow{x_2}\} \in \mathcal{P}_S \subseteq \mathcal{F}$. This contradicts the fact that O avoids $\mathcal{F}$. □

Lemma 2.3.5 has an interesting corollary. Suppose, that, in a graph, every star of separations in some given consistent orientation O of S_k points to some k-block. Is there *one* k-block to which all these stars – and hence every separation in O – point? This is indeed the case:

Corollary 2.3.6. *If every star of separations in some strongly consistent orientation of S is contained in some profile of S, then there exists one profile of S that contains all these stars.*

Proof. In Lemma 2.3.5, take $\mathcal{F} := \mathcal{P}_S$. An orientation O of S whose stars each lie in a profile of S cannot contain a star from $\mathcal{P}_S^*$. But if O is regular, consistent, and avoids $\mathcal{P}_S^*$, then by Lemma 2.3.5 it also avoids $\mathcal{P}_S$ and hence is a $\mathcal{P}_S$-tangle. □

Before we can apply Theorem 2.2.1 to our newly found set $\mathcal{F}^*$ of stars, we have to overcome another problem: $\vec{S}$ may fail to be $\mathcal{F}^*$-separable. To address this problem, let us briefly recall what it means for a family to be closed under shifting. Suppose we have a a pair of separations $\vec{r} \leqslant \vec{s}$ such that $\vec{r}$ is nontrivial, nondegenerate, and not forced by $\mathcal{F}$. Suppose further that $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$, and that we have a star $\tau = \{\vec{x_1}, \vec{x_2}, \dots, \vec{x_n}\} \subseteq \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$ that contains some separation $\vec{x_1} \geqslant \vec{r}$. Then the image τ' of τ under $f\downarrow^{\vec{r}}_{\vec{s}}$ is

$$\tau' = \{\vec{x_1} \vee \vec{s}, \vec{x_2} \wedge \overleftarrow{s}, \dots, \vec{x_n} \wedge \overleftarrow{s}\},$$

where the fact that $\vec{s}$ emulates $\vec{r}$ guarantees that $\tau' \subseteq \vec{S}$. Let us call τ' a *shift of* τ, and more specifically the *shift of τ from $\vec{r}$ to $\vec{s}$*. (See [51] for why this is well defined.)

For a family $\mathcal{F}$ to be closed under shifting it is sufficient that it contains all shifts of its elements: that for every $\tau \in \mathcal{F}$, every $\vec{x_1} \in \tau$, every nontrivial and nondegenerate $\vec{r} \leqslant \vec{x_1}$ not forced by $\mathcal{F}$, and every $\vec{s}$ emulating $\vec{r}$, the shift of τ from $\vec{r}$ to $\vec{s}$ is in $\mathcal{F}$.

The idea for making Theorem 2.2.1 applicable to $\mathcal{F}^*$ will be to close $\mathcal{F}^*$ by adding any missing shifts. Let us define a family $\hat{\mathcal{F}}^*$ as follows: Let $\mathcal{G}_0 = \mathcal{F}^*$, define $\mathcal{G}_{n+1}$ inductively as the set of shifts of elements of $\mathcal{G}_n$, and put $\hat{\mathcal{F}}^* := \bigcup_n \mathcal{G}_n$. Clearly $\hat{\mathcal{F}}^*$ is closed under shifting.

Next, let us show that a strongly consistent orientation of S avoids $\mathcal{F}^*$ if and only if it avoids $\hat{\mathcal{F}}^*$. We first note the following lemma.

Lemma 2.3.7. *Let O be a regular $\mathcal{P}$-tangle of S. Let $\sigma \subseteq \vec{S}$ be a star, and let σ' be a shift of σ from some $\vec{r}$ to some $\vec{s} \in \vec{S}$. Then $\sigma' \subseteq O$ implies that $\sigma \subseteq O$.*

Proof. Let $\sigma = \{\vec{x_1}, \vec{x_2}, \dots, \vec{x_n}\}$, with $\vec{r} \leqslant \vec{x_1}$. Then $\vec{x_1} \vee \vec{s} \in \sigma' \subseteq O$. Since O is strongly consistent, this implies that $\vec{x_1}$ and $\vec{s}$ lie in O. Also, for any $i \geqslant 2$, as $x_i \in S$, either $\vec{x_i}$ or $\overleftarrow{x_i}$ lies in O. However, since $\vec{s} \in O$ and $\vec{x_i} \wedge \overleftarrow{s} \in \sigma' \subseteq O$, and O avoids $\mathcal{P}$, it cannot be the case that $\overleftarrow{x_i} \in O$. Hence $\vec{x_i} \in O$ for all $i \geqslant 2$, and so $\sigma \subseteq O$. □

Lemma 2.3.8. *Let $\mathcal{F}^*$ be an uncrossing of $\mathcal{F}$. Then an orientation O of S is a regular $\hat{\mathcal{F}}^*$-tangle if and only if it is a regular $\mathcal{F}$-tangle.*

Proof. Recall that O is a regular $\mathcal{F}$-tangle if and only if it is a regular $\mathcal{F}^*$ tangle (Lemma 2.3.5). Clearly every regular $\hat{\mathcal{F}}^*$-tangle also avoids $\mathcal{F}^*$, and hence is also a regular $\mathcal{F}^*$-tangle, and $\mathcal{F}$-tangle.

Conversely, every regular $\mathcal{F}$-tangle avoids both $\mathcal{F}^* = \mathcal{G}_0$ (Lemma 2.3.5) and hence, by Lemma 2.3.7 and $\mathcal{F} \supseteq \mathcal{P}_S$, also $\mathcal{G}_1$. Proceeding inductively we see that O avoids $\mathcal{G}_n$ for each n, and so avoids $\bigcup_n \mathcal{G}_n = \hat{\mathcal{F}}^*$. Hence O is a regular $\hat{\mathcal{F}}^*$-tangle. □

Before we can, at last, apply Theorem 2.2.1 to our set $\hat{\mathcal{F}}^*$, we have to make one final adjustment: Theorem 2.2.1 requires its set $\mathcal{F}$ of stars to be standard, i.e., to contain all singletons stars $\{\vec{r}\}$ with $\overleftarrow{r}$ trivial in $\vec{S}$. Since trivial separations are small, it will suffice to add to $\hat{\mathcal{F}}^*$ all singleton stars $\{\vec{x}\}$ such that $\overleftarrow{x}$ is small; we denote the resulting superset of $\hat{\mathcal{F}}^*$ by $\overline{\mathcal{F}^*}$. Then $\overline{\mathcal{F}^*}$-tangles contain all small separations, so they are precisely the regular $\hat{\mathcal{F}}^*$-tangles.

Clearly, $\overline{\mathcal{F}^*}$ is a standard set of stars. Also, the shift of any singleton star $\{\vec{x}\}$ is again a singleton star $\{\vec{y}\}$ such that $\vec{x} \leqslant \vec{y}$. Moreover, if $\overleftarrow{x}$ is small then so is $\overleftarrow{y} \leqslant \overleftarrow{x}$, so if $\{\vec{x}\}$ lies in $\overline{\mathcal{F}^*}$ then so does $\{\vec{y}\}$. Therefore, since $\hat{\mathcal{F}}^*$ is closed under shifting, $\overline{\mathcal{F}^*}$ too is closed under shifting. Hence, we get the following duality theorem for abstract profiles:

Theorem 2.3.9. *Let $\vec{S}$ be a submodular separation system in some universe of separations, let $\mathcal{F} \subseteq 2^{\vec{S}}$ contain $\mathcal{P}_S$, and let $\mathcal{F}^*$ be any uncrossing of $\mathcal{F}$. Then the following are equivalent:*

- *There is no regular $\mathcal{F}$-tangle of S;*
- *There is no $\overline{\mathcal{F}^*}$-tangle of S;*
- *There is an S-tree over $\overline{\mathcal{F}^*}$.*

Proof. By Lemmas 2.3.5 to 2.3.8 the regular $\mathcal{F}$-tangles of S are precisely its regular $\hat{\mathcal{F}}^*$-tangles, and by Lemma 2.3.1 these are precisely its $\overline{\mathcal{F}^*}$-tangles. Hence the first two statements are equivalent. Since $\overline{\mathcal{F}^*}$ is a standard set of stars which is closed under shifting, Theorem 2.2.1 implies that the second two statements are equivalent. $\square$

2.4 Duality for special tangles, blocks, and profiles

Let us now apply Theorem 2.3.9 to prove Theorems 2.1.1–2.1.3 from the Introduction. We shall first state the latter two results by specifying $\mathcal{F}$, and then deduce their tree-decomposition formulations as in Theorems 2.1.2 and 2.1.3, along with Theorem 2.1.1.

Theorem 2.4.1. *For every finite graph G and every integer $k > 0$ exactly one of the following statements holds:*

- *G contains a k-block;*
- *G has an S_k-tree over $\overline{\mathcal{B}_k^*}$, where $\mathcal{B}_k^*$ is any uncrossing of $\mathcal{B}_k$.*

Proof. In Theorem 2.3.9, let $S = S_k$ be the set of separations of order $< k$ in G, and let $\mathcal{F} := \mathcal{B}_k \cap 2^{\vec{S_k}}$. Then $\vec{S_k}$ is a submodular separation system in the universe of all separations of G, and the regular $\mathcal{F}$-tangles in G are precisely the orientations $O(b)$ of S_k for k-blocks b in G. Hence G has a k-block if and only if it has a regular $\mathcal{F}$-tangle for this $\mathcal{F}$. The assertion now follows from Theorem 2.3.9. $\square$

As for profiles, every graph G has a regular 1-profile – just orient every 0-separation towards some fixed component –and a regular 2-profile (Lemma 2.3.3). So we need a duality theorem only for $k > 2$. Recall that, for $k > 2$, all k-profiles of graphs are regular (Lemma 2.3.2).

Theorem 2.4.2. *For every finite graph G and every integer $k > 2$ exactly one of the following statements holds:*

- *G has a k-profile;*
- *G has an S_k-tree over $\overline{\mathcal{P}_k^*}$, where $\mathcal{P}_k^*$ is any uncrossing of $\mathcal{P}_k$.*

Proof. In Theorem 2.3.9, let $S = S_k$ be the set of separations of order $< k$ in G, and let $\mathcal{F} := \mathcal{P}_k$. Then $\vec{S_k}$ is a submodular separation system in the universe of all separations of G, and the regular $\mathcal{F}$-tangles in G are precisely its k-profiles. The assertion now follows from Theorem 2.3.9, as before. □

Theorems 2.1.1–2.1.3 now follow easily: we just have to translate S-trees into tree-decompositions.

Proof of Theorems 2.1.1–2.1.3. Given a set S of separations of G and an S-tree (T,α), with $\alpha(\vec{e}) =: (A_\alpha(\vec{e}), B_\alpha(\vec{e}))$ say, we obtain a tree-decomposition $(T,\mathcal{V}_\alpha)$ of G with $\mathcal{V}_\alpha = (V_t)_{t\in T}$ by letting

$$V_t := \bigcap \{\, B_\alpha(\vec{e}) \mid \vec{e} = (s,t) \in \overrightarrow{E(T)} \,\}.$$

Note that (T,α) can be recovered from this tree-decomposition: given just T and $\mathcal{V} = (V_t)_{t\in T}$, we let α map each oriented edge $\vec{e} = (t_1,t_2)$ of T to the oriented separation of G it induces: the separation $\big(\bigcup_{t\in T_1} V_t\,,\, \bigcup_{t\in T_2} V_t\big)$ where T_i is the component of $T - e$ containing t_i.

Recall that the set $\mathcal{T}$ defined in (2.2.2) contains $\mathcal{P}$. Hence so does any $\mathcal{F} \supseteq \mathcal{T}$. For such $\mathcal{F}$, therefore, every $\mathcal{F}$-tangle of S_k is a k-profile, and hence is regular by Lemma 2.3.2 if $k > 2$. For $\mathcal{F}_k := \mathcal{F} \cap 2^{\vec{S_k}}$ and

$$\mathcal{T}_{\mathcal{F}}(k) := \{\, (T,\mathcal{V}_\alpha) \mid (T,\alpha) \text{ is an } S_k\text{-tree over } \overline{\mathcal{F}_k^*} \,\}$$

we thus obtain Theorem 2.1.1 directly from Theorem 2.3.9. Similarly, letting

$$\begin{aligned} \mathcal{T}_{\mathcal{B}}(k) &:= \{\, (T,\mathcal{V}_\alpha) \mid (T,\alpha) \text{ is an } S_k\text{-tree over } \overline{\mathcal{B}_k^*} \,\} \\ \mathcal{T}_{\mathcal{P}}(k) &:= \{\, (T,\mathcal{V}_\alpha) \mid (T,\alpha) \text{ is an } S_k\text{-tree over } \overline{\mathcal{P}_k^*} \,\} \end{aligned}$$

yields Theorems 2.1.2 and 2.1.3 as corollaries of Theorems 2.4.1 and 2.4.2. □

2.5 Width parameters

In this section we derive some bounds for the block and profile width of a graph that follow easily from our main results combined with those of [51, 52] and [110].

Given a star σ of separations in a graph, let us call the set $\bigcap\{\, B \mid (A,B) \in \sigma\}$ the *interior* of σ in G. For example, every star in $\mathcal{P}_k^*$ is of the form

$$\{(A,B), (B\cap C, A\cup D), (B\cap D, A\cup C)\} \subseteq S_k,$$

and hence every vertex of its interior lies in at least two of the separators $A\cap B$, $(B\cap C)\cap(A\cup D)$ and $(B\cap D)\cap(A\cup C)$. Since all these separations are in S_k, the interior of any star in $\mathcal{P}_k^*$ thus has size at most $3(k-1)/2$.

We can apply this observation to obtain the following upper bound on the profile-width $\mathrm{pw}(G)$ of a graph G in terms of its tree-width $\mathrm{tw}(G)$:

Theorem 2.5.1. *For every graph G,*

$$\mathrm{pw}(G) \leqslant \mathrm{tw}(G) + 1 \leqslant \tfrac{3}{2}\mathrm{pw}(G). \tag{2.5.1}$$

Proof. For the first inequality, note that $\mathrm{tw}(G)+1$ is the largest integer k such that G has no tree-decomposition into parts of order $< k$. By the duality theorem for tree-width from [52], having no such tree-decomposition is equivalent to admitting an $\mathcal{S}^k$-tangle of S_k, where

$$\mathcal{S}^n = \{\, \tau \subseteq \vec{U} : \tau \text{ is a star and } \big|\bigcap\{\, B : (A,B) \in \tau \,\}\big| < n \,\}$$

and $\overrightarrow{U}$ is the universe of all separations of G. But among the $\mathcal{S}^k$-tangles of S_k are all the k-profiles of G. (An easy induction on $|\tau|$ shows that a regular k-profile has no subset $\tau \in \mathcal{S}^k$; cf. Lemma 2.3.2 and [48, Prop. 3.4].) Therefore G has no k-profile for $k > \mathrm{tw}(G) + 1$, which Corollary 2.1.4 translates into $\mathrm{pw}(G) \leqslant \mathrm{tw}(G) + 1$.

For the second inequality, recall that if $k := \mathrm{pw}(G)$ then G has a tree-decomposition in $\mathcal{T}_{\mathcal{P}}(k+1)$. The parts of this tree-decomposition are interiors of stars in $\overline{\mathcal{P}^*_{k+1}}$, so they have size at most $3k/2$.This tree-decomposition, therefore, has width at most $(3k/2) - 1$, which thus is an upper bound for $\mathrm{tw}(G)$. □

We can also relate the profile-width of a graph to its branch-width, as follows. In order to avoid tedious exceptions for small k, let us define the *adjusted branch-width* of a graph G as

$$\mathrm{brw}(G) := \min\{\, k \mid G \text{ has no } S_{k+1}\text{-tree over } \mathcal{T}^*\}.$$

By [52, Theorem 4.4], this is equivalent to the *tangle number* of G, the greatest k such that G has a k-tangle. For $k \geq 3$ it coincides with the original branch-width of G as defined by Robertson and Seymour [110].[4]

Robertson and Seymour [110] showed that the adjusted branch-width of a graph is related to its tree-width in the same way as we found that the profile-width is:

$$\mathrm{brw}(G) \leqslant \mathrm{tw}(G) + 1 \leqslant \tfrac{3}{2}\mathrm{brw}(G). \tag{2.5.2}$$

Together, these inequalities imply the following relationship between branch-width and profile-width:

Corollary 2.5.2. *For every graph G,*

$$\mathrm{brw}(G) \leqslant \mathrm{pw}(G) \leqslant \tfrac{3}{2}\mathrm{brw}(G). \tag{2.5.3}$$

Proof. The first inequality follows from the fact that k-tangles are k-profiles, and that the largest k for which G has a k-tangle equals the adjusted branch-width: thus,

$$\mathrm{brw}(G) = k \leq \pi(G) = \mathrm{pw}(G)$$

by Corollary 2.1.4.

For the second inequality, notice that $\mathrm{pw}(G) \underset{(2.5.1)}{\leqslant} \mathrm{tw}(G) + 1 \underset{(2.5.2)}{\leqslant} \tfrac{3}{2}\mathrm{brw}(G)$. □

Since every k-block defines a k-profile, Corollary 2.1.4 implies that the block-width of a graph is a lower bound for its profile-width, and hence by (2.5.1) also for it tree-width (plus 1). Conversely, however, no function of the tree-width of a graph can be a lower bound for its block-width. Indeed, the tree-width of the $k \times k$-grid H_k is at least k (see [43]) but H_k contains no 5-block: in every set of $\geqslant 5$ vertices there are two non-adjacent vertices, and the neighbourhood of either vertex is then a set of size 4 which separates the two vertices. Therefore there exist graphs with bounded block-width and arbitrarily high tree-width.

Since large enough tree-width forces both large profile-width (2.5.1) and large branch-width (2.5.2), the grid example shows further that making these latter parameters large cannot force the block-width of a graph above 4.

[4]Our *adjusted branch-width* is the dual parameter to the tangle number for all k, while the original branch-width from [110] achieves this only for $k \geq 3$: it deviates from the tangle number for some graphs and $k \leq 2$. See [52, end of Section 4] for a discussion.

Chapter 3

A unified treatment of linked and lean tree-decompositions

3.1 Introduction

Given a tree T and vertices $t_1, t_2 \in V(T)$ let us denote by t_1Tt_2 the unique path in T between t_1 and t_2. Given a graph G a *tree-decomposition* of G is a pair $(T, \mathcal{V})$ consisting of a tree T, together with a collection of subsets of vertices $\mathcal{V} = \{V_t \subseteq V(G) : t \in V(T)\}$, called *bags*, such that:

- $V(G) = \bigcup_{t \in T} V_t$;
- For every edge $e \in E(G)$ there is a t such that e lies in V_t;
- $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever $t_2 \in V(t_1Tt_3)$.

The *width* of this tree-decomposition is the quantity $\max\{|V_t| - 1 : t \in V(T)\}$ and its *adhesion* is $\max\{|V_t \cap V_{t'}| : (t, t') \in E(T)\}$. Given a graph G its *tree-width* $\mathrm{tw}(G)$ is the smallest k such that G has a tree-decomposition of width k.

Definition 3.1.1. *A tree decomposition $(T, \mathcal{V})$ of a graph G is called* linked *if for all $k \in \mathbb{N}$ and every $t, t' \in V(T)$, either G contains k disjoint V_t-$V_{t'}$ paths or there is an $s \in V(tTt')$ such that $|V_s| < k$.*

Robertson and Seymour showed the following:

Theorem 3.1.2 (Robertson and Seymour [113])**.** *Every graph G admits a linked tree-decomposition of width $< 3 \cdot 2^{tw(G)}$.*

This result was an essential part of their proof that for any k the set of graphs with tree-width less than k is well-quasi-ordered by the minor relation. Thomas gave a new proof of Theorem 3.1.2, improving the bound on the tree-width of the linked tree-decomposition from $3 \cdot 2^{\mathrm{tw}(G)} - 1$ to the best possible value of $\mathrm{tw}(G)$.

Theorem 3.1.3 (Thomas [124])**.** *Every graph G admits a linked tree-decomposition of width $tw(G)$.*

In fact he showed a stronger result.

Definition 3.1.4. *A tree decomposition $(T, \mathcal{V})$ of a graph G is called* lean *if for all $k \in \mathbb{N}$, $t, t' \in V(T)$ and vertex sets $Z_1 \subseteq V_t$ and $Z_2 \subseteq V_{t'}$ with $|Z_1| = |Z_2| = k$, either G contains k disjoint $Z_1 - Z_2$ paths or there exists an edge $\{s, s'\} \in E(tTt')$ with $|V_s \cap V_{s'}| < k$.*

Thomas showed that every graph has a lean tree-decomposition of width $\mathrm{tw}(G)$. It is relatively simple to make a lean tree-decomposition linked without increasing its width: we subdivide each edge $(s, s') \in E(T)$ by a new vertex u and add $V_u := V_s \cap V_{s'}$. The real strength of Definition 3.1.4 in comparison to Definition 3.1.1 is the case $t = t'$. Broadly, Definition 3.1.1 tells us that the 'branches' in the tree-decomposition are no larger than the connectivity of the graph requires, if the separators $V_s \cap V_{s'}$ along a path in T are large, then G is highly connected along this branch. The case $t = t'$ of Definition 3.1.4 tells us that the bags are also no larger than their 'external connectivity' in G requires.

Bellenbaum and Diestel [17] used some of the ideas from Thomas' paper to give short proofs of two known results concerning tree-decompositions, the first Theorem 3.1.3 and the other the tree-width duality theorem of Seymour and Thomas [118]. Very similar ideas appear in the literature in multiple proofs of the existence of 'linked' or 'lean' tree-decompositions of minimum width, for many different width parameters. For example θ-tree-width [34, 66], path-width [90] and directed path-width [85]. Similar ideas also appear in the proof of Geelen, Gerards and Whittle [68] that every matroid M admits a 'linked' branch-decomposition of width the branch width of M.

The existence of linked decompositions has been used multiple times in the literature to prove that collections of structures of bounded width are well-quasi-ordered. As mentioned before, the result of Thomas was used by Robertson and Seymour [108] to show that the set of graphs with bounded tree-width is well-quasi-ordered under the minor relation. Geelen, Gerards and Whittle [68] used their aforementioned result to extend this to matroids with bounded branch-width, and Oum [106] used the same result to show that the set of graphs with bounded rank-width is well-quasi-ordered under the vertex-minor relation. Recently similar ideas have also been applied to tournaments (see [38, 85]).

In this paper we will prove generalizations of Theorem 3.1.3 in a general framework introduced by Diestel [44], which give unifying proofs of the existence of 'linked' or 'lean' tree-decompositions for a broad variety of width parameters. In particular these theorems will imply all the known results for undirected graphs and matroids from the introduction, as well as many new results by applying them to other width parameters expressible in this framework. In particular we prove a generalization of Theorem 3.1.3 to matroids[1]:

Theorem 3.5.14. *Every matroid M admits a lean tree-decomposition of width $tw(M)$.*

In order to prove a result broad enough to cover both graph and matroid tree-width, it will be necessary to combine some of the ideas from the proof of Bellenbaum and Diestel [17] with that of Geelen, Gerards and Whittle [68]. In particular, we note that, when interpreted in terms of tree-decompositions of graphs, our proof will give a slightly different proof of Theorem 3.1.3 than appears in [17].

In the next section we will introduce the necessary background material. In Sections 3.3 and 3.4 respectively we will give proofs of our theorems on the existence of linked and lean tree-decompositions. Finally, in Section 3.5 we will use these theorems to deduce results about different width parameters of graphs and matroids.

3.2 Background material

3.2.1 Notation

Recall that, given a tree T and vertices $t_1, t_2 \in V(T)$, $t_1 T t_2$ is the unique path in T between t_1 and t_2. Similarly, given edges $e_1, e_2 \in E(T)$ we will denote by $e_1 T e_2$ the unique path in T

[1]See Section 3.5.2 for the relevant definitions.

which starts at e_1 and ends at e_2. If $\vec{e}$ is a oriented edge in T, say $\vec{e} = (x, y)$, then we will write $\overleftarrow{e} = (y, x)$, and we denote by $T(\vec{e})$ the subtree of T formed by taking the component C of $T - x$ which contains y, and adding the edge (x, y). We will write $\vec{E}(T)$ for the set $\{(x, y) : \{x, y\} \in E(T)\}$ of oriented edges of T.

We will write $\mathbb{N}$ to mean the set of natural numbers, which for us will include 0. For general graph theoretic notation we will follow [43].

3.2.2 Separation systems

Our objects of study will be the separation systems of Diestel and Oum [51]. A more detailed introduction to these structures can be found in [44].

A *separation system* $(\vec{S}, \leqslant, *)$ as defined in [44] consists of a partially ordered set $\vec{S}$, whose elements are called *oriented separations*, together with an order reversing involution $*$. One particular example of a separation system is the following: Given a set V, the set of ordered pairs

$$\mathrm{sep}(V) = \{(A, B) : A \cup B = V\}$$

together with the ordering

$$(A, B) \leqslant (C, D) \text{ if and only if } A \subseteq C \text{ and } B \supseteq D,$$

and the involution $(A, B)^* = (B, A)$ forms a separation system. Furthermore, any subset of $\mathrm{sep}(V)$ which is closed under involutions forms a separation system, and we call such a separation system a *set separation system* (see [52]). We will use V to denote the ground set of a set separation system unless otherwise specified. We will work with set separation systems rather than in full generality since the existence of an underlying ground set will be necessary to define our concept of 'leaness'. Some of the results in the paper, in particular Theorem 3.3.3, will carry over to the more abstract setting, but only with so many technical restrictions that little, if anything, is lost by this restriction.

The elements (A, B) of $\vec{S}$ are called *separations* and the corresponding *separation* is $\{A, B\}$. (A, B) and (B, A) are then the *orientations* of $\{A, B\}$[2]. The separation $\{V, V\}$ is called *degenerate*.

Two separations $\{A, B\}$ and $\{C, D\}$ are *nested* if they have $\leqslant$-comparable orientations. A set of separations is *nested* if its elements are pairwise nested[3]. A (multi)-set $\{(A_i, B_i) : i \in I\}$ of non-degenerate separations is a *multistar of separations* if $(A_i, B_i) \leqslant (B_j, A_j)$ for all $i, j \in I$.

There are binary operations on $\mathrm{sep}(V)$ such that $(A, B) \wedge (C, D)$ is the infimum and $(A, B) \vee (C, D)$ is the supremum of (A, B) and (C, D) in $\mathrm{sep}(V)$ given by:

$$(A, B) \wedge (C, D) = (A \cap C, B \cup D) \text{ and } (A, B) \vee (C, D) = (A \cup C, B \cap D).$$

If a set separation system $\vec{S}$ is closed under $\wedge$ and $\vee$ we say it is a *universe* of set separations. The following lemma, whose elementary proof we omit, is often useful.

Lemma 3.2.1. *If (A, B) is nested with (C, D) and (E, F) then it is also nested with $(C, D) \vee (E, F)$, $(C, D) \wedge (E, F)$, $(C, D) \vee (F, E)$ and $(C, D) \wedge (F, E)$.*

We call a function $(A, B) \mapsto |A, B|$ on a universe of set separations an *order function* if it is non-negative, symmetric and submodular. That is, if $0 \leqslant |A, B| = |B, A|$ for all $(A, B) \in \vec{S}$ and

$$|A \cap C, B \cup D| + |A \cup C, B \cap D| \leqslant |A, B| + |C, D|$$

[2]When the context is clear we will often refer to both oriented and unoriented separations as merely 'separations' to improve the flow of the text.

[3]We will often use terms defined for separations in reference to its orientations, and vice versa.

for all $(A, B), (C, D) \in \vec{S}$. Given a universe of set separations $\vec{S}$ with any order function we will often write $\vec{S}_k$ for the separation system

$$\vec{S}_k := \{(A, B) \in \vec{S} \,:\, |A, B| < k\}.$$

If $r : 2^V \to \mathbb{N}$ is a non-negative submodular function then it is easy to verify that

$$|X, Y|_r = r(X) + r(Y) - r(V)$$

is an order function on any universe contained in $\mathrm{sep}(V)$. We note the following lemma, which will be useful later.

Lemma 3.2.2. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|_r$ for some non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$. If $(B, A) \leqslant (C, D)$, then $|A \cup C, B \cap D|_r \leqslant |C, D|_r$.*

Proof. Note that $A \cup C = V$ and $B \cap D \subset C \cap D$. Hence,

$$\begin{aligned} |A \cup C, B \cap D|_r &= r(A \cup C) + r(B \cap D) - r(V) \\ &= r(B \cap D) \leqslant r(C \cap D) \\ &\leqslant r(C) + r(D) - r(C \cup D) = |C, D|_r. \end{aligned}$$

□

Our main application of separation systems will be to separations of graphs and matroids. If G is a graph, we say that an ordered pair (A, B) of subsets of $V(G)$ is an oriented separation if $A \cup B = V(G)$ and there is no edge between $A \setminus B$ and $B \setminus A$. It is not hard to show that the set of oriented separations $\vec{S}$ of a graph G forms a universe of set separations as a subset of $\mathrm{sep}(V(G))$, and if we consider the non-negative non-decreasing submodular function $r : 2^{V(G)} \to \mathbb{N}$ given by $r(A) := |A|$, then the order function given by r is $|X, Y|_r = |X| + |Y| - |V| = |X \cap Y|$. Throughout this paper, whenever we consider separation systems of graph separations we will use this order function.

Similarly, if $M = (E, \mathcal{I})$ is a matroid with rank function r, we say that an ordered bipartition $(A, E \setminus A)$ of the ground set is an oriented separation. Note that $r : 2^E \to \mathbb{N}$ is a non-negative non-decreasing submodular function. Again, the set of oriented separations $\vec{S}$ of a matroid M forms a universe of set separations, with an order function given by $|X, Y|_r = r(X)+r(Y)-r(E)$. Again, whenever we consider separations systems of matroid separations we will use this order function.

3.2.3 S-trees over multistars

Given a tree T and $t \in V(T)$ we write

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

Definition 3.2.3. *Let $\vec{S}$ be a set separation system. An S-*tree *is a pair (T, α) where T is a finite tree and*

$$\alpha : \vec{E}(T) \to \vec{S} \setminus \{(V, V)\}$$

is a map from the set of oriented edges of T to $\vec{S} \setminus \{(V, V)\}$ such that, for each edge $e \in E(T)$, if $\alpha(\vec{e}) = (A, B)$ then $\alpha(\overleftarrow{e}) = (B, A)$.

Given a finite set X we will write $\mathbb{N}^X$ for the set of all finite submultisets of X. If $\mathcal{F} \subseteq \mathbb{N}^{\vec{S}}$ is a family of multistars we say that an S-tree (T, α) is *over* $\mathcal{F}$ if for all $t \in V(T)$, the multiset $\sigma_t := \{\alpha(\vec{e}) : \vec{e} \in \vec{F}_t)\} \in \mathcal{F}$. We will often write $\alpha(\vec{F}_t)$ to denote this multiset. Sometimes, for notational convenience, we will refer to an S-tree over $\mathcal{F}$ when $\mathcal{F} \not\subseteq \mathbb{N}^{\vec{S}}$, by which it should be taken to mean an S-tree over $\mathcal{F} \cap \mathbb{N}^{\vec{S}}$. If (T, α) is an S-tree over a family of multistars, it is easy to check that α preserves the natural ordering on $\vec{E}(T)$.

It is observed in [52] that many existing ways of decomposing graphs can be expressed in this framework if we take $\vec{S}$ to be the universe of separations of a graph. For example given a multistar $\sigma = \{(A_1, B_1), (A_2, B_2), \dots, (A_n, B_n)\}$ let us write $\langle \sigma \rangle := |\bigcap_i B_i|$. If we let

$$\mathcal{F}_k = \{\sigma \in \mathbb{N}^{\vec{S}} : \sigma \text{ a multistar}, \langle \sigma \rangle < k\}$$

then it is shown in [52] that a graph G has a tree-decomposition of width $< k - 1$ if and only if there exists an S_k-tree over $\mathcal{F}_k$. More examples will be given later in Section 3.5 when we apply Theorems 3.3.3 and 3.4.1 to existing types of tree-decompositions.

Given two separations $(A, B) \leqslant (C, D)$ in a universe of set separations $\vec{S}$ with an order function $|\cdot|$ let us define

$$\lambda\left((A, B), (C, D)\right) = \min\{|X, Y| : (A, B) \leqslant (X, Y) \leqslant (C, D)\}.$$

We will think of λ as being a measure of connectivity, and indeed in the case of separation systems of graphs and matroids it will coincide with the normal connectivity function.

Definition 3.2.4. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|$, $k \in \mathbb{N}$. We say that an S_k-tree (T, α) over some family of multistars is* linked *if for every pair of edges $\vec{e} \leqslant \vec{f}$ in $\vec{E}(T)$ the following holds:*

$$\min\{|\alpha(\vec{g})| : \vec{g} \in \vec{E}(eTf)\} = \lambda\left(\alpha(\vec{e}), \alpha(\vec{f})\right).$$

This definition more closely resembles the definition of linked from Geelen, Gerards and Whittle [68] than Definition 3.1.1. However, as in the introduction, if we take a tree-decomposition of a graph which is linked in this sense and add the adhesion sets as parts, it will be linked in the sense of Thomas. For any family of multistars $\mathcal{F}$ we also introduce the following concept of leanness.

Definition 3.2.5. *Let $\vec{S}$ be a set separation system and let (T, α) be an S-tree over some family of multistars $\mathcal{F}$. Given a vertex $t \in T$ we say that a separation (A, B) is* addable *at t if $\sigma_t \cup \{(A, B)\} \in \mathcal{F}$.*

Definition 3.2.6. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|_r$ for some non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$, $k \in \mathbb{N}$. We say that an S_k-tree (T, α) over a family of multistars $\mathcal{F} \subseteq \mathbb{N}^{\vec{S}_k}$ is* $\mathcal{F}$-lean *if for every pair of vertices t and t' in T and every pair of separations $(A, B) \leqslant (B', A')$ such that (A, B) is addable at t and (A', B') is addable at t', either*

$$\lambda\left((A, B), (B', A')\right) \geqslant \min\{r(A), r(A')\}$$

or

$$\min\{|\alpha(\vec{g})| : \vec{g} \in \vec{E}(tTt')\} = \lambda\left((A, B), (B', A')\right).$$

We note that, unlike in the case of Definitions 3.1.1 and 3.1.4, it is not true in general that the existence of an $\mathcal{F}$-lean S_k-tree over $\mathcal{F}$ will imply the existence of a linked S_k-tree over $\mathcal{F}$. Indeed, for many families $\mathcal{F}$ of multistars, the set of separations addable at any multistar may be very limited, if not empty. In this case the property of being $\mathcal{F}$-lean may be vacuously true of any S_k-tree over $\mathcal{F}$.

3.2.4 Shifting S-trees

Let $\vec{S}$ be a universe of set separations. Our main tool will be a method for taking an S-tree (T,α) and a separation (X,Y) and building another S-tree (T',α') which contains (X,Y) as the image of some edge.

Given an S-tree over some family of multistars (T,α), a separation $(A,B) = \alpha(\vec{e}) \in \alpha(\vec{E}(T))$, and another, non-trivial separation $(A,B) \leqslant (X,Y)$ the *shift of* (T,α) *onto* (X,Y) *(with respect to* $\vec{e}$*)* is the S-tree $(T(\vec{e}),\alpha')$ where α' is defined as follows:

For every edge $f \in E(T(\vec{e}))$ there is a unique orientation of f such that $\vec{e} \leqslant \vec{f}$. We define

$$\alpha'(\vec{f}) := \alpha(\vec{f}) \vee (X,Y) \text{ and } \alpha'(\overleftarrow{f}) := \alpha'(\vec{f})^* = \alpha(\overleftarrow{f}) \wedge (Y,X).$$

Note that, since $(A,B) \leqslant (X,Y)$ we have that $\alpha'(\vec{e}) = \alpha(\vec{e}) \vee (X,Y) = (A,B) \vee (X,Y) = (X,Y)$, and so $(X,Y) \in \alpha(\vec{E}(T(\vec{e})))$.

Our tool will be in essence [Lemma 4.2, [51]], however, for technical reasons we will take a slightly different statement. Explicitly, for the results in [51] it is only ever necessary to use the shifting operation when $\alpha(\vec{e})$ is a non-trivial separation, whereas we may need to do so when $\alpha(\vec{e})$ is trivial. This allows the authors of [51] to assume that the S-trees they consider are of a particularly simple kind, which we are not able to assume.

Lemma 3.2.7 (Lemma 4.1, [51])**.** *Let S be a universe of set separations, (T,α) an S-tree over some family of multistars, $\vec{e} \in \vec{E}(T)$ and $(X,Y) \geqslant \alpha(\vec{e})$ be a non-trivial separation. Then the shift of (T,α) onto (X,Y) with respect to $\vec{e}$ is an S-tree over some family of multistars.*

In general, if $\vec{S}$ is not a universe but just a set separation system, it may not be true that the shift of an S-tree is still an S-tree. However, when $\vec{S} = \vec{S}_k$ for some universe there is a natural condition on (X,Y) which guarantees that the shift will still be an S_k-tree.

Definition 3.2.8. *Let $\vec{S}$ be a set separation system living in some universe with an order function $|\cdot|$ and let $(A,B),(X,Y) \in \vec{S}$. We say that (X,Y) is* linked *to (A,B) if $(A,B) \leqslant (X,Y)$ and*

$$|X,Y| = \lambda\left((A,B),(X,Y)\right).$$

Lemma 3.2.9. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|$, $k \in \mathbb{N}$, (T,α) an S_k-tree over some family of multistars and $(A,B) = \alpha(\vec{e}) \in \alpha(\vec{E}(T))$. Let (X,Y) be a non-trivial separation that is linked to (A,B) and let $(T(\vec{e}),\alpha')$ be the shift of (T,α) onto (X,Y) with respect to $\vec{e}$. Then $(T(\vec{e}),\alpha')$ is an S_k-tree over some family of multistars.*

Proof. By Lemma 3.2.7 the shift is an S-tree over some family of multistars, so it will be sufficient to show that $\alpha'(\vec{E}(T(\vec{e}))) \subseteq \vec{S}_k$. Suppose that $(C,D) = \alpha'(\vec{f}) \in \alpha'(\vec{E}(T(\vec{e})))$. By symmetry we may assume that $\vec{e} \leqslant \vec{f}$, and so $(C,D) = \alpha(\vec{f}) \vee (X,Y)$ by definition.

Since $\vec{e} \leqslant \vec{f}$, $\alpha(\vec{e}) = (A,B) \leqslant \alpha(\vec{f})$ and also $(A,B) \leqslant (X,Y)$. Hence,

$$(A,B) \leqslant \alpha(\vec{f}) \wedge (X,Y) \leqslant (X,Y).$$

Therefore, since (X,Y) is linked to (A,B), it follows that

$$|\alpha(\vec{f}) \wedge (X,Y)| \geqslant |X,Y|$$

and so by the submodularity of the order function

$$|C,D| = |\alpha(\vec{f}) \vee (X,Y)| \leqslant |\alpha(\vec{f})| < k,$$

and so $(C,D) \in \vec{S}_k$. $\square$

Finally we will need a condition that guarantees that the shift of an S_k-tree over $\mathcal{F}$ is still over $\mathcal{F}$.

Definition 3.2.10. *Let $\vec{S}$ be a set separation system living in some universe with an order function $|\cdot|$ and let $\mathcal{F} \subseteq \mathbb{N}^{\vec{S}}$ be a family of multistars. We say that $\mathcal{F}$ is* fixed under shifting *if whenever $\sigma \in \mathcal{F}$, (T, α) is an S-tree over some family of multistars with $\vec{e} = (t, t') \in \vec{E}(T)$, $\sigma = \alpha(\vec{F}_s)$ for some $s \in V(T(\vec{e})) \setminus \{t\}$, and (X, Y) is a non-trivial separation that is linked to $\alpha(\vec{e})$, then in the shift of (T, α) onto (X, Y) with respect to $\vec{e}$, $\sigma' = \alpha'(\vec{F}_s) \in \mathcal{F}$.*

It is essentially this property of a family of multistars $\mathcal{F}$ that is used in [52] to prove a duality theorem for S_k-trees over $\mathcal{F}$. It may seem like a strong property to hold, but in fact it is seen to hold in a large number of cases corresponding to known width-parameters of graphs such as branch-width, tree-width (see Corollary 3.2.16), and path-width.

Lemma 3.2.11. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|$, $k \in \mathbb{N}$ and let $\mathcal{F} \subset \mathbb{N}^{\vec{S}}$ be a family of multistars which is fixed under shifting. Let (T, α) be an S_k-tree over some family of multistars and $\vec{e} = (t, t') \in \vec{E}(T)$ be such that*

$$\text{for every } s \in V(T(\vec{e})) \setminus \{t\},\ \alpha\left(\vec{F}_s\right) \in \mathcal{F}.$$

Let (X, Y) be a non-trivial separation that is linked to $\alpha(\vec{e})$ and let $(T(\vec{e}), \alpha')$ be the shift of (T, α) onto (X, Y) with respect to $\vec{e}$. Then $(T(\vec{e}), \alpha')$ is an S_k-tree over $\mathcal{F} \cup \{(Y, X)\}$. Furthermore, if $\{(Y, X)\} \notin \mathcal{F}$ then t is the unique vertex of $T(\vec{e})$ such that $\alpha'\left(\vec{F}_t\right) = \{(Y, X)\}$.

Proof. By Lemma 3.2.9 the shift is an S_k-tree, it remains to show that it is over $\mathcal{F} \cup \{(Y, X)\}$. Firstly, we note that by definition $\alpha'(\vec{e}) = (X, Y)$ and so $\alpha'(\vec{F}_t) = \{\alpha'(\overleftarrow{e})\} = \{(Y, X)\} \in \mathcal{F} \cup \{(Y, X)\}$.

Suppose then that $s \in V(T(\vec{e})) \setminus \{t\}$. By assumption, $\alpha\left(\vec{F}_s\right) \in \mathcal{F}$ and so, since $\mathcal{F}$ is fixed under shifting, $\alpha'(\vec{F}_s) \in \mathcal{F}$. □

Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|_r$ for some non-decreasing submodular function $r : 2^V \to \mathbb{N}$. For any multistar

$$\sigma = \{(A_0, B_0), (A_1, B_1), \ldots, (A_n, B_n)\}$$

let us define the *size* of σ

$$\langle\sigma\rangle_r = \sum_{i=0}^{n} r(B_i) - n \cdot r(V).$$

Note that, when $r(\cdot) = |\cdot|$ is the cardinality function then $\langle\sigma\rangle_r = |\bigcap_{i=0}^{n} B_i|$. We define

$$\mathcal{F}_p = \{\sigma \subset \mathbb{N}^{\vec{S}} : \sigma \text{ a multistar}, \langle\sigma\rangle_r < p\}.$$

We will collect a few properties of $\langle\cdot\rangle_r$ that we will need. The following lemma will be useful for these proofs and later.

Lemma 3.2.12. *Let $r : 2^V \to \mathbb{N}$ be a non-negative non-decreasing submodular function, $X \subseteq V$ and let $\{(A_0, B_0), (A_1, B_1), \ldots, (A_n, B_n)\}$ be a multistar of separations in sep(V). If we write $B_i^* = \bigcap_{j=0}^{i} B_j$, then*

$$\sum_{i=0}^{n} r(B_i \cap X) \geqslant r(B_n^* \cap X) + n \cdot r(X).$$

Proof. Note that, since $A_j \subseteq B_i$ for each $i \neq j$ it follows that $B_i^* \cup B_{i+1} \supseteq A_{i+1} \cup B_{i+1} = V$ for each i. Hence, by submodularity

$$r(B_i^* \cap X) + r(B_{i+1} \cap X) \geqslant r(B_{i+1}^* \cap X) + r(X).$$

Since $B_0^* = B_0$, if we add these inequalities for $i = 0, \dots, n-1$, we get

$$\sum_{i=0}^{n} r(B_i \cap X) \geqslant r(\bigcap_{i=0}^{n} B_i \cap X) + n \cdot r(X).$$

□

Lemma 3.2.13. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|_r$ for some non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$. If σ is a multistar and $(A, B) \in \sigma$ then $\langle \sigma \rangle_r \geqslant |A, B|_r$.*

Proof. Let us write $\sigma = \{(A, B), (A_1, B_1), \dots, (A_n, B_n)\}$. Since σ is a multistar, $(A, B) \leqslant (B_i, A_i)$ for all i and so $A \subset \bigcap_{i=1}^{n} B_i$. Hence $r(A) \leqslant r(\bigcap_{i=1}^{n} B_i)$. Therefore, applying Lemma 3.2.12 to the multistar $\sigma \setminus \{(A, B)\}$ with $X = V$, we see that

$$\begin{aligned}
\langle \sigma \rangle_r &= r(B) + \sum_{i=1}^{n} r(B_i) - n \cdot r(V) \\
&\geqslant r(B) + r(\bigcap_{i=1}^{n} B_i) - r(V) \\
&\geqslant r(B) + r(A) - r(V) \\
&= |A, B|_r
\end{aligned}$$

□

Lemma 3.2.14. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|_r$ for some non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$. If σ is a multistar and $\sigma \cup \{(A, B)\}$ is a multistar then $\langle \sigma \rangle_r \geqslant r(A)$.*

Proof. Let us write $\sigma = \{(A_0, B_0), (A_1, B_1), \dots, (A_n, B_n)\}$. If $\sigma \cup \{(A, B)\}$ is a multistar then $(A, B) \leqslant (B_i, A_i)$ and so $A \subset B_i$ for all i. Hence, by Lemma 3.2.12 applied to the multistar σ with $X = V$,

$$\begin{aligned}
\langle \sigma \rangle_r &= \sum_{i=0}^{n} r(B_i) - n \cdot r(V) \\
&\geqslant r(\bigcap_{i=0}^{n} B_i) \\
&\geqslant r(A).
\end{aligned}$$

□

The following lemma can be seen as an analogue of [Lemma 2, [17]], in that it gives a condition for when the shifting operation does not increase the 'size' of any part in the tree-decomposition, when interpreted as $\langle \sigma_t \rangle_r$. The proof of this lemma follows closely the proofs of [Lemmas 6.1 and 8.3, [52]].

Lemma 3.2.15. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|_r$ for some non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$ and let $p \in \mathbb{N}$. Let (T,α) be an S-tree over a set of multistars and let $\vec{e} = (t,t') \in \vec{E}(T)$, $s \in V(T(\vec{e})) \setminus \{t\}$ and $\sigma = \{(A_0,B_0),(A_1,B_1),\dots,(A_n,B_n)\} = \alpha(\vec{F}_s)$, with $(A_0,B_0) = \alpha(\vec{g})$ where $\vec{g}$ is the unique $\vec{g} \in \vec{F}_s$ with $\vec{e} \leqslant \vec{g}$. Let (X,Y) be a non-trivial separation that is linked to $\alpha(\vec{e})$, (T',α') be the shift of (T,α) onto (X,Y) with respect to $\vec{e}$ and write $\sigma' = \alpha'(\vec{F}_s)$. Then $\langle\sigma'\rangle_r \leqslant \langle\sigma\rangle_r$ and if equality holds, then*

$$|\Big(\bigcap_{i\neq 0} B_i\Big) \cap X, \Big(\bigcup_{i\neq 0} A_i\Big) \cup Y|_r = |X,Y|_r$$

Proof. By definition of α',

$$\sigma' = \alpha'(\vec{F}_s) = \{(A_0 \cup X, B_0 \cap Y), (A_1 \cap Y, B_1 \cup X), \dots, (A_n \cap Y, B_n \cup X)\}.$$

So, we wish to show that

$$\langle\sigma'\rangle_r = r(B_0 \cap Y) + \sum_{i=1}^{n} r(B_i \cup X) - n \cdot r(V) \leqslant \sum_{i=0}^{n} r(B_i) - n \cdot r(V) = \langle\sigma\rangle_r. \tag{3.2.1}$$

Since r is submodular, it follows that

$$r(B_0 \cap Y) + r(B_0 \cup Y) \leqslant r(B_0) + r(Y)$$

and, for each $i = 1,\dots,n$,

$$r(B_i \cap X) + r(B_i \cup X) \leqslant r(B_i) + r(X).$$

By re-arranging these questions and adding them together, we see that

$$\begin{aligned} r(B_0 \cup Y) + \sum_{i=1}^{n} r(B_i \cap X) &\leqslant \sum_{i=0}^{n} r(B_i) - \left(r(B_0 \cap Y) + \sum_{i=1}^{n} r(B_i \cup X) \right) + r(Y) + n \cdot r(X) \\ &= (\langle\sigma\rangle_r - \langle\sigma'\rangle_r) + r(Y) + n \cdot r(X). \end{aligned}$$

Therefore, in order to show (3.2.1) it will be sufficient to show that

$$r(B_0 \cup Y) + \sum_{i=1}^{n} r(B_i \cap X) \geqslant r(Y) + n \cdot r(X)$$

By Lemma 3.2.12 applied to the multistar $\sigma \setminus \{(A_0,B_0)\}$ and X,

$$\sum_{i=1}^{n} r(B_i \cap X) \geqslant r\left(\bigcap_{i=1}^{n} B_i \cap X\right) + (n-1) \cdot r(X).$$

Let us write $A^* = \bigcup_{j=1}^{n} A_j$ and $B^* = \bigcap_{j=1}^{n} B_j$. Note that $A \subseteq A_0 \subseteq B^*$ and $B \supseteq B_0 \supseteq A^*$. Since (X,Y) and $(B^*,A^*) \in \vec{S}$, so is $(X \cap B^*, Y \cup A^*)$ and $(A,B) \leqslant (X \cap B^*, Y \cup A^*) \leqslant (X,Y)$. Hence, since (X,Y) is linked to (A,B), it follows that $|X \cap B^*, Y \cup A^*|_r \geqslant |X,Y|_r$. Therefore by definition of $|\cdot|_r$,

$$r(B^* \cap X) + r(A^* \cup Y) \geqslant r(X) + r(Y).$$

Since $A^* \subset B_0$ and r is non-decreasing, it follows that $r(B_0 \cup Y) \geqslant r(A^* \cup Y)$ and so we can conclude that

$$\begin{aligned} r(B_0 \cup Y) + \sum_{i=1}^{n} r(B_i \cap X) &\geqslant r(B_0 \cup Y) + r(B^* \cap X) + (n-1) \cdot r(X) \\ &\geqslant r(A^* \cup Y) + r(B^* \cap X) + (n-1) \cdot r(X) \\ &\geqslant r(Y) + n \cdot r(X). \end{aligned}$$

Finally, if equality holds in (3.2.1) then it holds throughout the argument, and so in particular

$$|X \cap B^*, Y \cup A^*|_r = |X, Y|_r.$$

□

The following is a simple corollary of Lemma 3.2.15.

Corollary 3.2.16. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|_r$ for some non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$ and let $p \in \mathbb{N}$. Then $\mathcal{F}_p$ is fixed under shifting.*

3.3 Linked S_k-trees

In this section we will prove a general theorem on the existence of linked tree-decompositions. The proof follows closely the proof of Geelen, Gerards and Whittle [68], extending their result to a broader class of tree-decompositions. They consider branch decompositions of integer-valued symmetric submodular functions. Given such a function λ on the subsets of a set V we can consider λ as an order function on the universe $\vec{S} = \{(A, V \setminus A)\}$ of bipartitions of V by taking $|A, V \setminus A| = \lambda(A)$, where by scaling by an additive factor we may assume λ is a non-negative function. They defined a notion of 'linkedness' for these decompositions and showed the following theorem:

Theorem 3.3.1. *[[68], Theorem 2.1] An integer-valued symmetric submodular function with branch-width n has a linked branch-decomposition of width n.*

A direct application of Theorem 3.3.1 gives analogues of Theorem 3.1.3 for branch decompositions of matroids or graphs [68], and also rank-decompositions of graphs [106].

It can be shown that a branch-decomposition of λ is equivalent to an S-tree over a set $\mathcal{T}$ of multistars which is fixed under shifting (See the proof of [52, Lemma 4.3], and that the width of this branch-decomposition is the smallest k such that the S-tree is an S_k-tree. Furthermore, in this way the definition of linkedness given in [68] coincides with Definition 3.2.4.

Let us say that a universe of set separations $\vec{S}$ with an order function $|\cdot|$ is *grounded* if it satisfies the conclusion of Lemma 3.2.2, that is, if for every pair of separations $(B, A) \leqslant (C, D)$ we have that $|A \cup C, B \cap D| \leqslant |C, D|$. We will need the following short lemma.

Lemma 3.3.2. *Let $\vec{S}$ be a grounded universe of set separations with an order function $|\cdot|$. If $(A, B), (X, Y) \in \vec{S}$ are such that (X, Y) is linked to (A, B) and (X, Y) is trivial, then $|X, Y| = |A, B|$.*

Proof. Since $(A, B) \leqslant (X, Y)$ and (X, Y) is trivial, $Y = B = V$, where V is ground set of $\vec{S}$. Furthermore, since $(A, B) \leqslant (X, Y)$ and $\vec{S}$ is grounded, it follows that

$$|A, B| = |V, A| = |B \cup X, A \cap Y| \leqslant |X, Y|.$$

Since (X, Y) is linked to (A, B), $|X, Y| \leqslant |A, B|$ and hence $|A, B| = |X, Y|$. □

Lemma 3.2.2 says that $\vec{S}$ is grounded whenever $|\cdot| = |\cdot|_r$ for some non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$. We note further that, if $\vec{S}$ is the universe of bipartitions of a set then $\vec{S}$ is grounded for any order function $|\cdot|$. Indeed, since $|\cdot|$ is symmetric and submodular, it follows that for any bipartition (X, Y),

$$2|X,Y| = |X,Y| + |Y,X| \geqslant |X \cup Y, X \cap Y| + |X \cap Y, X \cup Y| = 2|V, \emptyset|$$

and so, if $(B, A) \leqslant (C, D)$ then $|A \cup C, B \cap D| = |V, \emptyset| \leqslant |C, D|$.

In this way the following theorem implies, and generalises Theorem 3.3.1.

Theorem 3.3.3. *Let $\vec{S}$ be a grounded universe of set separations with an order function $|\cdot|$, $k \in \mathbb{N}$, and let $\mathcal{F} \subset \mathbb{N}^{S_k}$ be a family of multistars which is fixed under shifting. If there exists an S_k-tree over $\mathcal{F}$, then there exists a linked S_k-tree over $\mathcal{F}$.*

Proof. Let (T, α) be an S_k-tree over $\mathcal{F}$. For an integer p, we write T_p for the subforest of T where $e \in E(T_p)$ if and only if $|\alpha(\vec{e})| \geqslant p$ for either orientation of e. Let us write $e(T_p)$ for the number of edges of T_p and $c(T_p)$ for the number of components of T_p.

We define an order on the set

$$\mathcal{T} = \{(T, \alpha) : (T, \alpha) \text{ an } S_k\text{-tree over } \mathcal{F}\}$$

as follows. We say that $(T, \alpha) \prec (S, \alpha')$ if there exists a $p \in \mathbb{N}$ such that for all $p' > p$, $e(T_{p'}) = e(S_{p'})$ and $c(T_{p'}) = c(S_{p'})$ and either:

- $e(T_p) < e(S_p)$, or
- $e(T_p) = e(S_p)$ and $c(T_p) > c(S_p)$.

Let (T, α) be a $\prec$-minimal element of $\mathcal{T}$. We claim that (T, α) is linked.

Suppose not, that is, there are two edges $\vec{e} \leqslant \vec{f}$ such that there is no $\vec{g} \in \vec{E}(T)$ with $\vec{e} \leqslant \vec{g} \leqslant \vec{f}$ and

$$|\alpha(\vec{g})| = \lambda\left(\alpha(\vec{e}), \alpha(\vec{f})\right).$$

Let $\alpha(\vec{e}) := (A, B)$ and $\alpha(\vec{f}) := (D, C)$. Now, there is some separation $(A, B) \leqslant (X, Y) \leqslant (D, C)$ such that

$$|X, Y| = \lambda((A, B), (D, C)) =: \ell.$$

Let us choose such an (X, Y) which is nested with a maximum number of separations in $\alpha\left(\vec{E}(T)\right)$. Note that (X, Y) is linked to (A, B) and (Y, X) is linked to (C, D) and, since $\vec{S}$ is grounded and $\ell < \min\{|A, B|, |C, D|\}$, by Lemma 3.3.2 (X, Y) is non-trivial.

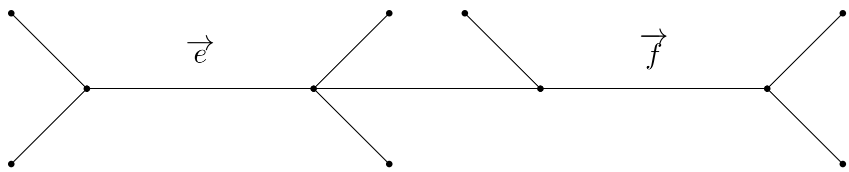

Figure 3.1: The tree T.

Let (T_1, α_1) be the S_k-tree given by shifting (T, α) onto (X, Y) with respect to $\vec{e}$ and let (T_2, α_2) be the S_k-tree given by shifting (T, α) onto (Y, X) with respect to $\overleftarrow{f}$. By Lemmas 3.2.9 and 3.2.11, (T_1, α_1) and (T_2, α_2) are S_k-trees over $\mathcal{F} \cup \{(Y, X)\}$ and $\mathcal{F} \cup \{(X, Y)\}$ respectively. For each vertex and edge in T let us write v_1 and v_2 or e_1 and e_2 for the copy of v or e in T_1 and T_2 respectively, if it exists (Note that, since $T_1 = T(\vec{e})$ and $T_2 = T(\overleftarrow{f})$, not every vertex or edge will appear in both trees). We let $(\hat{T}, \hat{\alpha})$ be the following S_k-tree:

- $\hat{T}$ is the tree formed by taking the disjoint union of T_1 and T_2, and identifying the edge $\overrightarrow{e}_1$ with the edge $\overrightarrow{f}_2$;
- $\hat{\alpha}$ is formed by taking the union of α_1 and α_2 on the domain $\overrightarrow{E}(\hat{T})$.

We remark that, since $\alpha_1(\overrightarrow{e}_1) = (X,Y) = \alpha_2(\overrightarrow{f}_2)$, the map α' is well defined. By Lemma 3.2.11 $(\hat{T}, \hat{\alpha})$ is an S_k-tree over $\mathcal{F}$.

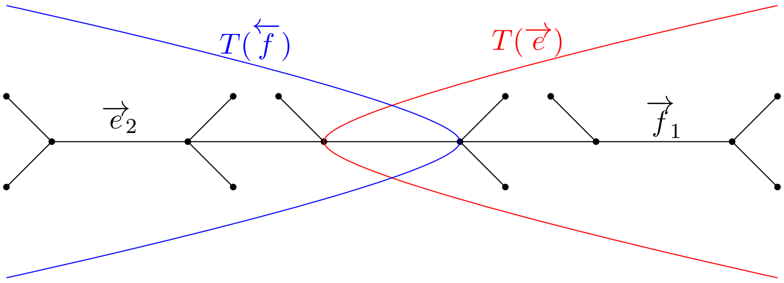

Figure 3.2: The tree $\hat{T}$.

Note that, since (X,Y) is linked to (A,B) and (Y,X) is linked to (C,D) it follows from Lemma 3.2.9 that for every $w \in E(T(\overrightarrow{e})) \cap E(T(\overleftarrow{f}))$

$$|\alpha(\overrightarrow{w})| \geqslant \max\{|\hat{\alpha}(\overrightarrow{w}_1)|, |\hat{\alpha}(\overrightarrow{w}_2)|\}.$$

Claim 3.3.4. *If*

$$|\hat{\alpha}(\overrightarrow{w}_i)| = |\alpha(\overrightarrow{w})| > \ell$$

then,

$$|\hat{\alpha}(\overrightarrow{w}_{2-i})| \leqslant \ell.$$

Remark. *For ease of exposition, we consider this to be vacuously satisfied if $\overrightarrow{w}_{2-i}$ does not exist.*

Proof. Indeed, suppose without loss of generality that $|\alpha(\overrightarrow{w})| = |\hat{\alpha}(\overrightarrow{w}_1)|$ and that $\overrightarrow{e} \leqslant \overrightarrow{w}$. Note that, $\alpha(\overrightarrow{e}) = (A,B) \leqslant \alpha(\overrightarrow{w}) =: (E,F)$. Then,

$$\hat{\alpha}(\overrightarrow{w}_1) = (E,F) \vee (X,Y).$$

Since $|\hat{\alpha}(\overrightarrow{w}_1)| = |E,F|$, by submodularity of the order function $|(E,F) \wedge (X,Y)| \leqslant |X,Y|$. However,

$$(A,B) \leqslant (E,F) \wedge (X,Y) \leqslant (D,C)$$

and so $(E,F) \wedge (X,Y)$ was a candidate for (X,Y). Moreover, since (E,F) is nested with every separation in $\alpha\left(\overrightarrow{E}(T)\right)$, it follows from Lemma 3.2.1 that $(E,F) \wedge (X,Y)$ is nested with every separation in $\alpha\left(\overrightarrow{E}(T)\right)$ that (X,Y) is, and also with (E,F) itself. Hence, by our choice of (X,Y), it follows that (X,Y) was already nested with (E,F).

We may suppose that $w \in E(T(\overrightarrow{e})) \cap E(T(\overleftarrow{f}))$, since otherwise there is no $\overrightarrow{w}_2$. Hence, since $\overrightarrow{e} \leqslant \overrightarrow{w}$, there are two cases to consider, either w is on the path in T from e to f or not. Let us suppose we are in the first case, and so $\overrightarrow{w} \leqslant \overrightarrow{f}$. Then, by definition, $\alpha_1(\overrightarrow{w}_1) = (E,F) \vee (X,Y) = (E \cup X, F \cap Y)$ and $\alpha_2(\overleftarrow{w}_2) = (F,E) \vee (Y,X) = (F \cup Y, E \cap X)$.

There are now four cases as to how (E,F) and (X,Y) are nested. If $(E,F) \leqslant (X,Y)$ then $\alpha_1(\overrightarrow{w}_1) = (X,Y)$, contradicting our assumption that $|\hat{\alpha}(\overrightarrow{w}_1)| > \ell$. If $(F,E) \leqslant (Y,X)$ then $\alpha_2(\overleftarrow{w}_2) = (Y,X)$, and so $|\hat{\alpha}(\overrightarrow{w}_2)| = \ell$.

If $(F,E) \leqslant (X,Y)$ then, since $\vec{S}$ is grounded, $|E \cup X, F \cap Y| = |\alpha_1(\vec{w}_1)| \leqslant |X,Y| = \ell$, again a contradiction, and if $(E,F) \leqslant (Y,X)$ then $|F \cup Y, E \cap X| = |\alpha_2(\overleftarrow{w}_2)| \leqslant |X \cap Y| = \ell$.

Suppose then that w is not on the path in T from e to f, and so $\overleftarrow{w} \leqslant \vec{f}$. Again, $\alpha_1(\vec{w}_1) = (E,F) \vee (X,Y) = (E \cup X, F \cap Y)$ and in this case $\alpha_2(\vec{w}_2) = (E,F) \vee (Y,X) = (E \cup Y, F \cap X)$.

As before, there are four cases as to how (E,F) and (X,Y) are nested. If $(E,F) \leqslant (X,Y)$ then $\alpha_1(\vec{w}_1) = (X,Y)$, a contradiction, and if $(E,F) \leqslant (Y,X)$ then $\alpha_2(\vec{w}_2) = (Y,X)$, and $|\hat{\alpha}(\vec{w}_2)| = \ell$.

If $(F,E) \leqslant (X,Y)$ then, since $\vec{S}$ is grounded, $|E \cup X, F \cap Y| = |\alpha_1(\vec{w}_1)| \leqslant |X \cap Y| = \ell$, a contradiction, and finally if $(F,E) \leqslant (Y,X)$ then $|E \cup Y, F \cap X| = |\alpha_2(\vec{w}_2)| \leqslant |X,Y| = \ell$. $\square$

Claim 3.3.5. *For every $p > \ell$ and every $\vec{w} \in \vec{E}(T)$ with $|\alpha(\vec{w})| = p$ exactly one of $\hat{\alpha}(\vec{w}_1), \hat{\alpha}(\vec{w}_2)$ has order p, and the other has order $\leqslant \ell$. Furthermore, for each component C of T_p if $|\hat{\alpha}(\vec{w}_i)| = |\alpha(\vec{w})|$ for some $\vec{w} \in \vec{E}(C)$, then $|\hat{\alpha}(\vec{w'}_i)| = |\alpha(\vec{w'})|$ for every $\vec{w'} \in \vec{E}(C)$.*

Proof of Claim. Suppose for contradiction that $p > \ell$ is the largest integer where the claim fails to hold. It follows that $e(\hat{T}_{p'}) = e(T_{p'})$ and $c(\hat{T}_{p'}) = c(T_{p'})$ for all $p' > p$. Hence, by $\prec$-minimality of T, $e(\hat{T}_p) \geqslant e(T_p)$. However, by assumption, for every separation $\alpha(\vec{w})$ of order $> p$ exactly one of $\hat{\alpha}(\vec{w}_1), \hat{\alpha}(\vec{w}_2)$ has the same order, and the other has order $\leqslant \ell$. Also, if $|\alpha(\vec{w})| = p > \ell$, then $|\alpha(\vec{w})| \geqslant \max\{|\hat{\alpha}(\vec{w}_1)|, |\hat{\alpha}(\vec{w}_2)|\}$ and by Claim 3.3.4 if one of $\hat{\alpha}(\vec{w}_1), \hat{\alpha}(\vec{w}_2)$ has the same order as $\alpha(\vec{w})$ then the other has order $\leqslant \ell$.

Hence, $e(\hat{T}_p) \leqslant e(T_p)$, and so by $\prec$-minimality of T it follows that $e(\hat{T}_p) = e(T_p)$, and for every $\vec{w}$ with $|\alpha(\vec{w})| = p$ exactly one of $\hat{\alpha}(\vec{w}_1), \hat{\alpha}(\vec{w}_2)$ is of order p, and the other has order $\leqslant \ell$.

Recall that $\hat{T}$ was formed by joining a copy of $T(\vec{e})$ and $T(\overleftarrow{f})$ along a separation (X,Y) of order $\ell < p$. It follows from the first part of the claim that $c(\hat{T}_p) \geqslant c(T_p)$, and so by $\prec$-minimality of T, $c(\hat{T}_p) = c(T_p)$. Hence, for each component C of T_p if $|\hat{\alpha}(\vec{w}_i)| = |\alpha(\vec{w})|$ for some $\vec{w} \in \vec{E}(C)$, then $|\hat{\alpha}(\vec{w'}_i)| = |\alpha(\vec{w'})|$ for every $\vec{w'} \in \vec{E}(C)$. Therefore the claim holds for p, contradicting our assumption. $\square$

By assumption, $\vec{e}$ and $\vec{f}$ lie in the same component of $T_{\ell+1}$, since for every $\vec{e} \leqslant \vec{g} \leqslant \vec{f}$

$$|\alpha(\vec{g})| > |X,Y| = \ell.$$

However, $\alpha_2(\vec{e}_2) = (A,B)$ and $\alpha_1(\overleftarrow{f}_1) = (C,D)$, contradicting Claim 3.3.5. $\square$

3.4 $\mathcal{F}$-lean S_k-trees

Theorem 3.4.1. *Let $\vec{S}$ be a universe of set separations with an order function $|\cdot|_r$ for some non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$, $k \in \mathbb{N}$, and let $\mathcal{F} \subset \mathbb{N}^{S_k}$ be a family of multistars which is fixed under shifting. If there exists an S_k-tree over $\mathcal{F}$, then there exists an $\mathcal{F}$-lean S_k-tree over $\mathcal{F}$.*

Proof. Let (T,α) be an S_k-tree over $\mathcal{F}$. We write T^p for the induced subforest of T on the set of vertices

$$V(T^p) = \{t \in V(T) : \langle \sigma_t \rangle_r \geqslant p\}.$$

Let us write $v(T^p)$ for the number of vertices of T^p and $c(T^p)$ for the number of components.

We define an order on the set

$$\mathcal{T} = \{(T,\alpha) : (T,\alpha) \text{ an } S_k\text{-tree over } \mathcal{F}\}$$

as follows. We say that $(T,\alpha) \prec (S,\alpha')$ if there exists an $p \in \mathbb{N}$ such that for all $p' > p$, $v(T^{p'}) = v(S^{p'})$ and $c(T^{p'}) = c(S^{p'})$ and either:

- $v(T^p) < v(S^p)$, or
- $v(T^p) = v(S^p)$ and $c(T^p) > c(S^p)$.

Let (T, α) be a $\prec$-minimal element of $\mathcal{T}$. We claim that (T, α) is $\mathcal{F}$-lean.

Suppose not. Then there is some pair of vertices $t, t' \in V(T)$ and a pair of separations $(A, B) \leqslant (B', A')$ such that (A, B) is addable at t and (A', B') is addable at t', with

$$\lambda\big((A, B), (B', A')\big) < \min\{r(A), r(A')\}$$

and

$$|\alpha(\vec{g})| > \lambda\big((A, B), (B', A')\big) \text{ for all } \vec{g} \in \vec{E}(tTt').$$

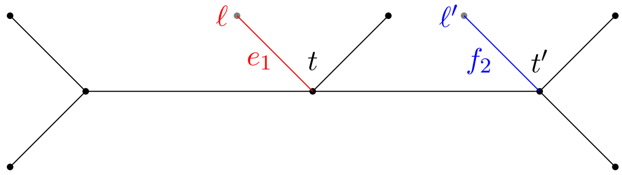

Figure 3.3: The tree T.

Let (T_1, α_1) be the S_k-tree formed in the following way:

- T_1 is the tree formed by adding an extra leaf ℓ to T at t;
- The restriction of α_1 to $\vec{E}(T)$ is α and $\alpha_1(\ell, t) = (A, B)$.

Similarly let (T_2, α_2) be the S_k-tree formed in the following way:

- T_2 is the tree formed by adding an extra leaf ℓ' to T at t';
- The restriction of α_2 to $\vec{E}(T)$ is α and $\alpha_2(\ell', t') = (A', B')$.

Note that, since (A, B) and (A', B') were addable at t and t', (T_1, α_1) and (T_2, α_2) are S_k-trees over $\mathcal{F} \cup \{(B, A)\}$ and $\mathcal{F} \cup \{(B', A')\}$ respectively. Let us denote by $\vec{e_1}$ the edge $(\ell, t) \in \vec{E}(T_1)$ and $\overleftarrow{f_2}$ the edge $(\ell', t') \in \vec{E}(T_2)$. For each vertex v and edge g in T let us write v_1 and v_2 or g_1 and g_2 for the copy of v or g in T_1 and T_2 respectively, and similarly $\sigma^i_{v_i}$ for the multistars $\alpha_i(\vec{F}_{v_i})$. Note that, for every $v \neq t, t'$ we have $\sigma^i_{v_i} = \sigma_v$, however $\sigma^1_{t_1} = \sigma_t \cup \{(A, B)\}$ and $\sigma^2_{t'_2} = \sigma_{t'} \cup \{(A', B')\}$.

Over all separations (X, Y) such that $(A, B) \leqslant (X, Y) \leqslant (B', A')$ and $|X, Y|_r = \lambda\big((A, B), (B', A')\big) =: \ell$ we pick (X, Y) which is nested with a maximum number of separations in $\alpha(\vec{E}(T))$. Note that (X, Y) is linked to (A, B) and (Y, X) is linked to (A', B'). Furthermore, if (X, Y) is trivial, then $Y = B = V$ and by Lemma 3.3.2 $|X, Y|_r = r(X) = r(A) = |A, B|_r$ contradicting our assumption on ℓ. Hence (X, Y) is non-trivial.

Let (T_1, β_1) be the S_k-tree given by shifting (T_1, α_1) onto (X, Y) with respect to $\vec{e_1}$ and let (T_2, β_2) be the S_k-tree given by shifting (T_2, α_2) onto (Y, X) with respect to $\overleftarrow{f}_2$. Note that, since $\mathcal{F}$ is fixed under shifting, by Lemmas 3.2.9 and 3.2.11, (T_1, α_1) and (T_2, β_2) are S_k-trees over $\mathcal{F} \cup \{(X, Y)\}$ and $\mathcal{F} \cup \{(Y, X)\}$ respectively and that, since ℓ and ℓ' were leaves, these are indeed S_k-trees with underlying trees T_1 and T_2 respectively. We let $(\hat{T}, \hat{\alpha})$ be the following S_k-tree:

- $\hat{T}$ is the tree formed by taking the disjoint union of T_1 and T_2 and identifying the edge $\vec{e}_1$ with the edge $\vec{f}_2$;

- $\hat{\alpha}$ is formed by taking the union of β_1 and β_2 on the domain $\vec{E}(\hat{T})$.

We remark that, since $\beta_1(\vec{e}_1) = (X,Y) = \beta_2(\vec{f}_2)$, the map $\hat{\alpha}$ is well defined and furthermore, by Lemma 3.2.11, $(\hat{T}, \hat{\alpha})$ is an S_k-tree over $\mathcal{F}$. For each vertex in $v \in V(T)$ there are now two copies v_1 and v_2 in $\hat{T}$. We will write $\hat{\sigma}_{v_i}$ for the multistars $\hat{\alpha}(\vec{F}_{v_i})$.

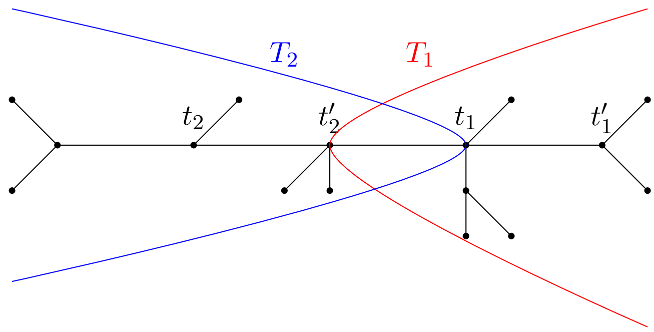

Figure 3.4: The tree $\hat{T}$.

Note that, since (X,Y) is linked to (A,B) and (Y,X) is linked to (A',B') it follows from Lemma 3.2.15 that for every $s_i \in V(T_i)$

$$\langle \hat{\sigma}_{s_i} \rangle_r \leqslant \langle \sigma^i_{s_i} \rangle_r \leqslant \langle \sigma_s \rangle_r, \tag{3.4.1}$$

where the final inequality holds trivially for all $s \neq t, t'$, and for $s = t, t'$ since adding a separation to a multistar only decreases the size.

Claim 3.4.2. *If*

$$\langle \hat{\sigma}_{s_i} \rangle_r = \langle \sigma_s \rangle_r > \ell$$

then

$$\langle \hat{\sigma}_{s_{i-2}} \rangle_r \leqslant \ell.$$

Proof. Indeed, let us assume without loss of generality that

$$\langle \hat{\sigma}_{s_1} \rangle_r = \langle \sigma^1_{s_1} \rangle_r,$$

and suppose first that $s \neq t, t'$, and so $\sigma^i_{s_i} = \sigma_s$. Let us write

$$\sigma_s = \{(A_0, B_0), \dots, (A_n, B_n)\}.$$

For each $s \in V(T) \setminus \{t, t'\}$ there is a unique edge $\vec{g} \in \vec{F}_s$ such that $\vec{e}_1 \leqslant \vec{g}_1$, and similarly a unique edge $\vec{h} \in \vec{F}_s$ such that $\overleftarrow{f}_2 \leqslant \vec{h}_2$. Let us suppose without loss of generality that $\alpha(\vec{g}) = (A_0, B_0)$ and $\alpha(\vec{h}) = (A_j, B_j)$, where perhaps $j = 0$. Then

$$\hat{\sigma}_{s_1} = \{(A_0 \cup X, B_0 \cap Y), \dots, (A_n \cap Y, B_n \cup X)\}$$

and

$$\hat{\sigma}_{s_2} = \{(A_j \cup Y, B_j \cap X), \dots, (A_n \cap X, B_n \cup Y)\}$$

Now, since $\langle \hat{\sigma}_{s_1} \rangle_r = \langle \sigma_s \rangle_r$, by Lemma 3.2.15

$$|B^* \cap X, A^* \cup Y|_r = |X, Y|_r,$$

where $B^* = \bigcap_{i \neq 0} B_i$ and $A^* = \bigcup_{i \neq 0} A_i$.

However, $(A,B) \leqslant (B^* \cap X, A^* \cup Y) \leqslant (B', A')$ and so, by our choice of (X,Y), (X,Y) is nested with at least as many separations in $\alpha(\overrightarrow{E}(T))$ as $(B^* \cap X, A^* \cup Y)$. However, since (B^*, A^*) is nested with $\alpha(\overrightarrow{E}(T))$ and $(B^* \cap X, A^* \cup Y) \leqslant (B_i, A_i)$ for each $i \neq 0$, it follows by Lemma 3.2.1 that (X,Y) was already nested with (B_i, A_i) for each $i \neq 0$.

We note that if $(B_i, A_i) \leqslant (Y,X)$ for any $i \neq 0$, then $(A,B) \leqslant (B_i, A_i) \leqslant (Y,X)$ and $(A,B) \leqslant (X,Y)$. Hence $A \subset X \cap Y$ and so $r(A) \leqslant r(X \cap Y) \leqslant |X,Y|_r$, contradicting our assumption on A. Similarly if $(B_i, A_i) \leqslant (X,Y)$ for $i \neq j$, then it contradicts our assumption.

Therefore, for each $i \neq j, 0$, either $(X,Y) \leqslant (B_i, A_i)$ or $(Y,X) \leqslant (B_i, A_i)$. Note that, for each $i \neq j, 0$, $\{B_i \cup X, B_i \cup Y\} = \{B_i, V\}$. Hence,

$$\begin{aligned}
\langle \hat{\sigma}_{s_1} \rangle_r + \langle \hat{\sigma}_{s_2} \rangle_r &= r(B_0 \cap Y) + r(B_j \cap X) + \sum_{i \neq 0} r(B_i \cup X) + \sum_{i \neq j} r(B_i \cup Y) - 2n \cdot r(V) \\
&= r(B_0 \cap Y) + r(B_j \cap X) + r(B_0 \cup Y) + r(B_j \cup X) + \sum_{i \neq 0,j} r(B_i) - (n+1) \cdot r(V) \\
&= \langle \sigma_s \rangle_r + r(B_0 \cap Y) + r(B_j \cap X) + r(B_0 \cup Y) - r(B_0) - r(B_j) - r(V) \\
&\leqslant \langle \sigma_s \rangle_r + r(Y) + r(X) - r(V) \\
&= \langle \sigma_s \rangle_r + |X,Y|_r.
\end{aligned}$$

From which is follows that, if $\langle \hat{\sigma}_{s_1} \rangle_r = \langle \sigma_s \rangle_r$, then $\langle \hat{\sigma}_{s_2} \rangle_r \leqslant |X,Y|_r$. If $s = t$ or t', then a similar calculation holds. □

Claim 3.4.3.

$$\langle \hat{\sigma}_{t_1} \rangle_r < \langle \sigma_t \rangle_r.$$

and

$$\langle \hat{\sigma}_{t'_2} \rangle_r < \langle \sigma_{t'} \rangle_r.$$

Proof. Let us write

$$\sigma_t = \{(A_1, B_1), \ldots, (A_n, B_n)\},$$

and so

$$\langle \hat{\sigma}_{t_1} \rangle_r = r(Y) + \sum_{i=1}^{n} r(B_i \cup X) - n \cdot r(V).$$

Suppose for contradiction that $\langle \hat{\sigma}_{t_1} \rangle_r = \langle \sigma_t \rangle_r$. As before it follows that we can split (A_i, B_i) into sets I and I' such that $(X,Y) \leqslant (B_j, A_j)$ for $j \in I$ and $(Y,X) \leqslant (B_j, A_j)$ for $j \in I'$. It follows that

$$\langle \hat{\sigma_{t_1}} \rangle_r = r(Y) + \sum_{j \in I} r(B_j) - |I| \cdot r(V).$$

However, since σ_t is a multistar it follows by Lemma 3.2.12 that

$$\sum_{j \in I'} r(B_j) \geqslant r(\bigcap_{j \in I'} B_j) + (|I'| - 1) \cdot r(V).$$

Now, $A \subset \bigcap_{i \in I'} B_i$ and also by definition of I', $Y \subset \bigcap_{j \in I'} B_j$ and hence $r(Y \cup A) \leqslant r(\bigcap_{j \in I'} B_j)$. However, by submodularity

$$r(X) + r(A \cup Y) \geqslant r(A) + r(V)$$

and by assumption $r(A) > |X,Y|_r = r(X) + r(Y) - r(V)$, and so

$$r(A \cup Y) > r(Y),$$

from which we conclude that

$$\begin{aligned} r(Y) &< r(Y \cup A) \\ &\leqslant r(\bigcap_{j \in I'} B_j) \\ &\leqslant \sum_{j \in I'} r(B_j) - (|I'| - 1) \cdot r(V). \end{aligned}$$

Therefore,

$$\begin{aligned} \langle \hat{\sigma}_{t_1} \rangle_r &= r(Y) + \sum_{j \in I} r(B_j) - |I| \cdot r(V) \\ &< \sum_{i=1}^{n} r(B_i) - (n-1) \cdot r(V) = \langle \sigma_t \rangle_r. \end{aligned}$$

A similar argument shows that $\langle \hat{\sigma}_{t'_2} \rangle_r < \langle \sigma_{t'} \rangle_r$. □

Claim 3.4.4. *For every $p > \ell$ and every $s \in V(T)$ with $\langle \sigma_s \rangle_r = p$ exactly one of $\hat{\sigma}_{s_1}, \hat{\sigma}_{s_2}$ has size p, and the other has size $\leqslant \ell$ and further for each component C of T^p, if $\langle \hat{\sigma}_{s_i} \rangle_r = \langle \sigma_s \rangle_r$ for some $s \in V(C)$, then $\langle \hat{\sigma}_{s'_i} \rangle_r = \langle \sigma_{s'} \rangle_r$ for every $s' \in V(C)$.*

Proof. Suppose for contradiction that $p > \ell$ is the largest integer where the claim fails to hold. We split into three cases. First let us suppose $p \geqslant \max\{\langle \sigma_t \rangle_r, \langle \sigma_{t'} \rangle_r\}$.

By assumption, for all $p' > p$, $v(\hat{T}^p) = v(T^p)$ and $c(\hat{T}^p) = c(T^p)$. Hence, by $\prec$-minimality of T $v(\hat{T}^p) \geqslant v(T^p)$. However, since $p \geqslant \langle \sigma_t \rangle_r \geqslant r(A) > \ell$, for every $s \in V(T)$ with $\langle \sigma_s \rangle_r = p$, by Claim 3.4.2 if one of $\hat{\sigma}_{s_1}, \hat{\sigma}_{s_2}$ has size $\langle \sigma_s \rangle_r$, then the other has size $\leqslant \ell$.

Therefore, it follows that $v(\hat{T}^p) \leqslant v(T^p)$, and so $v(\hat{T}^p) = v(T^p)$, and for every $s \in V(T)$ with $\langle \sigma_s \rangle_r = p$ exactly one of $\hat{\sigma}_{s_1}, \hat{\sigma}_{s_2}$ has size $\langle \sigma_s \rangle_r$, and the other has size $\leqslant \ell$. Recall that $\hat{T}$ is formed by joining two copies of T by an edge between t_1 and t'_2, and that by Claim 3.4.3 $\hat{\sigma}_{t_1}$ and $\hat{\sigma}_{t'_2}$ both have size $< \max\{\langle \sigma_t \rangle_r, \langle \sigma_{t'} \rangle_r\} \leqslant p$. It follows that $c(\hat{T}^p) \geqslant c(T^p)$, and so again by $\prec$-minimality of T, $c(\hat{T}^p) = c(T^p)$. Further, for each component C of T^p, if $\langle \hat{\sigma}_{s_i} \rangle_r = \langle \sigma_s \rangle_r$ for some $s \in V(C)$, then $\langle \hat{\sigma}_{s'_i} \rangle_r = \langle \sigma_{s'} \rangle_r$ for every $s' \in V(C)$. Therefore the claim holds for p, contradicting our assumption.

Suppose then that $p = \max\{\langle \sigma_t \rangle_r, \langle \sigma_{t'} \rangle_r\}$, say without loss of generality $p = \langle \sigma_t \rangle_r$. As before we conclude that for every $s \in V(T)$ with $\langle \sigma_s \rangle_r = p$ exactly one of $\hat{\sigma}_{s_1}, \hat{\sigma}_{s_2}$ has size $\langle \sigma_s \rangle_r$, and the other has size $\leqslant \ell$. In this case, since by Claim 3.4.3 $\langle \hat{\sigma}_{t_1} \rangle_r < \langle \sigma_t \rangle_r = p$, it follows that $\langle \hat{\sigma}_{t_1} \rangle_r \leqslant \ell$. This allows us to conclude that the copies of each component of T^p in $\hat{T}^p$ are separated by the vertex t_1, and thus $c(\hat{T}^p) \geqslant c(T^p)$. So, as before we can conclude that for each component C of T^p, if $\langle \hat{\sigma}_{s_i} \rangle_r = \langle \sigma_s \rangle_r$ for some $s \in V(C)$, then $\langle \hat{\sigma}_{s'_i} \rangle_r = \langle \sigma_{s'} \rangle_r$ for every $s' \in V(C)$. Therefore the claim holds for p, contradicting our assumption.

Finally, when $\ell < p < \max\{\langle \sigma_t \rangle_r, \langle \sigma_{t'} \rangle_r\}$ the same argument will hold, since one of $\langle \hat{\sigma}_{t_1} \rangle_r$ or $\langle \hat{\sigma}_{t_2} \rangle_r$ has size $\leqslant \ell$. □

Suppose that $t = t'$. Since (A, B) and (A', B') are addable at σ_t it follows by Lemma 3.2.14 that $\langle \sigma_t \rangle_r \geqslant r(A) \geqslant \min\{r(A), r(A')\} > \ell$, and so by Claim 3.4.4

$$\langle \sigma_t \rangle_r = \max\{\langle \hat{\sigma}_{t_1} \rangle_r, \langle \hat{\sigma}_{t_2} \rangle_r\},$$

contradicting Claim 3.4.3.

Suppose that $t \neq t'$. Then, since

$$|\alpha(\overrightarrow{g})| > \ell \text{ for all } \overrightarrow{g} \in tTt,$$

it follows by Lemma 3.2.13 that $\langle \sigma_s \rangle_r > \ell$ for all $s \in V(tTt')$, and so t and t' are in the same component of $T^{(\ell+1)}$. However, $\langle \hat{\sigma}_{t_1} \rangle_r < \ell$ and $\langle \hat{\sigma}_{t'_2} \rangle_r < \ell$, contradicting Claim 3.4.4.

□

Remark. *Unlike in the case of tree-width for graphs it is not true in general that the existence of a $\mathcal{F}$-lean S_k-tree over $\mathcal{F}$ implies the existence of a linked S_k-tree over $\mathcal{F}$. It seems likely that similar methods should prove the existence of an S_k-tree which is both linked and $\mathcal{F}$-lean. However, since Theorem 3.3.3 could be stated in slightly more generality, and to avoid lengthening an already quite technical proof, we have stated the results separately. Explicitly, one should consider a minimal element of the following order on the set of S_k-trees over $\mathcal{F}$: we let $T < S$ if there is some $p \in \mathbb{N}$ such that:*

- $e(T_p) < e(S_p)$; *or*
- $e(T_p) = e(S_p)$ *and* $c(T_p) > c(S_p)$; *or*
- $e(T_p) = e(S_p), c(T_p) = c(S_p)$ *and* $v(T^p) < v(S^p)$; *or*
- $e(T_p) = e(S_p), c(T_p) = c(S_p), v(T^p)v(S^p)$ *and* $c(T^p) > c(S^p)$,

and for all $p' > p$, all four quantities are equal.

3.5 Applications

3.5.1 Graphs

Throughout this subsection $\overrightarrow{S}$ will be the universe of separations of some graph G. Given a multistar

$$\sigma = \{(A_0, B_0), (A_1, B_1), \dots, (A_n, B_n)\}$$

we define $\text{int}(\sigma) := \bigcap_{i=0}^n B_i$. Given an S-tree (T, α) we can construct an associated tree-decompositon $(T, \mathcal{V})_\alpha$ by letting

$$V_t = \text{int}(\sigma_t) \text{ for each } t \in T.$$

Lemma 3.5.1. *If (T, α) is an $\overrightarrow{S}$-tree over some family of multistars then $(T, \mathcal{V})_\alpha$ is a tree-decompositon. Furthermore, if $\alpha(t, t') = (A, B)$ then $V_t \cap V_{t'} = A \cap B$.*

Proof. Given an edge $e \in E(G)$ we can define an orientation on $E(T)$ as follows. For each edge $f \in E(T)$ let us pick one of the orientations $\overrightarrow{f}$ such that $\alpha(\overrightarrow{f}) = (A, B)$ with $e \in E(G[B])$. Note that, since $\alpha(\overrightarrow{f}) = (A, B)$ is a separation, either $e \in E(G[A])$ or $e \in E(G[B])$. Any orientation has a sink t, and it is a simple check that $e \in V_t$. This proves that $(T, \mathcal{V})$ satisfies the first two properties of a tree-decomposition.

Finally suppose that $t_2 \in V(t_1Tt_3)$ and $x \in V_{t_1} \cap V_{t_3}$. There is a unique edge $\overrightarrow{e} \in \overrightarrow{F}_{t_2}$ such that $t_1 \in V(T(\overleftarrow{e}))$ and similarly a unique edge $\overrightarrow{f} \in \overrightarrow{F}_{t_2}$ such that $t_3 \in V(T(\overleftarrow{f}))$. Let $\alpha(\overrightarrow{e}) = (A, B)$ and $\alpha(\overrightarrow{f}) = (C, D)$ and let us write

$$\sigma_{t_2} = \{(A, B), (C, D), (A_1, B_1), \dots, (A_n, B_n)\}.$$

Since $x \in V_{t_1}$ and α preserves the tree-ordering, it follows that $x \in A$ and hence $x \in D \cap \bigcap_{i=1}^n B_i$. Similarly, since $x \in V_{t_3}$, $x \in C$ and hence $x \in B \cap \bigcap_{i=1}^n B_i$ and so

$$x \in B \cap D \cap \bigcap_{i=1}^n B_i = \text{int}(\sigma_{t_2}) = V_{t_2}.$$

Finally, suppose that $\alpha(t,t') = (A,B)$. Since $V_t = \mathrm{int}(\sigma_t) \subseteq A$ and $V_{t'} = \mathrm{int}(\sigma_{t'}) \subseteq B$ it follows that $V_t \cap V_{t'} \subseteq A \cap B$. Conversely it is a simple check that, since σ_t and $\sigma_{t'}$ are multistars, $A \cap B \subset \mathrm{int}(\sigma_t)$ and $A \cap B \subset \mathrm{int}(\sigma_{t'})$. It follows that $V_t \cap V_{t'} = A \cap B$ □

We note that, conversely, given a tree-decomposition $(T, \mathcal{V})$ we can build an S_k-tree over some family of multistars $(T,\alpha)_{\mathcal{V}}$ by letting

$$\alpha(\vec{e}) = \alpha(t,t') = \left(\bigcup_{s\,:\,t \in V(sTt')} V_s, \bigcup_{s\,:\,t' \in V(sTt)} V_s \right),$$

and in this way the two notions are equivalent.

In this way, by applying Theorems 3.3.3 and 3.4.1 to appropriate families $\mathcal{F}$ of multistars we can prove a number of results about tree-decompositions of graphs, some known and some new. Since the notion of linked from Definition 3.1.1 will not be appropriate for talking about every type of tree-decomposition we consider, we make the following definition:

Definition 3.5.2. *A tree decomposition $(T, \mathcal{V})$ is called* $\vec{S}$-linked *if for all $k \in \mathbb{N}$ and every $t, t' \in V(T)$, either G contains k disjoint V_t-$V_{t'}$ paths or there is an edge $\{s, s'\} \in E(tTt')$ such that $|V_s \cap V_{s'}| < k$.*

As in the introduction, it is easy to see that given an $\vec{S}$-linked tree-decomposition, by subdividing each edge and adding as a bag the separating set $V_t \cap V_{t'}$ we obtain a linked tree-decomposition in the sense of Definition 3.1.1.

Lemma 3.5.3. *Let $\vec{S}$ be the universe of graph separations for some graph G. If (T,α) is a linked S_k-tree over some family of multistars then $(T,\mathcal{V})_\alpha$ is $\vec{S}$-linked.*

Proof. Suppose that (T,α) is a linked S_k-tree over some family of multistars. Given $r \in \mathbb{N}$ and $t, t' \in V(T)$ such that G does not contain r disjoint V_t-$V_{t'}$ paths, we wish to show that there is an edge $\{s, s'\} \in E(tTt')$ such that $|V_s \cap V_{s'}| < r$. Let $\vec{e}$ be the unique edge adjacent to t such that $t' \in V(T(\vec{e}))$ and similarly let $\overleftarrow{f}$ be the unique edge adjacent to t' such that $t \in V(T(\overleftarrow{f}))$. Note that $\vec{e} \leqslant \vec{f}$.

Let us write $\alpha(\vec{e}) = (A,B)$ and $\alpha(\vec{f}) = (C,D)$. Since $A \cap B \subset V_t$ and $C \cap D \subset V_{t'}$, and G does not contain r disjoint $V_t - V_{t'}$ paths, it follows by Menger's theorem that $\lambda((A,B),(C,D)) < r$. Hence, since (T,α) is linked, there is some edge $\vec{e} \leqslant \vec{g} \leqslant \vec{f}$ such that $|\alpha(\vec{g})| =: |X,Y| < r$. Let us write $\vec{g} = (s,s')$. Note that, $\{s,s'\} \in E(tTt')$ by construction, and by Lemma 3.5.1 $|V_s \cap V_{s'}| = |X \cap Y| < r$, as claimed. □

For many families of multistars, being $\mathcal{F}$-lean will not tell us much about the tree-decomposition. Indeed, if $\mathcal{F}$ only contains multi-sets of size 3 or 1 (as in the case of branch decompositions), then there is never an addable separation for any multistar in $\mathcal{F}$. However, for certain families of multistars being $\mathcal{F}$-lean will imply leaness in the traditional sense.

Definition 3.5.4. *Let $\vec{S}$ be a separation system. A family of multistars $\mathcal{F} \subset 2^{\vec{S}}$ is* S-stable *if whenever $\sigma \in \mathcal{F}$ and $(A,B) \in \vec{S}$ is such that $\sigma \cup \{(A,B)\}$ is a multistar then $\sigma \cup \{(A,B)\} \in \mathcal{F}$.*

Lemma 3.5.5. *Let $\vec{S}$ be a universe of separations with an order function $|\cdot|_r$ for some non-decreasing submodular function $r : 2^V \to \mathbb{N}$, and let $p \in \mathbb{N}$. Then $\mathcal{F}_p$ is S-stable.*

Proof. If $(A,B) \in \vec{S}$ is such that $\sigma \cup \{(A,B)\}$ is a multistar then,

$$\langle \sigma \cup \{(A,B)\} \rangle_r = \langle \sigma \rangle_r + r(B) - r(V) \leqslant \langle \sigma \rangle_r < p,$$

and so $\sigma \cup \{(A,B)\} \in \mathcal{F}_p$. □

Lemma 3.5.6. *Let $\vec{S}$ be the universe of graph separations for some graph G and let $\mathcal{F} \subset \mathbb{N}^{\vec{S}}$ be an S-stable family of multistars. If (T,α) is an $\mathcal{F}$-lean S-tree over $\mathcal{F}$ then $(T,\mathcal{V})_\alpha$ is lean.*

Proof. Given $k \in \mathbb{N}$, $t,t' \in V(T)$ and vertex sets $Z_1 \subseteq V_t$ and $Z_2 \subseteq V_{t'}$ with $|Z_1| = |Z_2| = k$, such that G does not contain k disjoint Z_1-Z_2 paths, we wish to show that there exists an edge $\{s,s'\} \in E(tTt')$ with $|V_s \cap V_{s'}| < k$.

Since (T,α) is over $\mathcal{F}$, both σ_t and $\sigma_{t'} \in \mathcal{F}$. Furthermore, since $Z_1 \subset V_t$ it follows that $\{(Z_1,V)\} \cup \sigma_t$ forms a multistar, and similarly so does $\{(Z_2,V)\} \cup \sigma_{t'}$. Hence, since $\mathcal{F}$ is $\vec{S}$-stable, both of these multistars are in $\mathcal{F}$, and so (Z_1,V) is addable at σ_t and (Z_2,V) is addable at $\sigma_{t'}$. Since G does not contain k disjoint Z_1-Z_2 paths, it follows that $\lambda\left((Z_1,V),(V,Z_2)\right) < k$. Hence, since (T,α) is $\mathcal{F}$-lean, there exists an edge $\vec{g} \in \vec{E}(tTt')$ such that $|\alpha(\vec{g})| := |X,Y| = \lambda\left((Z_1,V),(V,Z_2)\right)$. As before let us write $\vec{g} = (s,s')$. Then $\{s,s'\} \in E(tTt')$ and $|V_s \cap V_{s'}| = |X \cap Y| < k$, as claimed. $\square$

Given a multistar $\sigma \in \mathbb{N}^{\vec{S}}$ let us write $n(\sigma)$ for the cardinality of the multiset σ. Diestel and Oum [52] showed that for the families of multistars

$$\mathcal{F}_k := \{\sigma \in \mathbb{N}^{\vec{S}_k} : \sigma \text{ a multistar}, \langle\sigma\rangle_r < k\},$$

$$\mathcal{P}_k := \{\sigma \subset \mathcal{F}_k : n(\sigma) \leqslant 2\}$$

and

$$\begin{aligned}\mathcal{T}_k :=&\{\sigma \in \mathbb{N}^{\vec{S}_k} : \sigma \text{ a multistar}, \sigma = \{(A_1,B_1),(A_2,B_2),(A_3,B_3)\} \text{ and } \bigcup_{i=1}^{3} G[A_i] = G\}\\ &\cup \{\sigma \in \mathcal{F}_k : n(\sigma) = 1\},\end{aligned}$$

the following statements are true:

- G admits a tree-decomposition of width $< k-1$ if and only if there is an S_k-tree over $\mathcal{F}_k$,
- G admits a path-decomposition of width $< k-1$ if and only if there is an S_k-tree over $\mathcal{P}_k$,
- G admits a branch-decomposition of width $< k$ if and only if there is an S_k-tree over $\mathcal{T}_k$[4].

Furthermore, they showed that in every case the family of multistars was fixed under shifting[5]. Also, they showed that if $\theta \leqslant p$ then $\mathcal{F}_p^\theta := \mathcal{F}_p \cap \mathbb{N}^{\vec{S}_\theta}$ is fixed under shifting, and that

- G has θ-tree-width $< p-1$ if and only if there is an S_θ-tree over $\mathcal{F}_p^\theta$,

where the θ-tree-width of G θ-tw(G) is defined, as in [66], to be the smallest p such that G admits a tree-decomposition of width p and adhesion $< \theta$.

Definition 3.5.7. *A tree decomposition $(T,\mathcal{V})$ is called* θ-lean *if for all $k < \theta$, $t,t' \in V(T)$ and vertex sets $Z_1 \subseteq V_{t_1}$ and $Z_2 \subseteq V_{t_2}$ with $|Z_1| = |Z_2| = k$, either G contains k disjoint Z_1-Z_2 paths or there exists an edge $\{s,s'\} \in E(tTt')$ with $|V_s \cap V_{s'}| < k$.*

We note that by the same arguments as Lemmas 3.5.5 and 3.5.6, if an S_θ-tree over $\mathcal{F}_p^\theta$ is $\mathcal{F}_p^\theta$-lean, then the corresponding tree-decomposition is θ-lean. Let us write pw(G), bw(G), and θ-tw(G) for the path-, branch-, and θ-tree-width of G respectively. Applying Theorem 3.3.3 to these families of multistars gives the following theorem:

[4] Due to a quirk in how branch-width is defined, this is only true for $k > 3$, see the comment in [52] after Theorem 4.4.

[5] Strictly they showed that a slightly weaker condition, which they called *closed under shifting*, holds (see the comment at the start of section 3.2.4). However the same proof shows they are also fixed under shifting.

Theorem 3.5.8. *Let G be a graph then the following statements are true:*

- *G admits an $\vec{S}$-linked path-decomposition of width $pw(G)$;*
- *G admits an $\vec{S}$-linked branch-decomposition of width $bw(G)$;*
- *G admits a lean tree-decomposition of width $tw(G)$;*
- *G admits a θ-lean tree-decomposition of width θ-$tw(G)$.*

We note that none of these results are in essence new. The first result appears in a paper of Lagergren [90], where it is attributed to Seymour and Thomas. The second result is implied by a broader theorem [[68], Theorem 2.1] on linked branch decompositions of submodular functions and the third is Thomas' theorem, Theorem 3.1.3 in this paper. The fourth is stated without proof in [[34], Theorem 2.3], and a slightly weaker result is claimed in [[66], Theorem 3.1] although the proof in fact shows the stronger statement.

As mentioned before, there is never an addable separation for any multistar $\sigma \in \mathcal{T}_k$, and so Theorem 3.4.1 gives us no insight into branch decompositions. However, for path-decompositions it tells us something about the bags at the two leaves.

Lemma 3.5.9. *Let (T, α) be an $\mathcal{P}_k$-lean S_k-tree over $\mathcal{P}_k$ and let T be a path $t_0, t_1, \dots, t_n$. Then $(T, \mathcal{V})_\alpha$ has the following properties:*

- *For all $Z_1, Z_2 \subset V_{t_0}$ or $Z_1, Z_2 \subset V_{t_n}$ with $|Z_1| = |Z_2| = r$ there are r-disjoint $Z_1 - Z_2$ paths in G.*
- *For all $Z_1 \subset V_{t_0}$ and $Z_2 \subset V_{t_n}$ with $|Z_1| = |Z_2| = r$ either there are r disjoint $Z_1 - Z_2$ paths in G or there is an edge $\{t_i, t_{i+1}\} \in E(T)$ with $|V_{t_i} \cap V_{t_{i+1}}| < r$.*

Proof. Let us show the first statement, the proof of the second is similar. Suppose without loss of generality that $Z_1, Z_2 \subset V_{t_0}$. Since $Z_i \subset V_{t_0}$ it follows that $|Z_i| = r < k$ and so $(Z_i, V) \in \vec{S}_k$. Let $\alpha(t_1, t_0) = (A, B)$. Since $\{(A, B)\} \in \mathcal{P}_k$, $|B| < k$ and since $B = V_{t_0}$ and $Z_i \subseteq V_{t_0}$ both (Z_i, V) are addable at $\{(A, B)\}$.

Therefore, since (T, α) is $\mathcal{P}_k$-lean, it follows that $\lambda\left((Z_0, V), (V, Z_1)\right) \geqslant \min\{|Z_i|\} = r$. Hence G contains r disjoint Z_1-Z_2 paths. □

Not only is this result broad enough to imply many known theorems, the framework is also flexible enough to encompass many other types of tree-decompositions. For example, more recently, Diestel, Eberenz and Erde [46] showed that there exist families of multistars $\mathcal{B}_k$ and $\mathcal{P}_k \subset \mathbb{N}^{\vec{S}_k}$, which are fixed under shifting, such that the existence of S_k-trees over $\mathcal{B}_k$ or $\mathcal{P}_k$ is dual to the existence of a k-block or a k-profile in the graph respectively (k-blocks and k-profiles are examples of what Diestel and Oum call "highly cohesive structures" which represent obstructions to low width, see [51]). They defined the *profile-width* and *block-width* of a graph G, which we denote by $\mathrm{blw}(G)$ and $\mathrm{prw}(G)$, to be the smallest k such that there is an S_k-tree over $\mathcal{B}_k$ or $\mathcal{P}_k$ respectively. Again, applying Theorems 3.3.3 and 3.4.1 to these families of multistars we get the following theorem:

Theorem 3.5.10. *Let G be a graph, then the following statements are true:*

- *G admits an $\vec{S}$-linked profile-decomposition of width $prw(G)$;*
- *G admits an $\vec{S}$-linked block-decomposition of width $blw(G)$.*

The family $\mathcal{B}_k$ of multistars is built from a family $\mathcal{B}_k^*$ of multistars by iteratively taking all possible multistars that appear as shifts of multistars in $\mathcal{B}_k^*$, in order to guarantee that $\mathcal{B}_k$ is fixed under shifting. The set $\mathcal{B}_k^*$ can be taken to be stable, but it is not clear if this property is maintained when moving to $\mathcal{B}_k$. It would be interesting to know if a lean block-decomposition of width $\mathrm{blw}(G)$ always exists.

3.5.2 Matroid tree-width

Hliněný and Whittle [78, 79] generalized the notion of tree-width from graphs to matroids. Let $M = (E, I)$ be a matroid with rank function r. Hliněný and Whittle defined a tree-decomposition of M to be a pair (T, τ) where $\tau : E \to V(T)$ is an arbitrary map. Every vertex $v \in V(T)$ separates the tree into a number of components $T_1, T_2, \ldots, T_d$ and we define the *width* of the bag $\langle \tau^{-1}(v) \rangle$ to be

$$\sum_{i=1}^{d} r(E \setminus \tau^{-1}(T_i)) - (d-1) \cdot r(E).$$

The *width* of a tree-decomposition is $\max\{\langle \tau^{-1}(v) \rangle \,:\, v \in V(T)\}$ and the *tree-width* of M is the smallest k such that M has a tree-decomposition of width k. This is a generalisation of the tree-width of graphs, and in particular Hliněný and Whittle showed that for any graph G with at least one edge, if $M(G)$ is the cycle matroid of G then the tree-width of G is the tree width of $M(G)$.

We can express their notion of a tree-decomposition of a matroid in the language of S_k-trees in the following way. Given any $X \subset E$ the *connectivity* of X is given by

$$\lambda(X) := r(X) + r(E \setminus X) - r(M),$$

where r is the rank function of M. If we consider the universe of separations $\vec{S}$ given by the bipartitions of E, that is, pairs of the form $(X, E \setminus X)$, it follows that $|X, E \setminus X|_r = \lambda(X)$ is an order function on $\vec{S}$.

Let us define, as before

$$\mathcal{F}_k = \{\sigma \in \mathbb{N}^{\vec{S}_k} \,:\, \sigma \text{ a multistar with } \langle \sigma \rangle < k\}.$$

Diestel and Oum [Lemma 8.4 [52]] showed that M has tree-width $< k$ in the sense of Hliněný and Whittle if and only if there is an S_k-tree over $\mathcal{F}_k$. Explicitly given an S_k-tree (T, α) there is a natural map $\tau : E \to T$ where $\tau(e) = v$ if and only if $e \in B$ for all $(A, B) \in \alpha(\vec{F}_v)$. Conversely given a tree-decomposition (T, τ) and an edge $\vec{f} = (t_1, t_2)$ consider the two subtrees T_1 and T_2 consisting of the component of $T - t_2$ containing t_1 and the component of $T - t_1$ containing t_2 respectively. We can then define $\alpha(\vec{f}) = \left(\tau^{-1}(V(T_1)), \tau^{-1}(V(T_2))\right)$. In this way we get an equivalence between S_k-trees and matroid tree-decompositions, and it is easy to check that the width of a bag $\langle \tau^{-1}(v) \rangle = \langle \sigma_v \rangle_r$. Let us say that a matroid tree-decomposition is *linked* if the corresponding S_k-tree is linked.

We first note that Azzato [15] showed that Theorem 3.3.1 implies the following theorem:

Theorem 3.5.11. *Every matroid has a linked tree-decomposition of width at most* $2tw(M)$.

It is a simple corollary of Theorem 3.3.3 that this bound can be improved to the best possible bound.

Corollary 3.5.12. *Every matroid has a linked tree-decomposition of width at most* $tw(M)$.

However we can also apply Theorem 3.4.1 to give us a generalization of Theorem 3.1.3 to matroids. If we wish to express this in the framework of Hliněný and Whittle we could make the following definition. Given disjoint subsets $Z_1, Z_2 \subseteq E$ let us write

$$\lambda(Z_1, Z_2) := \min\{\lambda(X) \,:\, Z_1 \subseteq X, Z_2 \subseteq E \setminus X\}.$$

Definition 3.5.13. *A matroid tree decomposition (T, τ) is called* lean *if for all $k \in \mathbb{N}$, $t, t' \in T$ and subsets $Z_1 \subseteq \tau^{-1}(t)$ and $Z_2 \subseteq \tau^{-1}(t')$ with $r(Z_1) = r(Z_2) = k$, either $\lambda(Z_1, Z_2) \geqslant k$ or there exists an edge $\{s, s'\} \in E(tTt')$ such that, if we let T_1 be the component of $T - s'$ containing s and T_2 be the component of $T - s$ containing s', then $\left(\tau^{-1}\left(V(T_1)\right), \tau^{-1}\left(V(T_2)\right)\right)$ is a $(< k)$-separation.*

It is a simple argument in the vein of Lemmas 3.5.5 and 3.5.6 that if an S_k-tree over $\mathcal{F}_k$ is $\mathcal{F}_k$-lean then the associated matroid tree-decomposition is lean. Then, Theorem 3.4.1 applied to $\mathcal{F}_k$ gives us the following generalization of Theorem 3.1.3 to matroids, the main new result in this paper.

Theorem 3.5.14. *Every matroid M admits a lean tree-decomposition of width $tw(M)$.*

We note that a non-negative non-decreasing submodular function $r : 2^V \to \mathbb{N}$ is, if normalised such that $r(\emptyset) = 0$, a polymatroid set function. So, in the broadest generality our results can be interpreted in terms of tree-decompositions of polymatroidal set functions. However, we are not aware of any references in the literature to such tree-decompositions.

Acknowledgement

The author would like to thank Nathan Bowler for useful contributions to the proof of Theorem 3.4.1.

Chapter 4

Structural submodularity and tangles in abstract separation systems

4.1 Introduction

This paper is, in a sense, the capstone of a comprehensive project [35, 33, 36, 37, 44, 45, 46, 48, 51, 52, 58, 59] whose aim has been to utilize the idea of tangles familiar from Robertson and Seymour's graph minors project as a way of capturing clusters in other contexts, such as image analysis [54], genetics [39], or the social sciences [40]. The idea is to use tangles, which in graphs are certain consistent ways of orienting their low-order separations, as an indirect way of capturing 'fuzzy' clusters – ones that cannot easily be described by simply listing their elements – by instead orienting all those low-order separations towards them. We can then think of these as a collection of signposts all pointing to that cluster, and of clusters as collective targets of such consistent pointers.

Once clusters have been captured by 'abstract tangles' in this way, one can hope to generalize to such clusters Robertson and Seymour's two fundamental results about tangles in graphs [110]. One of these is the *tree-of-tangles* theorem. It says that any set of distinguishable tangles – ones that pairwise do not contain each other – can in fact be distinguished pairwise by a small and nested set of separations: for every pair of tangles there is a separation in this small and nested collection that distinguishes them. Formally, this means that these two tangles orient it differently; informally it means that one of its two orientations points to one of the tangles, while its other orientation points to the other tangle. Since these separations are nested, they split the underlying structure in a tree-like way, giving it a rough overall structure.

The other fundamental result from [110], the *tangle-tree duality* theorem, tells us that if there are no tangles of a desired type then the entire underlying structure can be split in such a tree-like way, i.e. by some nested set of separations, so that the regions corresponding to a node of the structure tree are all small. (What exactly this means may depend on the type of tangle considered.)

This research programme required a number of steps, of which this paper constitutes the last.

The first step was to make the notion of tangles independent from their natural habitat of graphs. In a graph, tangles are ways of consistently orienting all its separations $\{A, B\}$ up to some given order, either as (A, B) or as (B, A). If we want to do this for another kind of underlying structure than a graph, this structure will have to come with a notion of 'separation', it must be possible to 'orient' these separations, and there must be a difference between doing this 'consistently' or 'inconsistently'. If we wish to express, and perhaps prove, the two fundamental tangle theorems in such an abstract context, we further need a notion of when two 'separations' are nested.

There are many structures that come with a natural notion of separation. For sets, for example, we might simply take bipartitions. The notion of nestedness can then be borrowed from the nestedness of sets and applied to the bipartition classes. Thinking of a bipartition as an unordered pair of subsets, we can also naturally orient it 'towards one or the other of these subsets' by ordering the pair. Finally, we have to come up with natural notions of when orientations of different separations are consistent: we think of this as 'roughly pointing the same way', and it is another prerequisite for defining tangles to make this formal. This is both trickier to do in an abstract context and one of our main sources of freedom; we shall address this question in Section 4.2.

The completion of the first step in our research programme thus consisted in abstracting from the various notions of separation, and of consistently orienting separations, a minimum set of requirements that might serve as axioms for an abstract notion of tangle applicable to all of them. This resulted in the concept of *separation systems* and their ('abstract') tangles [44].

The second step, then, was to generalize the proofs of the tree-of-tangles theorem and the tangle-tree duality theorem to the abstract setting of separation systems. This was done in [48] and [51], respectively.

In order to prove these theorems, or to apply them to concrete cases of abstract separation systems, e.g. as in [52, 54], one so far still needed a further ingredient of graph tangles: a submodular order function on the separation system considered. Our aim in this paper is to show that one can do without this: we shall prove that a structural consequence of the existence of a submodular order function, a consequence that can be expressed in terms of abstract separation systems, can replace the assumption that such a function exists in the proofs of the above two theorems. We shall refer to separation systems that satisfy this structural condition as *submodular separation systems*. [1]

With this third step, then, the programme sketched above will be complete: we shall have a notion of tangle for very general abstract separation systems, as well as a tree-of-tangles theorem and a tangle-tree duality theorem for these tangles that can be expressed and proved without the need for any submodulary order function on the separation systems considered.

Formally, our two main results read as follows:

Theorem 4.1.1. *Every submodular separation system $\vec{S}$ contains a tree set of separations that distinguishes all the abstract tangles of S.*

Theorem 4.1.2. *Let $\vec{S}$ be a submodular separation system without degenerate elements in a distributive universe $\vec{U}$. Then exactly one of the following holds:*

(cl.1) *S has an abstract tangle.*

(cl.2) *There exists an S-tree over $\mathcal{T}^*$ (witnessing that S has no abstract tangle).*

(See Section 4.2 for definitions.) Three further theorems, which partly strengthen or generalize the above two, will be stated in Section 4.2 (and proved later) when we have more terminology available.

[1] There is also a notion of submodularity for separation *universes*. Separation universes are special separation systems that are particularly large, and they are always submodular as separation systems. For separation universes, therefore, submodularity is used with the narrower meaning of being endowed with a submodular order function [44].

One may ask, of course, whether weakening the existence of a submodular order function to ‘structural submodularity’ in the premise of these two theorems is worth the effort. We believe it is. For a start, the entire programme of developing abstract separation systems, and a theory of tangles for them, served the purpose of identifying the few structural assumptions one has to make of a set of objects called ‘separations’ in order to capture the essence of tangles in graphs, and thereby make them applicable in much wider contexts. It would then seem oblivious of these aims to stop just short of the goal: to continue to make unnecessarily strong assumptions of an extraneous and non-structural kind when weaker structural assumptions can achieve the same.

However, there is also a technical advantange. As we shall see in Sections 4.5.2 and 4.5.4, there are interesting abstract separation systems that are structurally submodular but which do not come with a natural submodular order function that implies this.

4.2 Abstract separation systems

Abstract separation systems were first introduced in [44]; see there for a gentle formal introduction and any terminology we forgot to define below. Motivation for why they are interesting can be found in the introductory sections of [48, 51, 52] and in [54]. In what follows we provide a self-contained account of just the definitions and basic facts about abstract separation systems that we need in this paper.

A *separation system* $(\vec{S}, \leqslant, {}^*)$ is a partially ordered set with an order-reversing involution ${}^*\colon \vec{S} \to \vec{S}$. The elements of $\vec{S}$ are called *(oriented) separations*. The *inverse* of $\vec{s} \in \vec{S}$ is $\vec{s}^{\,*}$, which we usually denote by $\overleftarrow{s}$. An *(unoriented) separation* is a set $s = \{\vec{s}, \overleftarrow{s}\}$ consisting of a separation and its inverse and we then refer to $\vec{s}$ and $\overleftarrow{s}$ as the two *orientations* of s. Note that it may occur that $\vec{s} = \overleftarrow{s}$, we then call $\vec{s}$ *degenerate*. The set of all separations is denoted by S. When the context is clear, we often refer to oriented separations simply as separations in order to improve the flow of text.

If the partial order $(\vec{S}, \leqslant)$ is a lattice with join $\vee$ and meet $\wedge$, then we call $(\vec{S}, \leqslant, {}^*, \vee, \wedge)$ a *universe* of (oriented) separations. It is *distributive* if it is distributive as a lattice. Typically, the separation systems we are interested in are contained in a universe of separations. In most applications, one starts with a universe $(\vec{U}, \leqslant, {}^*, \vee, \wedge)$ and then defines $\vec{S}$ as a set of separations of low order with respect to some *order function* on $\vec{U}$, a map $|\cdot|\colon \vec{U} \to [0, \infty)$ that is *symmetric* in that $|\vec{s}| = |\overleftarrow{s}|$, and *submodular* in that $|\vec{s} \vee \vec{t}\,| + |\vec{s} \wedge \vec{t}\,| \leqslant |\vec{s}| + |\vec{t}\,|$ for all $\vec{s}, \vec{t} \in \vec{U}$. Submodularity of the order function in fact plays a crucial role in several arguments. One of its most immediate consequences is that whenever both $\vec{s}, \vec{t} \in \vec{S_k} := \{\vec{u} \in \vec{U} \colon |\vec{u}| < k\}$, then at least one of $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$ again lies in $\vec{S_k}$.

In order to avoid recourse to the external concept of an order function if possible, let us turn this last property into a definition that uses only the language of lattices. Let us call a subset M of a lattice $(L, \vee, \wedge)$ *submodular* if for all $x, y \in M$ at least one of $x \vee y$ and $x \wedge y$ lies in M. A separation system $\vec{S}$ contained in a given universe $\vec{U}$ of separations is *(structurally) submodular* if it is submodular as a subset of the lattice underlying $\vec{U}$.

We say that $\vec{s} \in \vec{S}$ is *small* (and $\overleftarrow{s}$ is co-small) if $\vec{s} \leqslant \overleftarrow{s}$. An element $\vec{s} \in \vec{S}$ is *trivial* in $\vec{S}$ (and $\overleftarrow{s}$ is co-trivial) if there exists $t \in S$ whose orientations $\vec{t}, \overleftarrow{t}$ satisfy $\vec{s} < \vec{t}$ as well as $\vec{s} < \overleftarrow{t}$. Notice that trivial separations are small.

Two separations $s, t \in S$ are *nested* if there exist orientations $\vec{s}$ of s and $\vec{t}$ of t such that $\vec{s} \leqslant \vec{t}$. Two oriented separations are nested if their underlying separations are. We say that two separations *cross* if they are not nested. A set of (oriented) separations is *nested* if any two of its elements are. A nested separation system without trivial or degenerate elements is a *tree set*. A set σ of non-degenerate oriented separations is a *star* if for any two distinct $\vec{s}, \vec{t} \in \sigma$ we have $\vec{s} \leqslant \overleftarrow{t}$. A family $\mathcal{F} \subseteq 2^{\vec{U}}$ of sets of separations is *standard for* $\vec{S}$ if for any trivial $\vec{s} \in \vec{S}$ we have $\{\overleftarrow{s}\} \in \mathcal{F}$. Given $\mathcal{F} \subseteq 2^{\vec{U}}$, we write $\mathcal{F}^*$ for the set of all elements of $\mathcal{F}$ that are stars.

An *orientation* of S is a set $O \subseteq \vec{S}$ which contains for every $s \in S$, exactly one of $\overleftarrow{s}, \vec{s}$. An orientation O of S is *consistent* if whenever $r, s \in S$ are distinct and $\vec{r} \leqslant \vec{s} \in O$, then $\overleftarrow{r} \notin O$. The idea behind this is that separations $\overleftarrow{r}$ and $\vec{s}$ are thought of as pointing away from each other if $\vec{r} \leq \vec{s}$. If we wish to orient r and s towards some common region of the structure which they are assumed to 'separate', as is the idea behind tangles, we should therefore not orient them as $\overleftarrow{r}$ and $\vec{s}$.

Tangles in graphs also satisfy another, more subtle, consistency requirement: they never orient three separations r, s, t so that the region to which they point collectively is 'small'.[2] This can be mimicked in abstract separation systems by asking that three oriented separations in an 'abstract tangle' must never have a co-small supremum; see [44, Section 5]. So let us implement this formally.

Given a family $\mathcal{F} \subseteq 2^{\vec{U}}$, we say that O *avoids* $\mathcal{F}$ if there is no $\sigma \subseteq O$ with $\sigma \in \mathcal{F}$. A consistent $\mathcal{F}$-avoiding orientation of S is called an $\mathcal{F}$*-tangle* of S. An $\mathcal{F}$-tangle for $\mathcal{F} = \mathcal{T}$ with

$$\mathcal{T} := \{\{\vec{r}, \vec{s}, \vec{t}\} \subseteq \vec{U} \colon \vec{r} \vee \vec{s} \vee \vec{t} \text{ is co-small}\}$$

is an *abstract tangle*.

A separation $s \in S$ *distinguishes* two orientations O_1, O_2 of S if $O_1 \cap s \neq O_2 \cap s$. Likewise, a set N of separations *distinguishes* a set $\mathcal{O}$ of orientations if for any two $O_1, O_2 \in \mathcal{O}$, there is some $s \in N$ which distinguishes them.

Let us restate our tree-of-tangles theorem for abstract tangles of submodular separation systems:

Theorem 4.1.1. *Every submodular separation system $\vec{S}$ contains a tree set of separations that distinguishes all the abstract tangles of S.*

We now introduce the structural dual to the existence of abstract tangles. An S*-tree* is a pair (T, α) consisting of a tree T and a map $\alpha : \vec{E}(T) \to \vec{S}$ from the set $\vec{E}(T)$ of orientations of edges of T to $\vec{S}$ such that $\alpha(y, x) = \alpha(x, y)^*$ for all $xy \in E(T)$. Given $\mathcal{F} \subseteq 2^{\vec{U}}$, we call (T, α) an S*-tree over* $\mathcal{F}$ if $\alpha(F_t) \in \mathcal{F}$ for every $t \in T$, where

$$F_t := \{(s, t) \colon st \in E(T)\}.$$

It is easy to see that if S has an abstract tangle, then there can be no S-tree over $\mathcal{T}$.

Our tangle-tree duality theorem for abstract tangles of submodular separation systems, which we now re-state, asserts a converse to this. Recall that $\mathcal{T}^*$ denotes the set of stars in $\mathcal{T}$.

Theorem 4.1.2. *Let $\vec{S}$ be a submodular separation system without degenerate elements in a distributive universe $\vec{U}$. Then exactly one of the following holds:*

(cl.1) *S has an abstract tangle.*

(cl.2) *There exists an S-tree over $\mathcal{T}^*$ (witnessing that S has no abstract tangle).*

Here, it really is necessary to exclude degenerate separations: a single degenerate separation will make the existence of abstract tangles impossible, although there might still be $\mathcal{T}^*$-tangles (and therefore no S-trees over $\mathcal{T}^*$). We will actually prove a duality theorem for $\mathcal{T}^*$-tangles without this additional assumption and then observe that $\mathcal{T}^*$-tangles are in fact already abstract tangles, unless $\vec{S}$ contains a degenerate separation.

[2]Formally: so that the union of their sides to which they do *not* point is the entire graph.

In applications, we do not always wish to consider all the abstract tangles of a given separation system. For example, if S consists of the bipartitions $\{A, B\}$ of some finite set V (see [44] for definitions), then every $v \in V$ induces an abstract tangle

$$\tau_v := \big\{(A, B) \in \vec{S} \colon v \in B\big\},$$

the *principal tangle* induced by v. In particular, abstract tangles trivially exist in these situations. In order to exclude principal tangles, we could require that every tangle τ of S must satisfy $(\{v\}, V \setminus \{v\}) \in \tau$ for every $v \in V$.

More generally, we might want to prescribe for some separations s of S that any tangle of S we consider must contain a particular one of the two orientations of s rather than the other. This can easily be done in our abstract setting, as follows. Given $Q \subseteq \vec{U}$, let us say that an abstract tangle τ of S *extends* Q if $Q \cap \vec{S} \subseteq \tau$. It is easy to see that τ extends Q if and only if τ is $\mathcal{F}_Q$-avoiding, where

$$\mathcal{F}_Q := \{\{\overleftarrow{s}\} \colon \vec{s} \in Q \text{ non-degenerate}\}.$$

We call $Q \subseteq \vec{U}$ *down-closed* if $\vec{r} \le \vec{s} \in Q$ implies $\vec{r} \in Q$ for all $\vec{r}, \vec{s} \in \vec{U}$.

Here, then, is our refined tangle-tree duality theorem for abstract tangles of submodular separation systems.

Theorem 4.2.1. *Let $\vec{S}$ be a submodular separation system without degenerate elements in a distributive universe $\vec{U}$ and let $Q \subseteq \vec{U}$ be down-closed. Then exactly one of the following assertions holds:*

(cl.1) *S has an abstract tangle extending Q.*

(cl.2) *There exists an S-tree over $\mathcal{T}^* \cup \mathcal{F}_Q$.*

Observe that Theorem 4.2.1 implies Theorem 4.1.2 by taking $Q = \emptyset$.

The abstract tangles in Theorem 4.2.1 are not the only $\mathcal{F}$-tangles for which such a statement holds. In [52] the tangle-tree duality theorem of [51] is used to prove such a statement for a broad class of $\mathcal{F}$-tangles, albeit under a stronger assumption: one needs there that $\vec{S}$ is not just structurally submodular, as is our assumption here throughout our paper, but that $\vec{U}$ has a submodular order function and $\vec{S}$ is the set separations up to some fixed order (and therefore, in particular, submodular).

In Section 4.4, however, we will show that the weaker assumption that $\vec{S}$ itself is submodular is in fact sufficient to establish the only property of S whose proof in [52] requires a submodular order function: this is the fact that S is 'separable'. (We shall repeat the definition of this in Section 4.4.)

The other ingredient one needs for all those applications of the tangle-tree duality theorem from [51] is a property of $\mathcal{F}$: that $\mathcal{F}$ is 'closed under shifting'. Sometimes, a submodular order function on $\vec{U}$ is needed also to establish this property of $\mathcal{F}$. But if it is not, we can now prove the same application without a submodular order function, assuming only that S itself is submodular:

Theorem 4.2.2. *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a submodular separation system. Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars which is standard for $\vec{S}$ and closed under shifting. Then exactly one of the following holds:*

(cl.1) *There exists an $\mathcal{F}$-tangle of S.*

(cl.2) *There exists an S-tree over $\mathcal{F}$.*

We shall prove Theorem 4.2.2 in Section 4.4.

Our last result is an example of Theorem 4.2.2 for a concrete $\mathcal{F}$, a tangle-tree duality theorem for $\mathcal{F}$-tangles of bipartitions of a set that are used particularly often in applications [39, 40]. Let $\vec{U}$ be the universe of oriented bipartitions (A, B) of a set V (see [44] for definitions). Let $m \geq 1$ and $n \geq 2$ be integers, and define

$$\mathcal{F}_m := \{\, F \subseteq \vec{U} : \big| \bigcap_{(A,B)\in F} B \,\big| < m \,\}$$

and

$$\mathcal{F}_m^n := \{\, F \in \mathcal{F}_m : |F| < n \,\}.$$

To subsume $\mathcal{F}_m$ under this latter notation we allow $n = \infty$, so that $\mathcal{F}_m^\infty = \mathcal{F}_m$. Given any collection $\mathcal{F} \subseteq 2^{\vec{U}}$ of sets of oriented separations, we write $\mathcal{F}^*$ for its subcollection of those sets $F \in \mathcal{F}$ that are stars (of oriented separations).

We shall prove in Section 4.4 that the set of stars in $\mathcal{F}_m$ is closed under shifting. Building on Theorem 4.2.2, we then use this in Section 4.5 to prove the following:

Theorem 4.2.3. *Let $\vec{S} \subseteq \vec{U}$ be a submodular separation system, let $1 \leq m \in \mathbb{N}$ and $2 \leq n \in \mathbb{N} \cup \{\infty\}$, and let $\mathcal{F} = \mathcal{F}_m^n$. Then exactly one of the following two statements holds:*

(cl.1) *S has an $\mathcal{F}$-tangle;*

(cl.2) *There exists an S-tree over $\mathcal{F}^*$.*

The bound n on the size of the sets in $\mathcal{F}$ is often taken to be 4. In (i) we could replace $\mathcal{F}$ with $\mathcal{F}^*$, since for these $\mathcal{F}$ the $\mathcal{F}$-tangles are precisely the $\mathcal{F}^*$-tangles; see Section 4.5.3.

4.3 The tree-of-tangles theorem

In this section we will prove Theorem 4.1.1. In fact, we are going to prove a slightly more general statement. Let $\mathcal{P} := \{\{\vec{s}, \vec{t}, (\vec{s} \vee \vec{t})^*\} \colon \vec{s}, \vec{t} \in \vec{U}\}$. The $\mathcal{P}$-tangles are known as *H11*. A profile of S is *regular* if it contains all the small separations in $\vec{S}$.

Theorem 4.3.1. *Let $\vec{S}$ be a submodular separation system and Π a set of profiles of S. Then $\vec{S}$ contains a tree set that distinguishes Π.*

This implies Theorem 4.1.1, by the following easy observation.

Lemma 4.3.2. *Every abstract tangle is a profile.*

Proof. Let $\vec{s}, \vec{t} \in \vec{U}$ and $\vec{r} := \vec{s} \vee \vec{t}$. Then

$$\vec{s} \vee \vec{t} \vee \overleftarrow{r} = \vec{r} \vee \overleftarrow{r}$$

is co-small, so $\{\vec{s}, \vec{t}, \overleftarrow{r}\} \in \mathcal{T}$. Therefore $\mathcal{P} \subseteq \mathcal{T}$ and every $\mathcal{T}$-tangle is also a $\mathcal{P}$-tangle. □

We first recall a basic fact about nestedness of separations. For $s, t \in S$, we define the *corners* $\vec{s} \wedge \vec{t}$, $\vec{s} \wedge \overleftarrow{t}$, $\overleftarrow{s} \wedge \vec{t}$ and $\overleftarrow{s} \wedge \overleftarrow{t}$.

Lemma 4.3.3 ([44])**.** *Let $\vec{S}$ be a separation system in a universe $\vec{U}$ of separations. Let s, t be two crossing separations and $\vec{r}$ one of the corners. Then every separation that is nested with both s and t is nested with r as well.*

In the proof of Theorem 4.3.1, we take a nested set $\mathbb{N}$ of separations that distinguishes some set Π_0 of regular profiles and we want to exchange one element of $\mathbb{N}$ by some other separation while maintaining that Π_0 is still distinguished. The following lemma simplifies this exchange.

Lemma 4.3.4. *Let $\vec{S}$ be a separation system, $\mathcal{O}$ a set of consistent orientations of S and $\mathbb{N} \subseteq S$ an inclusion-minimal nested set of separations that distinguishes $\mathcal{O}$. Then for every $t \in \mathbb{N}$ there is a unique pair of orientations $O_1, O_2 \in \mathcal{O}$ that are distinguished by t and by no other element of $\mathbb{N}$.*

Proof. It is clear that at least one such pair must exist, for otherwise $\mathbb{N} \setminus \{t\}$ would still distinguish $\mathcal{O}$, thus violating the minimality of $\mathbb{N}$.

Suppose there was another such pair, say O_1', O_2'. After relabeling, we may assume that $\vec{t} \in O_1 \cap O_1'$ and $\overleftarrow{t} \in O_2 \cap O_2'$. By symmetry, we may further assume that $O_1 \neq O_1'$. Since $\mathbb{N}$ distinguishes $\mathcal{O}$, there is some $r \in \mathbb{N}$ with $\vec{r} \in O_1$, $\overleftarrow{r} \in O_1'$.

As t is the only element of $\mathbb{N}$ distinguishing O_1, O_2, it must be that $\vec{r} \in O_2$ as well, and similarly $\overleftarrow{r} \in O_2'$. We hence see that for any orientation τ of $\{r, t\}$, there is an $O \in \{O_1, O_2, O_1', O_2'\}$ with $\tau \subseteq O$. Since $\mathbb{N}$ is nested, there exist orientations of r and t pointing away from each other. But then one of O_1, O_2, O_1', O_2' is inconsistent, which is a contradiction. $\square$

Proof of Theorem 4.3.1. Note that it suffices to show that there is a nested set $\mathbb{N}$ of separations that distinguishes Π: Every consistent orientation contains every trivial and every degenerate element, so any inclusion-minimal such set $\mathbb{N}$ gives rise to a tree-set.

We prove this by induction on $|\Pi|$, the case $|\Pi| = 1$ being trivial.

For the induction step, let $P \in \Pi$ be arbitrary and $\Pi_0 := \Pi \setminus \{P\}$. By the induction hypothesis, there exists a nested set $\mathbb{N}$ of separations that distinguishes Π_0. If some such set $\mathbb{N}$ distinguishes Π, there is nothing left to show. Otherwise, for every nested $\mathbb{N} \subseteq S$ which distinguishes Π_0 there is a $P' \in \Pi_0$ which $\mathbb{N}$ does not distinguish from P. Note that P' is unique. For any $s \in S$ that distinguishes P and P', let $d(\mathbb{N}, s)$ be the number of elements of $\mathbb{N}$ which are not nested with s.

Choose a pair $(\mathbb{N}, s)$ so that $d(\mathbb{N}, s)$ is minimum. Clearly, we may assume $\mathbb{N}$ to be inclusion-minimal with the property of distinguishing Π_0. If $d(\mathbb{N}, s) = 0$, then $\mathbb{N} \cup \{s\}$ is a nested set distinguishing Π and we are done, so we now assume for a contradiction that $d(\mathbb{N}, s) > 0$.

Since $\mathbb{N}$ does not distinguish P and P', we can fix an orientation of each $t \in \mathbb{N}$ such that $\vec{t} \in P \cap P'$. Choose a $t \in \mathbb{N}$ such that t and s cross and $\vec{t}$ is minimal. Let (P_1, P_2) be the unique pair of profiles in Π_0 which are distinguished by t and by no other element of $\mathbb{N}$, say $\overleftarrow{t} \in P_1$, $\vec{t} \in P_2$. Let us assume without loss of generality that $\overleftarrow{s} \in P_1$. The situation is depicted in Figure 4.1. Note that we do not know whether $\vec{s} \in P_2$ or $\overleftarrow{s} \in P_2$. Also, the roles of P and P' might be reversed, but this is insignificant.

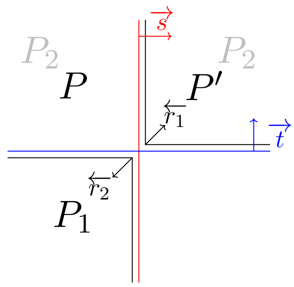

Figure 4.1: Crossing separations

Suppose first that $\vec{r_1} := \vec{s} \vee \vec{t} \in \vec{S}$. Let $Q \in \{P, P'\}$. If $\vec{s} \in Q$, then $\vec{r_1} \in Q$, since $\vec{t} \in P \cap P'$ and Q is a profile. If $\vec{r_1} \in Q$, then $\vec{s} \in Q$ since Q is consistent and $\vec{s} \leqslant \vec{r_1} \in Q$: it cannot be that $\vec{s} = \overleftarrow{r_1}$, since then s and t would be nested. Hence each $Q \in \{P, P'\}$ contains $\vec{r_1}$ if and only if it contains $\vec{s}$. In particular, r_1 distinguishes P and P'. By Lemma 4.3.3, every $u \in \mathbb{N}$

that is nested with s is also nested with r_1. Moreover, t is nested with r_1, but not with s, so that $d(\mathbb{N}, r_1) < d(\mathbb{N}, s)$. This contradicts our choice of s.

Therefore $\vec{s} \vee \vec{t} \notin \vec{S}$. Since $\vec{S}$ is submodular, it follows that $\vec{r_2} := \vec{s} \wedge \vec{t} \in \vec{S}$. Moreover, r_2 is nested with every $u \in \mathbb{N} \setminus \{t\}$. This is clear if $\vec{t} \leqslant \vec{u}$ or $\vec{t} \leqslant \overleftarrow{u}$, since $\vec{r_2} \leqslant \vec{t}$. It cannot be that $\overleftarrow{u} \leqslant \vec{t}$, because $\vec{u}, \vec{t} \in P$ and P is consistent. Since $\mathbb{N}$ is nested, only the case $\vec{u} < \vec{t}$ remains. Then, by our choice of $\vec{t}$, u and s are nested and it follows from Lemma 4.3.3 that u and r_2 are also nested. Hence $\mathbb{N}' := (\mathbb{N} \setminus \{t\}) \cup \{r_2\}$ is a nested set of separations.

To see that $\mathbb{N}'$ distinguishes Π_0, it suffices to check that r_2 distinguishes P_1 and P_2. We have $\vec{r_2} \in P_2$ since P_2 is consistent and $\vec{r_2} \leqslant \vec{t} \in P_2$: if $\vec{r_2} = \overleftarrow{t}$, then s and t would be nested. Since $\overleftarrow{r_2} = \overleftarrow{s} \vee \overleftarrow{t}$ and $\overleftarrow{s}, \overleftarrow{t} \in P_1$, we find $\overleftarrow{r_2} \in P_1$. Any element of $\mathbb{N}'$ which is not nested with s lies in $\mathbb{N}$. Since $t \in \mathbb{N} \setminus \mathbb{N}'$ is not nested with s, it follows that $d(\mathbb{N}', s) < d(\mathbb{N}, s)$, contrary to our choice of $\mathbb{N}$ and s. □

4.4 Tangle-tree duality

Our agenda for this section is first to prove Theorem 4.2.2, and then to derive from it Theorem 4.2.1, which as we have seen implies Theorem 4.1.2. Our proof will be an application of the basic tangle-tree duality theorem from [51].

For this we need to introduce the notion of separability, and then prove that submodular separation systems are separable (Lemma 4.4.4). This lemma not only lies at the heart of our proof of Theorem 4.2.2: it will also be central to any other result that asserts a tangle-tree type duality for separation systems $(\vec{S}, \leqslant, ^*)$ that are structurally submodular, but are not so simply as a corollary of the existence of a submodular order function on $\vec{S}$.

A separation $\vec{s} \in \vec{S}$ *emulates* $\vec{r}$ *in* $\vec{S}$ if $\vec{s} \geqslant \vec{r}$ and for every $\vec{t} \in \vec{S} \setminus \{\overleftarrow{r}\}$ with $\vec{t} \geqslant \vec{r}$ we have $\vec{s} \vee \vec{t} \in \vec{S}$. For $\vec{s} \in \vec{S}$, $\sigma \subseteq \vec{S}$ and $\vec{x} \in \sigma$, define

$$\sigma_{\vec{x}}^{\vec{s}} := \{\vec{x} \vee \vec{s}\} \cup \{\vec{y} \wedge \overleftarrow{s} : \vec{y} \in \sigma \setminus \{\vec{x}\}\}.$$

Lemma 4.4.1. *Suppose $\vec{s} \in \vec{S}$ emulates a non-trivial $\vec{r}$ in $\vec{S}$, and let $\sigma \subseteq \vec{S}$ be a star such that $\vec{r} \leqslant \vec{x} \in \sigma$. Then $\sigma_{\vec{x}}^{\vec{s}} \subseteq \vec{S}$ is a star.*

Proof. Note that for every $\vec{y} \in \sigma \setminus \{\vec{x}\}$ we have $\vec{r} \leqslant \overleftarrow{y}$. It is clear that for any two distinct $\vec{u}, \vec{v} \in \sigma_{\vec{x}}^{\vec{s}}$ we have $\vec{u} \leqslant \overleftarrow{v}$, so we only need to show that every element of $\sigma_{\vec{x}}^{\vec{s}}$ is non-degenerate and lies in $\vec{S}$. For every $\vec{u} \in \sigma_{\vec{x}}^{\vec{s}}$ there is a non-degenerate $\vec{t} \in \vec{S}$ with $\vec{r} \leqslant \vec{t}$ such that either $\vec{u} = \vec{t} \vee \vec{s}$ or $\overleftarrow{u} = \vec{t} \vee \vec{s}$.

Let $\vec{t} \in \vec{S}$ be non-degenerate with $\vec{r} \leqslant \vec{t}$. Since $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$, we find $\vec{t} \vee \vec{s} \in \vec{S}$. Assume for a contradiction that $\vec{t} \vee \vec{s}$ was degenerate. Since $\vec{t}$ is non-degenerate, we find that $\vec{t} < \vec{t} \vee \vec{s}$, so that $\vec{t}$ is trivial. But then so is $\vec{r}$, because $\vec{r} \leqslant \vec{t}$. This contradicts our assumption on $\vec{r}$. □

The separation system $\vec{S}$ is *separable* if for all non-trivial and non-degenerate $\vec{r_1}, \overleftarrow{r_2} \in \vec{S}$ with $\vec{r_1} \leqslant \vec{r_2}$ there exists an $\vec{s} \in \vec{S}$ which emulates $\vec{r_1}$ in $\vec{S}$ while simultaneously $\overleftarrow{s}$ emulates $\overleftarrow{r_2}$ in $\vec{S}$.

Given some $\mathcal{F} \subseteq 2^{\vec{U}}$, we say that $\vec{s}$ *emulates* $\vec{r}$ *in* $\vec{S}$ *for* $\mathcal{F}$ if $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$ and for every star $\sigma \subseteq \vec{S} \setminus \{\overleftarrow{r}\}$ with $\sigma \in \mathcal{F}$ and every $\vec{x} \in \sigma$ with $\vec{x} \geqslant \vec{r}$ we have $\sigma_{\vec{x}}^{\vec{s}} \in \mathcal{F}$.

The separation system $\vec{S}$ is $\mathcal{F}$*-separable* if for all non-trivial and non-degenerate $\vec{r_1}, \overleftarrow{r_2} \in \vec{S}$ with $\vec{r_1} \leqslant \vec{r_2}$ and $\{\overleftarrow{r_1}\}, \{\vec{r_2}\} \notin \mathcal{F}$ there exists an $\vec{s} \in \vec{S}$ which emulates $\vec{r_1}$ in $\vec{S}$ for $\mathcal{F}$ while simultaneously $\overleftarrow{s}$ emulates $\overleftarrow{r_2}$ in $\vec{S}$ for $\mathcal{F}$.

Theorem 4.4.2 ([51, Theorem 4.3])**.** *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a separation system. Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars, standard for $\vec{S}$. If $\vec{S}$ is $\mathcal{F}$-separable, then exactly one of the following holds:*

(cl.1) *There exists an $\mathcal{F}$-tangle of S.*

(cl.2) *There exists an S-tree over $\mathcal{F}$.*

In applications of Theorem 4.4.2 it is often easier to split the proof of the main premise, that $\vec{S}$ is $\mathcal{F}$-separable, into two parts: a proof that $\vec{S}$ is separable and one that $\mathcal{F}$ is *closed under shifting* in $\vec{S}$: that whenever $\vec{s} \in \vec{S}$ emulates (in $\vec{S}$) some nontrivial and nondegenerate $\vec{r} \leqslant \vec{s}$ not forced by $\mathcal{F}$, then it does so for $\mathcal{F}$. Indeed, the following is immediate from the definitions:

Lemma 4.4.3. *Let $\vec{U}$ be a universe of separations, $\vec{S} \subseteq \vec{U}$ a separation system, and $\mathcal{F} \subseteq 2^{\vec{U}}$ a set of stars. If $\vec{S}$ is separable and $\mathcal{F}$ is closed under shifting, then $\vec{S}$ is $\mathcal{F}$-separable.* $\square$

It is shown in [52] that if $\vec{U}$ is a universe of separations with an order function, then the sets $\vec{S_k}$ of all separations of order less than some fixed positive integer k are separable for all k, and virtually all the applications of Theorem 4.4.2 that are given in [51] involve a separation system of the form $\vec{S_k}$.

While many applications of the submodularity of an order function use only its structural consequence that motivated our abstract notion of submodularity, the use of submodularity in the proof that $\vec{S_k}$ is separable – see [52, Lemma 3.4] – uses it in a more subtle way. There, the orders of opposite corners of two crossing separations $\vec{s}$ and $\vec{t}$ are compared not with any fixed value of k but with the (possibly distinct) orders of s and t directly. This kind of argument is naturally difficult, if not impossible, to mimic in our set-up.

However, we can prove this nevertheless, choosing a different route. The following lemma is, in essence, the main result of this section:

Lemma 4.4.4. *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a submodular separation system. Then $\vec{S}$ is separable.*

We will actually prove a slightly more general statement about submodular lattices. Let $(L, \vee, \wedge)$ be a lattice and let $M \subseteq L$. Given $x, y \in M$, we say that x *pushes* y if $x \leqslant y$ and for any $z \in M$ with $z \leqslant y$ we have $x \wedge z \in M$. Similarly, we say that x *lifts* y if $x \geqslant y$ and for any $z \in M$ with $z \geqslant y$ we have $x \vee z \in M$. Observe that both of these relations are reflexive and transitive: Every $x \in M$ pushes (lifts) itself and if x pushes (lifts) y and y pushes (lifts) z, then x pushes (lifts) z. We say that M is *strongly separable* if for all $x, y \in M$ with $x \leqslant y$ there exists a $z \in M$ that lifts x and pushes y.

The definitions of lifting, pushing and separable extend verbatim to a separation system within a universe of separations when regarded as a subset of the underlying lattice. These notions are strengthenings of the notions of emulating and separable: If $\vec{s} \in \vec{S}$ lifts $\vec{r} \in \vec{S}$, then $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$, and $\vec{s}$ pushes $\vec{r}$ if and only if $\overleftarrow{s}$ lifts $\overleftarrow{r}$. Similarly, if $\vec{S}$ is strongly separable, then $\vec{S}$ is separable. Lemma 4.4.4 is then an immediate consequence of the following:

Lemma 4.4.5. *Let L be a finite lattice and $M \subseteq L$ submodular. Then M is strongly separable.*

Proof. Call a pair $(a, b) \in M \times M$ *bad* if $a \leqslant b$ and there is no $x \in M$ that lifts a and pushes b. Assume for a contradiction that there was a bad pair and choose one, say (a, b), such that $I(a, b) := \{u \in M : a \leqslant u \leqslant b\}$ is minimal.

We claim that a pushes every $z \in I(a, b) \setminus \{b\}$. Indeed, assume for a contradiction a did not push some such z. By minimality of (a, b), the pair (a, z) is not bad, so there is some $x \in M$ which lifts a and pushes z. By assumption, $x \neq a$ and so by minimality, the pair (x, b) is not bad, yielding a $y \in M$ which lifts x and pushes b. By transitivity, it follows that y lifts a. But then (a, b) is not a bad pair, which is a contradiction. An analogous argument establishes that b lifts every $z \in I(a, b) \setminus \{a\}$.

Since (a, b) is bad, a does not push b, so there is some $x \in M$ with $x \leqslant b$ for which $a \wedge x \notin M$. Similarly, there is a $y \in M$ with $y \geqslant a$ for which $b \vee y \notin M$. Since M is submodular, it follows that $a \vee x, b \wedge y \in M$. Note that $a \vee x, b \wedge y \in I(a, b)$. Furthermore, $x \leqslant a \vee x$ and $a \wedge x \notin M$, so a does not push $a \vee x$. We showed that a pushes every $z \in I(a, b) \setminus \{b\}$, so it follows that $a \vee x = b$. Similarly, we find that $b \wedge y = a$. But then

$$x \vee y = x \vee (a \vee y) = b \vee y \notin M,$$
$$x \wedge y = (x \wedge b) \wedge y = x \wedge a \notin M.$$

This contradicts the submodularity of M. □

As a result we obtain our tangle-tree duality theorem for $\mathcal{F}$-tangles of submodular separation systems, Theorem 4.2.2.

Proof of Theorem 4.2.2. Since $\vec{S}$ is submodular, Lemma 4.4.4 implies that $\vec{S}$ is separable. Since $\mathcal{F}$ is closed under shifting, it follows from Lemma 4.4.3 that $\vec{S}$ is $\mathcal{F}$-separable. The result then follows by Theorem 4.4.2. □

We will now use Theorem 4.2.2 to prove Theorem 4.2.1, which in turn implies Theorem 4.1.2. Recall that we are considering a downclosed subset $Q \subseteq \vec{U}$ of a distributive universe of separations, and a submodular separation system $\vec{S}$ without degenerate elements in $\subseteq \vec{U}$, and we wish to prove a tangle-tree duality theorem for abstract tangles of $\vec{S}$ extending Q. Note that these are precisely the $(\mathcal{T} \cup \mathcal{F}_Q)$-tangles of $\vec{S}$. However, since the family $\mathcal{F}$ in Theorem 4.2.2 is assumed to be a set of stars, we cannot work directly with $\mathcal{T}$. Instead we will work with $\mathcal{T}^*$, the set of stars in $\mathcal{T}$. It will turn out that, since $\vec{S}$ has no degenerate elements, this will not change the set of $\mathcal{T}$-tangles (cf. Lemma 4.4.8). So, we will first show that we can apply Theorem 4.2.2 with $\mathcal{F} = \mathcal{T}_Q$, where $\mathcal{T}_Q := \mathcal{T}^* \cup \mathcal{F}_Q$, and then show that the $\mathcal{T}^*$-tangles are precisely the $\mathcal{T}$-tangles. Theorem 4.2.1 will then follow.

Let us first prove the following simple fact, which will be useful in a few different situations.

Lemma 4.4.6. *Let $\vec{U}$ be a distributive universe of separations. Let $\vec{u}, \vec{v}, \vec{w} \in \vec{U}$. If $\vec{u} \leqslant \vec{v}$ and $\vec{v} \vee \vec{w}$ is co-small, then $\vec{v} \vee (\vec{w} \wedge \overleftarrow{u})$ is co-small.*

Proof. Let $\vec{x} := \vec{v} \vee (\vec{w} \wedge \overleftarrow{u})$. By distributivity of $\vec{U}$

$$\vec{x} = (\vec{v} \vee \vec{w}) \wedge (\vec{v} \vee \overleftarrow{u}) \geqslant (\vec{v} \vee \vec{w}) \wedge (\vec{u} \vee \overleftarrow{u}).$$

Let $\vec{s} := \vec{v} \vee \vec{w}$ and $\vec{t} := \vec{u} \vee \overleftarrow{u}$. Then $\overleftarrow{s} \leqslant \vec{s}$ by assumption and $\overleftarrow{s} \leqslant \overleftarrow{v} \leqslant \vec{t}$. Further $\overleftarrow{t} \leqslant \vec{u} \leqslant \vec{t}$ and $\overleftarrow{t} \leqslant \vec{u} \leqslant \vec{v}$. Therefore

$$\overleftarrow{x} \leqslant \overleftarrow{s} \vee \overleftarrow{t} \leqslant \vec{s} \wedge \vec{t} \leqslant \vec{x}.$$ □

In order to apply Theorem 4.2.2 with $\mathcal{F} = \mathcal{T}_Q$, we need to show that T_Q is closed under shifting.

Lemma 4.4.7. *If $Q \subseteq \vec{U}$ is down-closed and $\vec{U}$ is distributive, then $\mathcal{T}_Q$ is closed under shifting.*

Proof. Let $\vec{r} \in \vec{S}$ non-trivial and non-degenerate with $\{\overleftarrow{r}\} \notin \mathcal{F}$. Let $\vec{s} \in \vec{S}$ emulate $\vec{r}$ in $\vec{S}$, let $\mathcal{T}_Q \ni \sigma \subseteq \vec{S} \setminus \{\overleftarrow{r}\}$ and $\vec{r} \leqslant \vec{x} \in \sigma$. We have to show that $\sigma^{\vec{s}}_{\vec{x}} \in \mathcal{T}_Q$. From Lemma 4.4.1 we know that $\sigma^{\vec{s}}_{\vec{x}}$ is a star, so we only need to verify that $\sigma^{\vec{s}}_{\vec{x}} \in \mathcal{T}^* \cup \mathcal{F}_Q$.

Suppose first that $\sigma \in \mathcal{T}^*$. Let $\vec{w} := \bigvee(\sigma \setminus \{\vec{x}\})$. Applying Lemma 4.4.6 with $\vec{u} = \vec{s}$ and $\vec{v} = \vec{x} \vee \vec{s}$, we see that

$$\bigvee \sigma^{\vec{s}}_{\vec{x}} = (\vec{x} \vee \vec{s}) \vee (\vec{w} \wedge \overleftarrow{s})$$

is co-small. Since $\sigma^{\vec{s}}_{\vec{x}}$ has at most three elements, it follows that $\sigma^{\vec{s}}_{\vec{x}} \in \mathcal{T}$.

Suppose now that $\sigma \in \mathcal{F}_Q$. Then $\sigma = \{\vec{x}\}$ and $\overleftarrow{x} \in Q$. As Q is down-closed, we have $\overleftarrow{x} \wedge \overleftarrow{s} \in Q$. Since $\sigma_{\vec{x}}^{\vec{s}}$ is a star, $\overleftarrow{x} \wedge \overleftarrow{s}$ is non-degenerate and therefore

$$\sigma_{\vec{x}}^{\vec{s}} = \{\vec{x} \vee \vec{s}\} = \{(\overleftarrow{x} \wedge \overleftarrow{s})^*\} \in \mathcal{F}_Q.$$

$\square$

Lemma 4.4.8. *Let $\vec{U}$ be a distributive universe of separations and let $\vec{S} \subseteq \vec{U}$ be a submodular separation system without degenerate elements. Then the $\mathcal{T}^*$-tangles are precisely the abstract tangles.*

Proof. Since $\mathcal{T}^* \subseteq \mathcal{T}$, every abstract tangle is also a $\mathcal{T}^*$-tangle. We only need to show that, conversely, every $\mathcal{T}^*$-tangle in fact avoids $\mathcal{T}$.

For $\sigma \in \mathcal{T}$, let $d(\sigma)$ be the number of pairs $\vec{s}, \vec{t} \in \sigma$ which are not nested. Let O be a consistent orientation of S and suppose O was not an abstract tangle. Choose $\mathcal{T} \ni \sigma \subseteq O$ such that $d(\sigma)$ is minimum and, subject to this, σ is inclusion-minimal. We will show that σ is indeed a star, thus showing that O is not a $\mathcal{T}^*$-tangle.

If σ contained two comparable elements, say $\vec{s} \leqslant \vec{t}$, then $\sigma' := \sigma \setminus \{\vec{s}\}$ satisfies $\sigma' \in \mathcal{T}$, $\sigma' \subseteq O$ and $d(\sigma') \leqslant d(\sigma)$, violating the fact that σ is inclusion-minimal. Hence σ is an antichain. Since $\vec{S}$ has no degenerate elements, it follows from the consistency of O that any two nested $\vec{s}, \vec{t} \in \sigma$ satisfy $\vec{s} \leqslant \overleftarrow{t}$. To show that σ is a star, it thus suffices to prove that any two elements are nested.

Suppose that σ contained two crossing separations, say $\vec{s}, \vec{t} \in \sigma$. By submodularity of $\vec{S}$, at least one of $\vec{s} \wedge \overleftarrow{t}$ and $\overleftarrow{s} \wedge \vec{t}$ lies in $\vec{S}$. By symmetry we may assume that $\vec{r} := \vec{s} \wedge \overleftarrow{t} \in \vec{S}$. Let $\sigma' := (\sigma \setminus \{\vec{s}\}) \cup \{\vec{r}\}$. Since O is consistent, $\vec{r} \leqslant \vec{s}$ and $r \neq s$, it follows that $\vec{r} \in O$ and so $\sigma' \subseteq O$ as well.

Let $\vec{w} = \bigvee(\sigma \setminus \{\vec{t}\})$. As $\vec{t} \vee \vec{w} = \bigvee \sigma$ is co-small, we can apply Lemma 4.4.6 with $\vec{u} = \vec{v} = \vec{t}$ to deduce that $\vec{t} \vee (\vec{w} \wedge \overleftarrow{t})$ is co-small as well. But

$$\vec{t} \vee (\vec{w} \wedge \overleftarrow{t}) = \vec{t} \vee \bigvee_{\vec{x} \in \sigma \setminus \{\vec{t}\}} (\vec{x} \wedge \overleftarrow{t}) \leqslant \bigvee \sigma',$$

so $\bigvee \sigma'$ is also co-small and $\sigma' \in \mathcal{T}$.

We now show that $d(\sigma') < d(\sigma)$. Since s and t cross, while r and t do not, it suffices to show that every $\vec{x} \in \sigma \setminus \{\vec{s}\}$ which is nested with $\vec{s}$ is also nested with $\vec{r}$. But for every such $\vec{x}$ we have $\vec{s} \leqslant \overleftarrow{x}$. Since $\vec{r} \leqslant \vec{s}$, we get $\vec{r} \leqslant \overleftarrow{x}$ as well, showing that r and x are nested. So in fact $d(\sigma') < d(\sigma)$, which is a contradiction. This completes the proof that σ is nested and therefore a star. $\square$

We are now in a position to prove Theorem 4.2.1.

Proof of Theorem 4.2.1. By Lemma 4.4.4, $\vec{S}$ is separable, and by Lemma 4.4.7, $\mathcal{T}_Q$ is closed under shifting. Therefore, by Theorem 4.2.2, there is no S-tree over $\mathcal{T}^* \cup \mathcal{F}_Q$, if and only if S has a T_Q-tangle, that is, a $\mathcal{T}^*$-tangle extending Q.

However, since $\vec{U}$ is distributive and $\vec{S}$ contains no degenerate elements, Lemma 4.4.8 implies that S has a $\mathcal{T}^*$-tangle extending Q if and only if S has an abstract tangle extending Q. $\square$

4.5 Special cases and applications

4.5.1 Tangles in graphs and matroids

We briefly indicate how tangles in graphs and matroids can be seen as special cases of abstract tangles in separation systems. Tangles in graphs and hypergraphs were introduced by Robertson and Seymour in [110], but a good deal of the work is done in the setting of connectivity systems. Geelen, Gerards and Whittle [67] made this more explicit and defined tangles as well as the dual notion of branch-decompositions for connectivity systems, an approach that we will follow.

Let X be a finite set and $\lambda : 2^X \to \mathbb{Z}$ a map assigning integers to the subsets of X such that $\lambda(X \setminus A) = \lambda(A)$ for all $A \subseteq X$ and

$$\lambda(A \cup B) + \lambda(A \cap B) \leqslant \lambda(A) + \lambda(B)$$

for all $A, B \subseteq X$. The pair (X, λ) is then called a *connectivity system.*

Both graphs and matroids give rise to connectivity systems. For a given graph G, we can take $X := E(G)$ and define $\lambda(F)$ as the number of vertices of G incident with edges in both F and $E \setminus F$. Given a matroid M with ground-set X and rank-function r, we take λ to be the connectivity function $\lambda(A) := r(A) + r(X \setminus A) - r(X)$.

Now consider 2^X as a universe of separations with set-inclusion as the partial order and $A^* = X \setminus A$ as involution. For an integer k, the set $\vec{S_k}$ of all sets A with $\lambda(A) < k$ is then a submodular separation system. Let $Q := \{\emptyset\} \cup \{\{x\} \colon x \in X\}$ consist of the empty-set and all singletons of X and note that Q is down-closed.

A *tangle of order* k of (X, λ), as defined in [67], is then precisely an abstract tangle extending Q. It is easy to see that (X, λ) has a *branch-decomposition* of width $<k$ if and only if there exists an S_k-tree over $\mathcal{T}^* \cup \mathcal{F}_Q$. Theorem 4.2.1 then yields the classic duality theorem for tangles and branch-decompositions in connectivity systems, see [110, 67].

4.5.2 Clique separations

We now describe a submodular separation system that is not derived from a submodular order function, and provide a natural set of stars for which Theorem 4.2.2 applies.

Let $G = (V, E)$ be a finite graph and $\vec{U}$ the universe of all separations of G, that is, pairs (A, B) of subsets of V with $V = A \cup B$ such that there is no edge between $A \setminus B$ and $B \setminus A$. Here the partial order is given by $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ and $B \supseteq D$, and the involution is simply $(A, B)^* = (B, A)$. For $(A, B) \in \vec{U}$, we call $A \cap B$ the *separator* of (A, B). It is an *a-b-separator* if $a \in A \setminus B$ and $b \in B \setminus A$. We call $A \cap B$ a *minimal separator* if there exist $a \in A \setminus B$ and $b \in B \setminus A$ for which $A \cap B$ is an inclusion-minimal *a*-*b*-separator.

Recall that a *hole* in a graph is an induced cycle on more than three vertices. A graph is *chordal* if it has no holes.

Theorem 4.5.1 (Dirac [55])**.** *A graph is chordal if and only if every minimal separator is a clique.*

Let $\vec{S}$ be the set of all $(A, B) \in \vec{U}$ for which $G[A \cap B]$ is a clique. We call these the *clique separations.* Note that $\vec{S}$ is closed under involution and therefore a separation system. To avoid trivialities, we will assume that the graph G is not itself a clique. In particular, this implies that $\vec{S}$ contains no degenerate elements.

Lemma 4.5.2. *Let $s, t \in S$. At least three of the four corners of s and t are again in $\vec{S}$. In particular, $\vec{S}$ is submodular.*

Proof. Let $\vec{s} = (A, B)$ and $\vec{t} = (C, D)$. Since $G[A \cap B]$ is a clique and (C, D) is a separation, we must have $A \cap B \subseteq C$ or $A \cap B \subseteq D$, without loss of generality $A \cap B \subseteq C$. Similarly, it follows that $C \cap D \subseteq A$ or $C \cap D \subseteq B$; we assume the former holds. For each corner other than $\vec{s} \wedge \vec{t} = (A \cap C, B \cup D)$, the separator is a subset of either $A \cap B$ or $C \cap D$ and therefore the subgraph it induces is a clique. This proves our claim. $\square$

Suppose that the graph G contains a hole H. Then for every $(A, B) \in \vec{S}$, either $H \subseteq A$ or $H \subseteq B$. In this way, every hole H induces an orientation

$$O_H := \{(A, B) \in \vec{S} \colon H \subseteq B\}$$

of $\vec{S}$. We now describe these orientations as tangles over a suitable set of stars.

Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be the set of all sets $\{(A_1, B_1), \dots (A_n, B_n)\} \subseteq \vec{U}$ for which $G[\bigcap B_i]$ is a clique (note that the graph without any vertices is a clique). As usual, we denote by $\mathcal{F}^*$ the set of all elements of $\mathcal{F}$ which are stars.

Theorem 4.5.3. *Let O be an orientation of S. Then the following are equivalent:*

(cl.1) *O is an $\mathcal{F}^*$-tangle.*

(cl.2) *O is an $\mathcal{F}$-tangle.*

(cl.3) *There exists a hole H with $O = O_H$.*

It is easy to see that every orientation O_H induced by a hole H is an $\mathcal{F}$-tangle. To prove that, conversely, every $\mathcal{F}$-tangle is induced by a hole, we use Theorem 4.5.1 and an easy observation about clique-separators, Lemma 4.5.4 below. The proof that every $\mathcal{F}^*$-tangle is already an $\mathcal{F}$-tangle, the main content of Lemma 4.5.5 below, is similar to the proof of Lemma 4.4.8, but some care is needed to keep track of the separators of two crossing separations.

For a set $\tau \subseteq \vec{U}$, let $J(\tau) := \bigcap_{(A,B)\in\tau} B$ be the intersection of all the right sides of separations in τ, where $J(\emptyset) := V(G)$.

Lemma 4.5.4. *Let τ be a set of clique separations, $J = J(\tau)$ and $K \subseteq J$. Let $a, b \in J \setminus K$. If K separates a and b in $G[J]$, then it separates them in G.*

Proof. We prove this by induction on $|\tau|$, the case $\tau = \emptyset$ being trivial. Suppose now $|\tau| \geqslant 1$ and let $(X, Y) \in \tau$ arbitrary. Put $\tau' := \tau \setminus \{(X, Y)\}$ and $J' := J(\tau')$. Note that $J = J' \cap Y$. Let $G' := G[J']$ and $(X', Y') := (X \cap J', Y \cap J')$.

Then $K \subseteq J'$ and $a, b \in J' \setminus K$. Suppose K did not separate a and b in G' and let $P \subseteq J'$ be an induced a-b-path avoiding K. Since $G'[X' \cap Y']$ is a clique, P has at most two vertices in $X' \cap Y'$ and they are consecutive vertices along P. As $a, b \in Y'$ and (X', Y') is a separation of G', it follows that $P \subseteq Y'$. But then K does not separate a and b in $J = J' \cap Y$, contrary to our assumption.

Hence K separates a and b in G'. By inductive hypothesis applied to τ', it follows that K separates a and b in G. $\square$

Lemma 4.5.5. *Every $\mathcal{F}^*$-tangle is an $\mathcal{F}$-tangle and a regular profile.*

Proof. Let P be an $\mathcal{F}^*$-tangle. It is clear that P contains no co-small separation, since $\{(V, A)\} \in \mathcal{F}^*$ for every co-small $(V, A) \in \vec{S}$. Since P is consistent, it follows that P is in fact down-closed.

We now show that P is a profile. Let $(A, B), (C, D) \in P$ and assume for a contradiction that $(E, F) := ((A, B) \vee (C, D))^* \in P$. Recall that either $C \cap D \subseteq A$ or $C \cap D \subseteq B$.

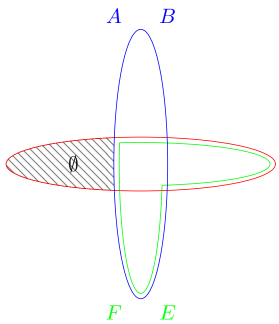

Figure 4.2: The case $C \cap D \subseteq B$

Suppose first that $C \cap D \subseteq B$; this case is depicted in Figure 4.2. Let $(X, Y) := (A, B) \wedge (D, C)$ and note that $X \cap Y \subseteq A \cap B$, so that $(X, Y) \in \vec{S}$. It follows from the consistency of P that $(X, Y) \in P$. Let $\tau := \{(C, D), (E, F), (X, Y)\}$ and observe that $\tau \subseteq P$ is a star. However

$$J(\tau) = D \cap (A \cup C) \cap (B \cup C) = (D \cap B) \cap (A \cup C),$$

which is the separator of (E, F). Since $(E, F) \in \vec{S}$, $G[J(\tau)]$ is a clique, thereby contradicting the fact that P is an $\mathcal{F}^*$-tangle.

Suppose now that $C \cap D \subseteq A$. Let $(X, Y) := (B, A) \wedge (C, D)$ and note that $X \cap Y \subseteq A \cap B$, so that $(X, Y) \in \vec{S}$. Since P is down-closed, it follows that $(X, Y) \in P$. Therefore $\tau := \{(A, B), (E, F), (X, Y)\} \subseteq P$. But τ is a star and

$$J(\tau) = B \cap (A \cup C) \cap (A \cup D) = B \cap (A \cup (C \cap D)) = B \cap A,$$

and so $G[J(\tau)]$ is a clique, which again contradicts our assumption that P is an $\mathcal{F}^*$-tangle. This contradiction shows that P is indeed a profile.

We now prove that for any $\tau \subseteq P$ there exists a star $\sigma \subseteq P$ with $J(\sigma) = J(\tau)$. It follows then, in particular, that P is an $\mathcal{F}$-tangle.

Given $\tau \subseteq P$, choose $\sigma \subseteq P$ with $J(\sigma) = J(\tau)$ so that $d(\sigma)$, the number of crossing pairs of elements of σ, is minimum and, subject to this, σ is inclusion-minimal. Then σ is an antichain: If $(A, B) \leqslant (C, D)$ and both $(A, B), (C, D) \in \sigma$, then $\sigma' := \sigma \setminus \{(A, B)\}$ satisfies $J(\sigma') = J(\sigma)$, thus violating the minimality of σ. Since $\sigma \subseteq P$ and P is consistent, no two elements of σ point away from each other. Therefore, any two nested elements of σ point towards each other. To verify that σ is a star, it suffices to check that σ is nested.

Assume for a contradiction that σ contained two crossing separations (A, B) and (C, D). If $(E, F) := (A, B) \vee (C, D) \in \vec{S}$, obtain σ' from σ by deleting (A, B) and (C, D) and adding (E, F). We have seen above that P is a profile, so $\sigma' \subseteq P$. By Lemma 4.3.3, every element of $\sigma \setminus \{(A, B), (C, D)\}$ that is nested with both (A, B) and (C, D) is also nested with (E, F). Since σ' misses the crossing pair $\{(A, B), (C, D)\}$, it follows that $d(\sigma') < d(\sigma)$. But $J(\sigma') = J(\sigma)$, contradicting the minimality of σ.

Hence it must be that $(E, F) \notin \vec{S}$, so $A \cap B \not\subseteq C$ and $C \cap D \not\subseteq A$. Therefore $(X, Y) := (A, B) \wedge (D, C) \in \vec{S}$. Let $\sigma' := (\sigma \setminus \{(A, B)\}) \cup \{(X, Y)\}$. Note that $(X, Y) \leqslant (A, B) \in P$, so $\sigma' \subseteq P$. Moreover $Y \cap D = (B \cup C) \cap D = B \cap D$, since $C \cap D \subseteq B$. Therefore $J(\sigma') = J(\sigma)$. As mentioned above, any $(U, W) \in \sigma \setminus \{(A, B)\}$ that is nested with (A, B) satisfies $(A, B) \leqslant (W, U)$. Therefore $(X, Y) \leqslant (A, B) \leqslant (W, U)$, so (X, Y) is also nested with (U, W). It follows that $d(\sigma') < d(\sigma)$, which is a contradiction. This completes the proof that σ is nested and therefore a star. $\square$

Proof of Theorem 4.5.3. (i) $\to$ (ii): See Lemma 4.5.5.

(ii) $\to$ (iii): Let O be an $\mathcal{F}$-tangle and $J := J(O)$. We claim that there is a hole H of G with $H \subseteq J$. Such a hole then trivially satisfies $O_H = O$.

Assume there was no such hole, so that $G[J]$ is a chordal graph. Since O is $\mathcal{F}$-avoiding, $G[J]$ itself cannot be a clique, so there exists a minimal set $K \subseteq J$ separating two vertices $a, b \in J \setminus K$ in $G[J]$. By Theorem 4.5.1, K induces a clique in G. By Lemma 4.5.4, K separates a and b in G, so there exists a separation $(A, B) \in \vec{S}$ with $A \cap B = K$, $a \in A \setminus B$ and $b \in B \setminus A$. As O orients $\vec{S}$, it must contain one of $(A, B), (B, A)$, say without loss of generality $(A, B) \in O$. But then $J \subseteq B$, contrary to $a \in J$. This proves our claim.

(iii) $\to$ (i): We have $H \subseteq J(O_H)$, so $J(O)$ does not induce a clique. Since every star $\sigma \subseteq O$ has $J(O) \subseteq J(\sigma)$ there is no star $\sigma \subseteq O$ such that $G[J(\sigma)]$ is a clique, and so O is $\mathcal{F}^*$-avoiding. Furthermore, O_H is clearly consistent, and so O is an $\mathcal{F}^*$-tangle. □

The upshot of Theorem 4.5.3 is that a hole in a graph, although a very concrete substructure, can be regarded as a tangle. This is in line with our general narrative, set forth e.g. in [52, 39, 40], that tangles arise naturally in very different contexts, and underlines the expressive strength of abstract separation systems and tangles.

What does our abstract theory then tell us about the holes in a graph? The results we will derive are well-known and not particularly deep, but it is nonetheless remarkable that the theory of abstract separation systems, emanating from the theory of highly connected substructures of a graph or matroid, is able to express such natural facts about holes.

Firstly, by Lemma 4.5.5, every hole induces a profile of S. Hence Theorem 4.3.1 applies and yields a nested set $\mathbb{N}$ of clique-separations distinguishing all holes which can be separated by a clique. This is similar to, but not the same as, the decomposition by clique separators of Tarjan [121]: the algorithm in [121] essentially produces a maximal nested set of clique separations and leaves 'atoms' that do not have any clique separations, whereas our tree set merely distinguishes the holes and leaves larger pieces that might allow further decomposition.

Secondly, we can apply Theorem 4.2.2 to find the structure dual to the existence of holes. It is clear that $\mathcal{F}^*$ is standard, since $\mathcal{F}^*$ contains $\{(V, A)\}$ for every $(V, A) \in \vec{S}$.

Lemma 4.5.6. *$\mathcal{F}^*$ is closed under shifting.*

Proof. Let $(X, Y) \in \vec{S}$ emulate a non-trivial $(U, W) \in \vec{S}$ with $\{(W, U)\} \notin \mathcal{F}^*$, let $\sigma = \{(A_i, B_i) \colon 0 \leqslant i \leqslant n\} \subseteq \vec{S}$ with $\sigma \in \mathcal{F}^*$ and $(U, W) \leqslant (A_0, B_0)$. Then

$$\sigma' := \sigma^{(X,Y)}_{(A_0,B_0)} = \{(A_0 \cup X, B_0 \cap Y)\} \cup \{(A_i \cap Y, B_i \cup X) \colon 1 \leqslant i \leqslant n\}.$$

By Lemma 4.4.1, $\sigma' \subseteq \vec{S}$ is a star. We need to show that $G[J(\sigma')]$ is a clique. Let $(A, B) := \bigvee_{i \geqslant 1}(A_i, B_i)$ and note that $(A, B) \leqslant (B_0, A_0)$, since σ is a star. Then

$$(B, A) \wedge (V, B_0) = (B, B_0) \in \vec{U}.$$

But $G[B \cap B_0] = G[J(\sigma)]$ is a clique, so in fact $(B, B_0) \in \vec{S}$. Since $(U, W) \leqslant (A_0, B_0) \leqslant (B, A)$, we see that $(U, W) \leqslant (B, B_0)$. As (X, Y) emulates (U, W) in $\vec{S}$, we find that $(E, F) := (X, Y) \vee (B, B_0) \in \vec{S}$. It thus follows that

$$J(\sigma') = (X \cup B) \cap (Y \cap B_0) = E \cap F$$

and so $G[J(\sigma')]$ is indeed a clique. Therefore $\sigma' \in \mathcal{F}^*$. □

Theorem 4.5.7. *Let G be a graph. Then the following are equivalent:*

(cl.1) *G has a tree-decomposition in which every part is a clique.*

(cl.2) *There exists an S-tree over $\mathcal{F}^*$.*

(cl.3) *S has no $\mathcal{F}^*$-tangle.*

(cl.4) *G is chordal.*

Proof. (i) → (ii): Let $(T, \mathcal{V})$ be a tree-decomposition of G in which every part is a clique. For adjacent $s, t \in T$, let $T_{s,t}$ be the component of $T - st$ containing t and let $V_{s,t}$ be the union of all V_u with $u \in T_{s,t}$. Define $\alpha : \vec{E}(T) \to \vec{U}$ as $\alpha(s,t) := (V_{t,s}, V_{s,t})$. Then $\alpha(s,t) = \alpha(t,s)^*$. The separator of $\alpha(s,t)$ is $V_s \cap V_t$, which is a clique by assumption. Hence (T, α) is in fact an S-tree. It is easy to see that $\alpha(F_t)$ is a star for every $t \in T$ and that $J(\alpha(F_t)) = V_t$. Therefore (T, α) is an S-tree over $\mathcal{F}^*$.

(ii) → (i): Given an S-tree (T, α) over $\mathcal{F}^*$, define $V_t := J(\alpha(F_t))$ for $t \in T$. It is easily verified that $(T, \mathcal{V})$ is a tree-decomposition of G. Each V_t is then a clique, since $\alpha(F_t) \in \mathcal{F}$.

(ii) ↔ (iii): Follows from Theorem 4.2.2, since $\mathcal{F}^*$ is standard for $\vec{S}$ and closed under shifting by Lemma 4.5.6.

(iii) ↔ (iv): Follows from Theorem 4.5.3. □

The equivalence of (i) and (iv) is a well-known characterization of chordal graphs that goes back to a theorem Gavril [65] which identifies chordal graphs as the intersection graphs of subtrees of a tree.

4.5.3 Tangle-tree duality in cluster analysis

Let us now apply Theorem 4.2.2 to a generic scenario in cluster analysis [39, 40], where V is thought of as a data set, S is a set of certain 'natural' bipartitions of V, and we are interested in certain $\mathcal{F}$-tangles as 'clusters'. The idea is that clusters should be described by these $\mathcal{F}$-tangles in the same way as the vertex set of a large grid in a graph is captured by the graph tangle τ it induces: although every oriented separation $\vec{s}$ in τ points to most of the vertices of the grid, the cluster can be 'fuzzy' in that these are not the same points for every $\vec{s} \in \tau$. Indeed, there need not be a single vertex to which all the $\vec{s} \in \tau$ point.

To mimic this idea, we want to choose $\mathcal{F}$ so that, whenever we consider just a few separations in S, any $\mathcal{F}$-tangle τ of S must orient these so that they all point to at least some m (say) points in V, while we do not require that the intersection of all the sets B for $(A, B) \in \tau$ must be large (or even non-empty).

Formally, then, let $\vec{U}$ be the universe of all oriented bipartitions (A, B) of some non-empty set V, including $(\emptyset, V)$ and $(V, \emptyset)$, with $\overleftarrow{s} = (B, A)$ for $\vec{s} = (A, B)$ and $(A, B) \wedge (C, D) := (A \cap C, B \cup D)$, and let $\vec{S} \subseteq \vec{U}$ be any submodular separation system in $\vec{U}$. Let $1 \leq m \in \mathbb{N}$ and $2 \leq n \in \mathbb{N} \cup \{\infty\}$. For these m and n, define

$$\mathcal{F}_m := \{\, F \subseteq \vec{U} : \Big| \bigcap_{(A,B) \in F} B \Big| < m \,\}$$

and

$$\mathcal{F}_m^n := \{\, F \in \mathcal{F}_m : |F| < n \,\}.$$

There is only one small separation in $\vec{U}$, the separation $(\emptyset, V)$. Hence regardless of what S may be, it has no trivial separation other than $(\emptyset, V)$. Since $\{(V, \emptyset)\} \in \mathcal{F}_m^n$ for all m and n, this makes $\mathcal{F}_m^n$ standard for S.

Recall that, for any $\mathcal{F} \subseteq 2^{\vec{S}}$, we write $\mathcal{F}^*$ for the set of stars in $\mathcal{F}$.

Lemma 4.5.8. *Let $1 \leq m \in \mathbb{N}$ and $2 \leq n \in \mathbb{N} \cup \{\infty\}$. For $\mathcal{F} = \mathcal{F}_m^n$, the $\mathcal{F}^*$-tangles of S are precisely its $\mathcal{F}$-tangles.*

Proof. Since $\mathcal{F}^* \subseteq \mathcal{F}$, it is clear that $\mathcal{F}$-tangles are $\mathcal{F}^*$-tangles. To show the converse, suppose there is an $\mathcal{F}^*$-tangle τ that fails to be an $\mathcal{F}$-tangle, because it contains some $F \in \mathcal{F}$ as a subset.

Clearly $F \in \mathcal{F} \setminus \mathcal{F}^*$, so F contains two crossing separations $\vec{r}$ and $\vec{s}$. Since $\vec{S}$ is submodular, one of their opposite corners $\vec{r} \wedge \overleftarrow{s}$ and $\vec{s} \wedge \overleftarrow{r}$ lies in $\vec{S}$; let us assume $\vec{r'} := \vec{r} \wedge \overleftarrow{s}$ does. Since $\vec{r'} \leq \vec{r} \in \tau$, the consistency of τ implies that $\vec{r'}$ lies in τ (rather than its inverse $\overleftarrow{r'}$). Indeed, this follows from the definition of consistency if $r' \neq r$. But if $r' = r$ then by $\vec{r'} \leq \vec{r}$ either $\vec{r'} = \vec{r} \in \tau$ as desired, or $\vec{r'} < \vec{r} = \overleftarrow{r'}$ with $\vec{r'}$ small but $\overleftarrow{r'} \in \tau$. Since $(\emptyset, V)$ is the only small separation in $\vec{S}$ and is in fact trivial, the consistency of τ once more implies that $\vec{r'} \in \tau$.

Let F' be obtained from F by replacing $\vec{r}$ with $\vec{r'}$. Note that $\bigcap_{(A,B)\in F} B = \bigcap_{(A',B')\in F'} B'$ has remained unchanged: although we replaced the set B from $\vec{r} = (A, B)$ in the first intersection with the bigger set B' from $\vec{r'} = (A', B')$ in the second, the additional $B' \setminus B$ has empty intersection with the set D from $\vec{s} = (C, D)$, and therefore does not increase the second intersection. Hence our assumption of $F \in \mathcal{F} = \mathcal{F}^n_m$ implies that also $F' \in \mathcal{F}^n_m = \mathcal{F}$.

Note that while $\vec{r}$ and $\vec{s}$ crossed, $\vec{r'}$ and $\vec{s}$ are nested; indeed, $\{\vec{r'}, \vec{s}\}$ is a star. Moreover, replacing an element of this star by a smaller separation will yield another star; in particular, it cannot result in another pair of crossing separations. This means that iterating the above uncrossing procedure of replacing in F an element $\vec{r} \in \tau$ with a smaller separation $\vec{r'} \in \tau$ in a way that keeps F in $\mathcal{F}$ will end after at most $\binom{|F|}{2}$ steps: for every 2-set $\{\vec{r}, \vec{s}\} \subseteq F$ we will consider only once in this iterated process a pair $\{\vec{r'}, \vec{s'}\}$ where $\vec{r'}$ is either $\vec{r}$ or a replacement of $\vec{r}$, and $\vec{s'}$ is either $\vec{s}$ or a replacement of $\vec{s}$.

Since the above process turns every pair of crossing separations from F into a 2-star of separations, and a set of separations is a star as soon as all its 2-subsets are stars, the set we turn F into will be a star in $\mathcal{F}$, an element of $\mathcal{F}^*$. As it will also still be a subset of τ, this contradicts our assumption that τ is an $\mathcal{F}^*$-tangle. $\square$

Lemma 4.5.9. *Let $1 \leq m \in \mathbb{N}$ and $2 \leq n \in \mathbb{N} \cup \{\infty\}$. The set $\mathcal{F}^*$ of stars in $\mathcal{F} = \mathcal{F}^n_m$ is closed under shifting.*

Proof. Suppose that $\vec{s} \in \vec{S}$ emulates in $\vec{S}$ some nontrivial $\vec{r}$ not forced by $\mathcal{F}$. We have to show that for every star $\sigma \subseteq \vec{S} \setminus \{\overleftarrow{r}\}$ with $\sigma \in \mathcal{F}^*$ and every $\vec{x} \in \sigma$ with $\vec{x} \geqslant \vec{r}$ we have $\sigma' := \sigma^{\vec{s}}_{\vec{x}} \in \mathcal{F}^*$.

Let $\vec{s} = (U, W)$, and for $(A, B) \in \sigma$ write $(A', B') \in \sigma'$ for the separation that (A, B) shifts to: if $(A, B) = \vec{x}$ then $(A', B') := (A \cup U, B \cap W)$, while if $(A, B) \in \sigma \setminus \{\vec{x}\}$ then $(A', B') := (A \cap W, B \cup U)$. From these explicit representations of the elements of σ' it is clear that

$$\bigcap_{(A',B')\in\sigma'} B' = \bigcap_{(A,B)\in\sigma} B' \subseteq \bigcap_{(A,B)\in\sigma} B\,,$$

since $B' \setminus B \subseteq U$ for every $(A, B) \in \sigma \setminus \{\vec{x}\}$ while $U \cap B' = \emptyset$ for $(A, B) = \vec{x}$, so that the overall intersection of all the B' equals that of all the B. And since these sets did not change, nor did their cardinality: as $\sigma \in \mathcal{F} = \mathcal{F}^n_m$ we also have $\sigma' \in \mathcal{F}^n_m = \mathcal{F}$. By Lemma 4.4.1, this implies $\sigma' \in \mathcal{F}^*$ as desired. $\square$

Together with Lemmas 4.5.8 and 4.5.9, Theorem 4.2.2 implies our last main theorem:

Theorem 4.5.10. *Let $\vec{S} \subseteq \vec{U}$ be a submodular separation system, let $1 \leq m \in \mathbb{N}$ and $2 \leq n \in \mathbb{N} \cup \{\infty\}$, and let $\mathcal{F} = \mathcal{F}^n_m$. Then exactly one of the following two statements holds:*

(cl.1) *S has an $\mathcal{F}$-tangle;*

(cl.2) *There exists an S-tree over $\mathcal{F}^*$.* $\square$

4.5.4 Phylogenetic trees from tangles of circle separations

Finally, let us describe an application of Theorem 4.1.1 and Theorem 4.2.3 in biology. Let V be a set, which we think of as a set of species, or of (possibly unknown) organisms, or of DNA samples. Our aim is to find their Darwinian 'tree of life': a way of dividing V recursively into ever-smaller subsets so that the leaves of this division tree correspond to the individual elements of V.

This tree can be formalized by starting with a root node labelled $\emptyset$ at level 0, a unique node labelled V at level 1, and then recursively adding children to each node, labelled $A \subseteq V$, say, corresponding to the subsets of A into which A is divided and labelling these children with the subsets to which they correspond. Furthermore, let us label the edges from level $k-1$ to k by k.

Every edge e of this tree naturally defines a bipartition of V, to which we assign the label k of e as it 'order'. Species that are fundamentally different are then separated by a bipartition of low order, while closely related but distinct species are only separated by bipartitions of higher order.

If we draw that tree in the plane, the leaves – and hence V – will be arranged in a circle, C say. The bipartitions $\{A, B\}$ of V defined by the edges of the tree are given by *circle separations*: A and B are covered by disjoint half-open segments of C whose union is C.

The way in which this tree is found in practice is roughly as follows. One first defines a metric on V in which species u, v are far apart if they differ in many respects. Then one applies some clustering algorithm, such as starting with the singletons $\{v\}$ for $v \in V$ as tiny clusters and then successively amalgamating close clusters (in terms of this distance function) into bigger ones. The bipartitions of V corresponding to pitching a single cluster against the rest of V will be nested and can therefore be represented as edge-separations of a tree as above, which is then output by the algorithm. If we draw the tree in the plane, the bipartitions then become circle separations of V as earlier.

A known problem with this approach is that, since every cluster found in this process defines a bipartition of V that ends up corresponding to an edge of the final tree, inaccuracies in the clustering process immediately affect this tree in an irreversible way. Bryant and Moulton [31] have suggested a more careful clustering process which produces not necessarily a tree but an outerplanar graph G on V, together with a set of particularly important bipartitions of V that are not necessarily nested but are still circle separations of V with respect to the outer face boundary of G. The task then is to select from the set S of these bipartition a nested subset that defines the desired phylogenetic tree, or perhaps to generate such a set from S in some other suitable way such as adding corners of crossing separations already selected.

This is where tangles can help: in a general way, but also in a rather specific way that finds the desired nested set from a set of circular separations of V.

Let us just briefly indicate the general way in which tangles can be used for finding phylogenetic trees in a novel way [39], without the need for any distance-based clustering. As input we need a collection of subsets of V to be used as similarity criteria, such as the set of species $v \in V$ that can fly or lay eggs, the set of DNA molecules that have base T in position 137, or the set of those organisms that respond to some test in a certain way. Then we define an order function on all the bipartitions of V, assigning low order to those bipartitions that do not cut accross many of our criteria sets so as to split them nearly in half. For example, we might count for $s = \{A, B\}$ the number of triples (a, b, c) such that $a \in A$ and $b \in B$ and both a and b satisfy (are elements of) the criterion c.

This order function is easily seen to be submodular on the universe $\vec{U}$ of all oriented bipartitions of V [40]. For suitable $\mathcal{F}$ whose $\mathcal{F}$-tangles are profiles, e.g. the $\mathcal{F} = \mathcal{F}_m^n$ with $n > 3$ considered in Section 4.5.3, we can then compute the (canonical) tree of tangles as in [48], or an S-tree over $\mathcal{F}$ as in [52, 51] if there are no tangles. In the first case, the tree of tangles for $\mathcal{F}_2^n$

has a good claim to be the phylogenetic tree for the species in V, see [39].

In the concrete scenario of [31], we further have the following specific application of tangles as studied in this paper. Let S be the set of circle separations of V, taken with respect to its circular ordering found by the current algorithm of [31]. We can now define $\mathcal{F}$-tangles on this S just as we did earlier on the set of all bipartitions of V, and consider the same order function as earlier.

This order function is not, however, submodular on our restricted set S: this would require that S is not just a separation system but a universe of separations, i.e., that corner separations of elements of S are again in S – which is not always the case. In particular, we do not get a tree-of-tangles theorem for S from [48], or a tangle-tree duality theorem from [52, 51]. [3]

However, we do get a tree-of-tangles theorem for S as a corollary of Theorem 4.1.1, and a tangle-tree duality theorem as a corollary of Theorem 4.2.3. For this we need the following easy lemma. Let $\vec{U}$ be the universe of all oriented bipartitions of V (see Section 4.5.3), equipped with any submodular order function under which $\{\emptyset, V\}$ has order 0. Let S be the set of all circle separations of V with respect to some fixed cyclic ordering of V.

Lemma 4.5.11. *For every $k > 0$, the set S_k of all circle separations of order $< k$ is submodular.*

Proof. Consider two oriented circle separations $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$ of V. Clearly, $\vec{r} \vee \vec{s}$ is again a circle separation unless the circle segments representing A and C are disjoint and both segments that join them on the circle meet V. In that case, however, the union of the segments representing B and D is the entire circle, so $\vec{r} \wedge \vec{s} = (\emptyset, V)$, which is a circle separation.

As $(\emptyset, V)$ has order $0 < k$ by assumption, and our order function is submodular on the set of all bipartitions of V, this implies that S_k is submodular: given $\vec{r}, \vec{s} \in \vec{S_k}$, either both $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{r}$ are in $\vec{S}$ and hence one of them is in $\vec{S_k}$, or one of them is $(\emptyset, V)$ or $(V, \emptyset)$ and therefore in $\vec{S_k}$. $\square$

Consider any $1 \leq m \in \mathbb{N}$ and $3 < n \in \mathbb{N} \cup \{\infty\}$. Let $\mathcal{F} = \mathcal{F}^n_m$ be as defined in Section 4.5.3. Here is our first tree-of-tangles theorem for circle separations:

Theorem 4.5.12. *For every $k > 0$, the set S_k of circle separations of V of order $< k$ contains a tree set of separations that distinguishes all the $\mathcal{F}$-tangles of S_k.*

Proof. In order to apply Theorem 4.1.1, we have to show that all $\mathcal{F}$-tangles of S_k are abstract tangles, i.e., that they contain no triple $(\vec{r}, \vec{s}, \vec{t})$ with $\vec{r} \vee \vec{s} \vee \vec{t} = (V, \emptyset)$ (which is the unique co-small separation in $\vec{U}$). But any such triple lies in $\mathcal{F}$ for the values of m and n we specified, so no $\mathcal{F}$-tangle of S_k contains it.

Lemma 4.5.11 and Theorem 4.1.1 thus imply the result. $\square$

Applying very recent work of Elbracht, Kneip and Teegen [57], we can unify the assertions of Theorem 4.5.12 over all k, as follows. Let us say that an element s of S *distinguishes* two orientations ϱ, τ of subsets of S if these orient s differently, i.e., if s has orientations $\vec{s} \in \varrho$ and $\overleftarrow{s} \in \tau$. If s has minimum order amongst the separations in S that distinguish ϱ from τ, we say that s distinguishes ϱ and τ *efficiently*. Similarly, a set $T \subseteq S$ *distinguishes* a set $\mathcal{T}$ of orientations of subsets of S *efficiently* if for every pair of distinct $\varrho, \tau \in \mathcal{T}$ there exists an $s \in T$ that distinguishes ϱ from τ efficiently.

Orientations ϱ and τ as above are called *distinguishable* if they are distinguished by some $s \in S$. Note that orientations ϱ of S_k and τ of S_ℓ for $k \leq \ell$ are indistinguishable if and only if $\varrho = \tau \cap \vec{S_k}$.

Here is our tree-of-tangles theorem for circle separations of mixed order:

[3] This is not to say that no submodular order function on S exists that returns the sets S_k we are interested in as sets $S_{k'}$ for some other k'. One can indeed construct such a function, but it is neither obvious nor natural.

Theorem 4.5.13. [57] *The set S of all circle separations of V contains a tree set that efficiently distinguishes all the distinguishable $\mathcal{F}$-tangles of subsets S_k of S.*

Elbracht, Kneip and Teegen [57] showed that the tree set in Theorem 4.5.13 can in fact be chosen *canonical*, i.e., so that every separation system automorphism (see [44]) of $\vec{S}$ acts on $\vec{T}$ as a set of automorphisms of T.

Finally, our tangle-tree duality theorem for circle separations:

Theorem 4.5.14. *For every $k > 0$, the set S_k of circle separations of V of order $< k$ satisfies exactly one of the following two assertions:*

(cl.1) *S_k has an $\mathcal{F}$-tangle;*

(cl.2) *There exists an S_k-tree over $\mathcal{F}^*$.*

Proof. Apply Lemma 4.5.11 and Theorem 4.2.3. □

Acknowledgement

Geoff Whittle has informed us that his student Jasmine Hall has obtained explicit proofs of Theorems 4.5.12 and 4.5.14 by re-working the theory of graph and matroid tangles given in [67, 110].

Chapter 5

Directed path-decompositions

5.1 Introduction

Given a tree T and vertices $t_1, t_2 \in V(T)$ let us denote by t_1Tt_2 the unique path in T between t_1 and t_2. Given a graph $G = (V, E)$ a *tree-decomposition* is a pair $(T, \mathcal{V})$ consisting of a tree T, together with a collection of subsets of vertices $\mathcal{V} = \{V_t \subseteq V(G) : t \in V(T)\}$, called *bags*, such that:

- $V(G) = \bigcup_{t \in T} V_t$;
- For every edge $e \in E(G)$ there is a t such that e lies in V_t;
- $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever $t_2 \in V(t_1Tt_3)$.

The *width* of this tree-decomposition is the quantity $\max\{|V_t| - 1 : t \in V(T)\}$ and its *adhesion* is $\max\{|V_t \cap V_{t'}| : (t, t') \in E(T)\}$. Given a graph G its *tree-width* $\mathrm{tw}(G)$ is the smallest k such that G has a tree-decomposition of width k. A *haven of order k* in a graph G is a function β which maps each set $X \subseteq V(G)$ of fewer than k vertices to some connected component of $G - X$ such that, for each such X and Y, $\beta(X)$ and $\beta(Y)$ touch. That is, either $\beta(X)$ and $\beta(Y)$ share a vertex, or there is an edge between $\beta(X)$ and $\beta(Y)$. Seymour and Thomas [118] showed that these two notions are dual to each other, in the following sense:

Theorem 5.1.1 (Seymour and Thomas)**.** *A graph has tree-width $\geqslant k - 1$ if and only if it has a haven of order $\geqslant k$.*

When T is a path we say that $(T, \mathcal{V})$ is a *path-decomposition* and the *path-width* $\mathrm{pw}(G)$ of a graph G is the smallest k such that G has a path-decomposition of width k. Whilst the path-width of a graph is clearly an upper bound for its tree-width, these parameters can be arbitrarily far apart. Bienstock, Robertson, Seymour and Thomas [21] showed that, as for tree-width, there is a structure like a haven, called a *blockage of order k*, such that the path-width of a graph is equal to the order of the largest blockage.

Explicitly, for any subset X of vertices in a graph G let us write $\partial(X)$ for the set of $v \in X$ which have a neighbour in $V(G) - X$. Two subsets $X_1, X_2 \subseteq V(G)$ are *complementary* if $X_1 \cup X_2 = V(G)$ and $\partial(X_1) \subseteq X_2$ (or, equivalently, $\partial(X_2) \subseteq X_1$). A blockage of order k is a set $\mathcal{B}$ such that:

- each $X \in \mathcal{B}$ is a subset of $V(G)$ with $|\partial(X)| \leqslant k$;
- if $X \in \mathcal{B}$, $Y \subseteq X$ and $|\partial(Y)| \leqslant k$, then $Y \in \mathcal{B}$;

- if X_1 and X_2 are complementary and $|X_1 \cap X_2| \leqslant k$, then $\mathcal{B}$ contains exactly one of X_1 and X_2.

Theorem 5.1.2 (Bienstock, Robertson, Seymour and Thomas)**.** *A graph has path-width $\geqslant k$ if and only if it has a blockage of order $\geqslant k$.*

The authors used the theorem to show the following result.

Theorem 5.1.3 (Bienstock, Robertson, Seymour and Thomas)**.** *For every forest F, every graph with pathwidth $\geqslant |V(F)| - 1$ has a minor isomorphic to F.*

In particular, since large binary trees have large path-width, it is a simple corollary of this result that if X is a graph, then the tree-width of graphs in

$$\mathrm{Forb}_{\preccurlyeq}(X) := \{G : X \text{ is not a minor of } G\}$$

is bounded if and only if X is a forest. A more direct proof of this fact, without reference to blockages, was later given by Diestel [41].

There have been numerous suggestions as to the best way to extend the concept of tree-width to digraphs. For example, directed tree-width [83, 82, 107], D-width [116], Kelly width [81] or DAG-width [105, 20]. In some of these cases generalizations of Theorem 5.1.1 have been studied, the hope being to find some structure in the graph whose existence is equivalent to having large 'width'. However these results have either not given an exact equivalence [83, Theorem 3.3], or only apply to certain classes of digraphs [116, Corollary 3].

In contrast, if we wish to decompose a digraph in a way that the model digraph is a directed path, it is perhaps clearer what a sensible notion of a 'directed path-decomposition' should be. The following definition appears in Barát [16] and is attributed to Robertson, Seymour and Thomas. We note that it agrees with the definition of a DAG-decomposition, in the case where the DAG is a directed path.

Given a digraph D a *directed path-decomposition* is a pair $(P, \mathcal{V})$ consisting of a path P, say with $V(P) = \{t_1, t_2, \dots, t_n\}$, together with a collection of subsets of vertices $\mathcal{V} = \{V_i \subseteq V(D) : i \in [n]\}$ such that:

- $V(D) = \bigcup_{t \in V(P)} V_t$;
- if $i < j < k$, then $V_i \cap V_k \subseteq V_j$;
- for every edge $e = (x, y) \in E(D)$ there exists $i \leqslant j$ such $x \in V_i$ and $y \in V_j$.

The *width* of a directed path-decomposition is $\max\{|V_i| - 1 : i \in [n]\}$ and its *adhesion* is $\max\{|V_i \cap V_{i+1}| : i \in [n-1]\}$. Given a digraph D its *directed path-width* $\mathrm{dpw}(D)$ is the smallest k such that D has a directed path-decomposition of width k.

Motivated by Theorem 5.1.2, Barát [16] defined a notion of a blockage in a digraph and showed that if the directed path width is at most $k-1$ then there is no 'blockage of order k'. However he suggested that it was unlikely that the existence of such a 'blockage of order k' was equivalent to having directed path-width at least k, at least for this particular notion of a blockage.

One of the problems with working with digraphs is that many of the tools developed for the theory of tree-decompositions of graphs do not work for digraphs. However, we noticed that in the case of directed path-decompositions some of the most fundamental tools do work almost exactly as in the undirected case. More precisely, Bellenbaum and Diestel [17] explicitly extracted a lemma from the work of Thomas [124], and used it to give short proofs of two theorems: The first, Theorem 5.1.1, and the second a theorem of Thomas on the existence of linked tree-decompositions. In [85] Kim and Seymour prove a similar theorem on the existence of linked directed path-decompositions for semi-complete digraphs. Their proof uses a tool analogous to the key lemma of Bellenbaum and Diestel.

This suggested that perhaps a proof of Theorem 5.1.2 could be adapted, to give an analogue for directed path-width. Indeed, in this paper we show that this is the case. Influenced by ideas from [51, 52] and [6] we will define a notion of *diblockage*[1] in terms of orientations of the set of directed separations of a digraph. We will use generalizations of ideas and tools developed by Diestel and Oum [51, 52] to prove the following.

Theorem 5.1.4. *A digraph has directed path-width $\geqslant k-1$ if and only if it has a diblockage of order $\geqslant k$.*

We also use these ideas to give a result in the spirit of Theorem 5.1.3. An *arborescence* is a digraph in which there is a unique vertex u, called the *root*, such that for every other vertex v there is exactly one directed walk from u to v. Equivalently, an arborescence is formed by taking a rooted tree T and orienting each edge in the tree away from the root. A forest of arborescences is a digraph in which each component is an arborescence. There is a generalisation of the notion of a minor to digraphs called a *butterfly minor*, whose precise definition we will defer until Section 5.4.

Theorem 5.1.5. *For every forest of arborescences F, every digraph D with directed path-width $\geqslant |V(F)|-1$ has a butterfly minor isomorphic to F.*

Kim and Seymour [85] considered the following property of a directed path-decomposition $(P,\mathcal{V})$:

$$\text{If } |V_k| \geqslant t \text{ for every } i \leqslant k \leqslant j, \text{ then there exists a collection of } t$$
$$\text{vertex-disjoint directed paths from } V_j \text{ to } V_i.$$

A directed path-decomposition $(P,\mathcal{V})$ satisfying this condition is said to be *linked*. Kim and Seymour showed that every semi-complete digraph D has a linked directed path-decomposition of width $\mathrm{dpw}(D)$ satisfying the above condition. We prove a slight generalisation of this, which in particular extends the result of Kim and Seymour to arbitrary digraphs.

Theorem 5.1.6. *Every digraph D has a linked directed path-decomposition of width $dpw(D)$.*

The paper is structured as follows. In Section 5.2 we define directed path-decompositions and introduce our main tool of shifting. In section 5.3 we define our notion of a directed blockage and prove Theorem 5.1.4. In Section 5.4 we prove Theorem 5.1.5. Finally, in Section 5.5 we discuss linked directed path-decompositions and prove Theorem 5.1.6.

5.2 Directed path-decompositions

Following the ideas of Diestel and Oum [51] it will be more convenient for us to rephrase the definition of a directed path-decomposition in terms of directed separations. Given a digraph $D=(V,E)$, a pair (A,B) of subsets of V is a *directed separation* if $A\cup B=V$ and there is no edge $(x,y)\in E$ with $x\in B\setminus A$ and $y\in A\setminus B$. Equivalently, every directed path which starts in B and ends in A must meet $A\cap B$.

For brevity, since in this paper we usually be considering digraphs, when the context is clear we will refer to directed separations simply as separations. The *order* of a separation (A,B),

[1] A precise definition will be given in Section 5.2.

which we will denote by $|A, B|$, is $|A \cap B|$ and we will write $\overrightarrow{S}_k$ for the set of separations of order $< k$ and define $\overrightarrow{S} := \bigcup_k \overrightarrow{S}_k$.

We define a partial order on $\overrightarrow{S}$ by

$$(A, B) \leqslant (C, D) \text{ if and only if } A \subseteq C \text{ and } B \supseteq D.$$

We will also define two operations $\wedge$ and $\vee$ such that, for $(A, B), (C, D) \in \overrightarrow{S}$

$$(A, B) \wedge (C, D) = (A \cap C, B \cup D) \text{ and } (A, B) \vee (C, D) = (A \cup C, B \cap D).$$

Lemma 5.2.1. *If (A, B) and (C, D) are separations then so are $(A, B) \wedge (C, D)$ and $(A, B) \vee (C, D)$.*

Proof. Let us prove the claim for $(A, B) \wedge (C, D)$, the proof for $(A, B) \vee (C, D)$ is similar. Firstly, since $A \cup B = V$ and $C \cup D = V$ it follows that

$$(A \cap C) \cup (B \cup D) = (A \cup B \cup D) \cap (C \cup B \cup D) = V \cap V = V.$$

Secondly, let $(x, y) \in E$ with $y \in (A \cap C) \setminus (B \cup D)$. Then, $y \in A \setminus B$ and $y \in C \setminus D$ and so, since (A, B) and (C, D) are separations, it follows that $x \notin B \setminus A$ and $x \notin D \setminus C$. Hence, $x \notin (B \cup D) \setminus (A \cap C)$ and so $(A \cap C, B \cup D)$ is a separation. □

Given some subset $\overrightarrow{S}' \subseteq \overrightarrow{S}$, we define an *$\overrightarrow{S}'$-path* to be a pair (P, α) where P is a path with vertex set $V(P) = \{t_1, t_2, \dots t_n\}$ and $\alpha : E(P) \to \overrightarrow{S}'$ is such that if $1 \leqslant i < j \leqslant n-1$ then $\alpha(t_i, t_{i+1}) \leqslant \alpha(t_j, t_{j+1})$. Note that, if we consider P as a directed path, with edges (t_i, t_{i+1}) for each i, then α preserves the natural order on the edges of P. In this way we can view this notion as a generalisation of the S-trees of Diestel and Oum [51].

We claim that, if $(P, \mathcal{V})$ is a directed path-decomposition then for every $1 \leqslant i \leqslant n-1$ the following is a separation

$$(\bigcup_{j \leqslant i} V_j, \bigcup_{j \geqslant i+1} V_j).$$

Indeed, by definition $\bigcup_\ell V_\ell = V(D)$ and if (x, y) were an edge from $\bigcup_{j \geqslant i+1} V_j \setminus \bigcup_{j \leqslant i} V_j$ to $\bigcup_{j \leqslant i} V_j \setminus \bigcup_{j \geqslant i+1} V_j$ then clearly there could be no $\ell < \ell'$ with $x \in V_\ell$ and $y \in V_{\ell'}$. In this way $(P, \mathcal{V})$ gives an $\overrightarrow{S}$-path by letting

$$\alpha(t_i, t_{i+1}) = (\bigcup_{j \leqslant i} V_j, \bigcup_{j \geqslant i+1} V_j).$$

Conversely, if (P, α) is an $\overrightarrow{S}$-path let us write $(A_i, B_i) := \alpha(t_i, t_{i+1})$, and let $B_0 = V = A_{n+1}$. Note that, if $i < j$ then $(A_i, B_i) \leqslant (A_j, B_j)$ and so $A_i \subseteq A_j$ and $B_i \supseteq B_j$.

Lemma 5.2.2. *Let (P, α) be an $\overrightarrow{S}$-path and A_i, B_i be as above. For $1 \leqslant i \leqslant n$ let $V_i = A_i \cap B_{i-1}$. Then $(P, \mathcal{V})$ is a directed path-decomposition.*

Proof. Firstly, we claim that for $1 \leqslant j \leqslant n-1$, $\bigcup_{i=1}^{j} V_i = A_j$. Indeed, $V_1 = A_1 \cap B_0 = A_1$ and, if the claim holds for $j-1$ then

$$\bigcup_{i=1}^{j} V_i = A_{i-1} \cup V_i = A_{i-1} \cup (A_i \cap B_{i-1}) = A_i,$$

since $A_{i-1} \cup B_{i-1} = V$ and $A_{i-1} \subseteq A_i$. As we will use it later, we note that a similar argument shows that $\bigcup_{i=j}^{n} V_i = B_{j-1}$. Hence,

$$\bigcup_{i=1}^{n} V_i = A_{n-1} \cup V_n = A_{n-1} \cup (A_n \cap B_{n-1}) = A_{n-1} \cup B_{n-1} = V.$$

Next, suppose that $i < j < k$. We can write

$$V_i \cap V_k = (A_i \cap B_{i-1}) \cap (A_k \cap B_{k-1}).$$

Since $(A_i, B_i) \leqslant (A_j, B_j)$, it follows that $A_i \subseteq A_j$ and since $(A_{j-1}, B_{j-1}) \leqslant (A_{k-1}, B_{k-1})$ and so $B_{k-1} \subseteq B_{j-1}$. Hence

$$V_i \cap V_k = (A_i \cap B_{i-1}) \cap (A_k \cap B_{k-1}) \subseteq A_j \cap B_{j-1} = V_j.$$

Finally, suppose for a contradiction there is some edge (x, y) such that there is no $i \leqslant j$ with $x \in V_i$ and $y \in V_j$. Since $\bigcup V_i = V$ there must be $i > j$ with $x \in V_i$ and $y \in V_j$. Pick such a pair with i minimal. It follows that $x \in V_i \setminus \bigcup_{k=1}^{i-1} V_k = (A_i \cap B_{i-1}) \setminus A_{i-1} \subseteq B_{i-1} \setminus A_{i-1}$ by the previous claim. Also, $y \in V_j \setminus \bigcup_{k \geqslant i} V_k = (A_j \cap B_{j-1}) \setminus B_{i-1} \subseteq A_{i-1} \setminus B_{i-1}$ since $A_j \subseteq A_{i-1}$. However, this contradicts the fact that (A_{i-1}, B_{i-1}) is a separation. $\square$

In this way the two notions are equivalent. We say that the *width* of an $\overrightarrow{S}$-path (P, α) is the width of the path-decomposition $(P, \mathcal{V})$ given by $V_i = A_i \cap B_{i-1}$. The following observation will be useful.

Lemma 5.2.3. *If (P, α) is an $\overrightarrow{S}$-path of width $< k - 1$ with $\alpha(t_i, t_{i+1}) = (A_i, B_i)$ and $B_0 = V = A_n$, then*

- *(P, α) is an $\overrightarrow{S}_k$-path;*
- *$(A_i, B_{i-1}) \in \overrightarrow{S}_k$ for each $1 \leqslant i \leqslant n$.*

Proof. For the first we note that, since $(A_{i-1}, B_{i-1}) \leqslant (A_i, B_i)$, it follows that $B_{i-1} \supseteq B_i$ and hence

$$|A_i, B_i| = |A_i \cap B_i| \leqslant |A_i \cap B_{i-1}| < k.$$

For the second, note that, since $A_{i-1} \subseteq A_i$, $B_{i-1} \backslash A_i \subseteq B_{i-1} \backslash A_{i-1}$, and since $A_{i-1} \cup B_{i-1} = V$, $A_i \backslash B_{i-1} \subseteq A_{i-1} \backslash B_{i-1}$. Hence, there is no edge from $B_{i-1} \backslash A_i$ to $A_i \backslash B_{i-1}$ and so $(A_i, B_{i-1}) \in \overrightarrow{S}$. Finally, since $|A_i \cap B_{i-1}| < k$, $(A_i, B_{i-1}) \in \overrightarrow{S}_k$.

$\square$

5.2.1 Shifting an $\overrightarrow{S}_k$-path

One of the benefits of thinking of a directed path-decomposition in terms of the separations it induces, rather than the bags, is that it allows one to easily describe some of the operations that one normally performs on tree-decompositions.

Given an $\overrightarrow{S}$-path (P, α) with $V(P) = \{t_1, \ldots, t_n\}$, let us write $\alpha(t_i, t_{i+1}) = (A_i, B_i)$. We call (A_1, B_1) the *initial leaf separation* of (P, α) and (A_{n-1}, B_{n-1}) the *terminal leaf separation*. It will be useful to have an operation which transforms an $\overrightarrow{S}$-path into one with a given initial/terminal leaf separation.

Let (P, α) be as above and let $(A_i, B_i) \leqslant (X, Y) \in \overrightarrow{S}$. The *up-shift of* (P, α) *onto* (X, Y) *with respect to* (A_i, B_i) is the $\overrightarrow{S}$-path (P', α') where $V(P') = \{t'_i, t'_{i+1}, \ldots, t'_n\}$ and α' is given by

$$\alpha'(t'_j, t'_{j+1}) := (A_j, B_j) \vee (X, Y) = (A_j \cup X, B_j \cap Y).$$

It is simple to check that, if $i \leqslant j < k \leqslant n-1$ then $\alpha'(t'_j, t'_j + 1) \leqslant \alpha'(t'_k, t'_{k+1})$, and so (P', α') is an $\overrightarrow{S}$-path. We note that the initial leaf separation of (P', α') is $\alpha'(t'_i, t'_{i+1}) = (A_i, B_i) \vee (X, Y) = (X, Y)$

Similarly if $(X,Y) \leqslant (A_i,B_i)$ the *down-shift of* (P,α) *onto* (X,Y) *with respect to* (A_i,B_i) is the $\vec{S}$-path (P',α') where $V(P') = \{t'_1, t'_2, \dots, t'_{i+1}\}$ and α' is given by

$$\alpha'(t'_j, t'_{j+1}) := (A_j,B_j) \wedge (X,Y) = (A_j \cap X, B_j \cup Y).$$

We note that the terminal leaf separation of (P',α') is $\alpha'(t'_i,t'_{i+1}) = (A_i,B_i) \wedge (X,Y) = (X,Y)$.

If (P,α) is an $\vec{S}_k$-path and (P',α') is an up/down-shift of (P,α) then, whilst (P',α') is an $\vec{S}$-path, it is not always the case that it will also be an $\vec{S}_k$-path, since the order of some of the separations could increase. We note that if (A,B) and (C,D) are separations then $\wedge$ and $\vee$ satisfy the following equality

$$|A,B| + |C,D| = |A \cup C, B \cap D| + |A \cap C, B \cup D|. \tag{5.2.1}$$

Given a pair of separations $(A,B) \leqslant (C,D) \in \vec{S}$ let us denote by

$$\lambda\big((A,B),(C,D)\big) := \min\{|X,Y| \,:\, (X,Y) \in \vec{S} \text{ and } (A,B) \leqslant (X,Y) \leqslant (C,D)\}.$$

We say that (X,Y) is *up-linked* to (A,B) if $(A,B) \leqslant (X,Y)$ and $|X,Y| = \lambda\big((A,B),(X,Y)\big)$. Similarly (X,Y) is *down-linked* to (A,B) if $(X,Y) \leqslant (A,B)$ and $|X,Y| = \lambda\big((X,Y),(A,B)\big)$.

Lemma 5.2.4. *Let (P,α) be an $\vec{S}_k$-path with $V(P) = \{t_1,t_2,\dots,t_n\}$, with $\alpha(t_j,t_{j+1}) = (A_j,B_j)$ for each j. If (X,Y) is up-linked to (A_i,B_i) then the up-shift of (P,α) onto (X,Y) with respect to (A_i,B_i) is an $\vec{S}_k$-path and if (X,Y) is down-linked to (A_i,B_i) then the down-shift of (P,α) onto (X,Y) with respect to (A_i,B_i) is an $\vec{S}_k$-path.*

Proof. By the discussion above it is clear that both are $\vec{S}$-paths, so it remains to show that the set of separations in the paths lie in $\vec{S}_k$. We will show the first claim, the proof of the second follows along similar lines. Let (P',α') be the up-shift of (P,α) onto (X,Y) with respect to (A_i,B_i), with $P' = \{t'_i, t'_{i+1}, \dots, t'_n\}$.

Let $i \leqslant j \leqslant n-1$. We wish to show that $\alpha'(t'_j,t'_{j+1}) = (A_j \cup X, B_j \cap Y)$ is in $\vec{S}_k$. We note that since $(A_i,B_i) \leqslant (A_j,B_j)$ and $(A_i,B_i) \leqslant (X,Y)$, it follows that

$$(A_i,B_i) \leqslant (A_j \cap X, B_j \cup Y) \leqslant (X,Y)$$

and so, since (X,Y) is up-linked to (A_i,B_i),

$$|A_j \cap X, B_j \cup Y| \geqslant |X,Y|.$$

Therefore, by (5.2.1),

$$|A_j \cup X, B_j \cap Y| \leqslant |A_j,B_j| < k.$$

Hence $\alpha'(t'_j,t'_{j+1}) = (A_j \cup X, B_j \cap Y) \in \vec{S}_k$. □

Recall that the width of an $\vec{S}$-path (P,α) is $\max_{1 \leqslant i \leqslant n} |A_i \cap B_{i-1}| - 1$. We would like to claim that, apart from the bag at the initial leaf in an up-shift, or the bag at the terminal leaf in a down-shift, shifting does not increase the size of the bags.

Again, this will not be true for general shifts however, if the assumptions of Lemma 5.2.4 hold, then it will hold. The proof of the fact follows the proof of [52, Lemma 6.1], which itself plays the role of the key lemma of Thomas from [17]. The fact that this lemma remains true for directed path decompositions is what allows us to prove our Theorems 5.1.4 and 5.1.6.

Lemma 5.2.5. *Let (P, α) be an $\vec{S}_k$-path with $V(P) = \{t_1, t_2, \dots, t_n\}$, let $\alpha(t_j, t_{j+1}) = (A_j, B_j)$ for each j, and let $B_0 = V = A_n$. Let $\omega_j = |A_j \cap B_{j-1}|$ be the size of the bags for each j, and let $i \in [n]$ be fixed. Suppose that (X, Y) is up-linked to (A_i, B_i) and (P', α') is the up-shift of (P, α) onto (X, Y) with respect to (A_i, B_i), with $\alpha'(t'_j, t'_{j+1}) = (A'_j, B'_j)$ for each j, and $B'_{i-1} = V = A'_n$. If $\omega'_j = |A'_j \cap B'_{j-1}|$, then for each $i + 1 \leqslant j \leqslant n$, $\omega'_j \leqslant \omega_j$.*

Similarly, suppose that (X, Y) is down-linked to (A_i, B_i) and (P', α') is the down-shift of (P, α) onto (X, Y) with respect to (A_i, B_i), with $\alpha'(t'_j, t'_{j+1}) = (A'_j, B'_j)$ for each j, and $B'_0 = V = A'_{i+1}$. If $\omega'_j = |A'_j \cap B'_{j-1}|$, then for each $1 \leqslant j \leqslant i$, $\omega'_j \leqslant \omega_j$.

Proof. Again, we will just prove the first statement, as the proof of the second is analogous. Recall that $(A'_j, B'_j) = (A_j \cup X, B_j \cap Y)$ for each $i \leqslant j \leqslant n$. Hence, for $i + 1 \leqslant j \leqslant n$

$$\omega'_j = |(A_j \cup X) \cap (B_{j-1} \cap Y)| = |A_j \cup X, B_{j-1} \cap Y|.$$

Here we have used the fact that (A_j, B_{j-1}) is a separation, by Lemma 5.2.3, to deduce that $(A_j \cup X, B_{j-1} \cap Y)$ is also a separation.

Now, since $(A_i, B_i) \leqslant (A_{j-1}, B_{j-1}) \leqslant (A_j, B_j)$, it follows that $(A_i, B_i) \leqslant (A_j, B_{j-1})$. Therefore, since also $(A_i, B_i) \leqslant (X, Y)$,

$$(A_i, B_i) \leqslant (A_j \cap X, B_{j-1} \cup Y) \leqslant (X, Y).$$

Hence, since (X, Y) is up-linked to (A_i, B_i),

$$|A_j \cap X, B_{j-1} \cup Y| \geqslant |X, Y|.$$

Therefore, by (5.2.1),

$$\omega'_j = |A_j \cup X, B_{j-1} \cap Y| \leqslant |A_j, B_{j-1}| = \omega_j.$$

□

5.3 ω-Diblockages

In [52] the structures which are dual to the existence of S-trees are defined as orientations of the set of separations in a graph. That is a subset of the separations which, for each separation given by an unordered pair $\{A, B\}$, contains exactly one of (A, B) or (B, A). Heuristically, one can think of these orientations as choosing, for each separation $\{A, B\}$ one of the two sides A or B to designate as 'large'. This idea generalises in some way the concept of tangles introduce by Robertson and Seymour [110].

In our case, since the directed separations we consider already have a defined 'direction', we will define an orientation of the set of directed separations to be just a bipartition of the set of directed separations. However we will still think of a bipartition $\vec{S}_k = \mathcal{O}^+ \dot\cup \mathcal{O}^-$ as designating for each directed separation $(A, B) \in \vec{S}_k$ one side as being 'large', the side B when $(A, B) \in \mathcal{O}^+$ and the side A when $(A, B) \in \mathcal{O}^-$. Our notion of a directed blockage will then be defined as some way to make these choices for each $(A, B) \in \vec{S}_k$ in a consistent manner.

Let us make the preceding discussion more explicit. We define a *partial orientation* of $\vec{S}_k$ to be a pair of disjoint subsets $\mathcal{O} = (\mathcal{O}^+, \mathcal{O}^-)$ such that $\mathcal{O}^+, \mathcal{O}^- \subseteq \vec{S}_k$. A partial orientation is an *orientation* if $\mathcal{O}^+ \dot\cup \mathcal{O}^- = \vec{S}_k$. Given a partial orientation $\mathcal{P}$ of $\vec{S}_k$ let us write

$$\vec{S}_{\mathcal{P}} = \vec{S}_k \setminus (\mathcal{P}^+ \cup \mathcal{P}^-).$$

We say a partial orientation $\mathcal{P} = (\mathcal{P}^+, \mathcal{P}^-)$ is *consistent* if

- if $(A,B) \in \mathcal{P}^+$, $(A,B) \geqslant (C,D) \in \vec{S}_k$ then $(C,D) \in \mathcal{P}^+$;
- if $(A,B) \in \mathcal{P}^-$, $(A,B) \leqslant (C,D) \in \vec{S}_k$ then $(C,D) \in \mathcal{P}^-$.

In the language of the preceeding discussion this formalises the intuitive idea that if B is the large side of (A,B) and $B \subseteq D$ then D should be the large side of (C,D) and similarly if A is the large side of (A,B) and $A \supseteq C$ then C should be the large side of (C,D).

An orientation $\mathcal{O}$ *extends* a partial orientation $\mathcal{P}$ if $\mathcal{P}^+ \subseteq \mathcal{O}^+$ and $\mathcal{P}^- \subseteq \mathcal{O}^-$. Given $\omega \geqslant k$, let us define $\mathcal{P}_\omega = (\mathcal{P}^+_\omega, \mathcal{P}^-_\omega)$ by

$$\mathcal{P}^+_\omega = \{(A,B) \in \vec{S}_k \,:\, |A| < \omega\} \text{ and } \mathcal{P}^-_\omega = \{(A,B) \in \vec{S}_k \,:\, |B| < \omega\}.$$

In order to ensure that $\mathcal{P}^+_\omega \cap \mathcal{P}^-_\omega = \emptyset$, and so $\mathcal{P}_\omega$ is a partial orientation, we will insist that the digraph D we are considering has at least $2\omega - k$ many vertices.

An *ω-diblockage* (of $\vec{S}_k$) is an orientation $\mathcal{O} = (\mathcal{O}^+ \cup \mathcal{O}^-)$ of $\vec{S}_k$ such that:

- $\mathcal{O}$ extends $\mathcal{P}_\omega$;
- $\mathcal{O}$ is consistent;
- if $(A,B) \in \mathcal{O}^+$ and $(A,B) \leqslant (C,D) \in \mathcal{O}^-$ then $|B \cap C| \geqslant \omega$.

We will show, for $\omega \geqslant k$, a duality between the existence of an $\vec{S}_k$-path of width $< \omega - 1$ and that of an ω-diblockage of $\vec{S}_k$. In the language of tree-decompositions, an $\vec{S}_k$-path of width $< \omega - 1$ is a directed path-decomposition of width $< \omega - 1$ in which all the *adhesion sets* have size $< k$, where the adhesion sets in a tree-decomposition $(T, \mathcal{V})$ are the sets $\{V_t \cap V_{t'} \,:\, (t,t') \in E(T)\}$. Tree-decompositions of undirected graphs with adhesion sets of bounded size have been considered by Diestel and Oum [52] and Geelen and Joeris [66].

We require one more definition for the proof. Given a partial orientation $\mathcal{P}$ of $\vec{S}_k$ and $\omega \geqslant k$, we say that an $\vec{S}_k$-path (P, α) with $V(P) = \{t_1, \ldots, t_n\}$ and $\alpha(t_j, t_{j+1}) = (A_j, B_j)$ for each j is *$(\omega, \mathcal{P})$-admissable* if

- for each $2 \leqslant i \leqslant n-1$, $\omega_i := |A_i \cap B_{i-1}| < \omega$;
- $(A_1, B_1) \in \mathcal{P}^+ \cup \mathcal{P}^+_\omega$;
- $(A_{n-1}, B_{n-1}) \in \mathcal{P}^- \cup \mathcal{P}^-_\omega$;

When $\omega = k$ we call an ω-diblockage of $\vec{S}_k$ a *diblockage of order k*. In this way, when $\omega = k$ the following theorem implies Theorem 5.1.4.

Theorem 5.3.1. *Let $\omega \geqslant k \in \mathbb{N}$ and let $D = (V,E)$ be a digraph with $|V| \geqslant 2\omega - k$. Then exactly one of the following holds:*

- *D has an $\vec{S}_k$-path of width $< \omega - 1$;*
- *There is an ω-diblockage of $\vec{S}_k$.*

Proof. We will instead prove a stronger statement. We claim that for every consistent partial orientation $\mathcal{P}$ of $\vec{S}_k$ which extends $\mathcal{P}_\omega$ exactly one of the following holds:

- either there exists an $(\omega, \mathcal{P})$-admissable $\vec{S}_k$-path; or
- there is an ω-diblockage of $\vec{S}_k$ extending $\mathcal{P}$.

Note that, $\mathcal{P}_\omega$ is a consistent partial orientation of $\overrightarrow{S}_k$ which extends $\mathcal{P}_\omega$, and an $(\omega, \mathcal{P}_\omega)$-admissable $\overrightarrow{S}_k$-path is an $\overrightarrow{S}_k$-path of width $< \omega - 1$. Hence the theorem follows from the above claim applied to $\mathcal{P} = \mathcal{P}_\omega$.

Let us first show that both cannot happen. Suppose for contradiction that there exists an $(\omega, \mathcal{P})$-admissable $\overrightarrow{S}_k$-path (P, α) with $V(P) = \{t_1, \ldots, t_n\}$, $\alpha(t_i, t_{i+1}) = (A_i, B_i)$ for each i and an ω-diblockage $\mathcal{O}$ of $\overrightarrow{S}_k$ extending $\mathcal{P}$.

Since $\mathcal{O}$ extends $\mathcal{P}$, and (P, α) is $(\omega, \mathcal{P})$-admissable, it follows that $(A_1, B_1) \in \mathcal{P}^+ \cup \mathcal{P}^+_\omega \subseteq \mathcal{O}^+$ and $(A_{n-1}, B_{n-1}) \in \mathcal{P}^- \cup \mathcal{P}^-_\omega \subseteq \mathcal{O}^-$. Let $j = \max\{t : (A_t, B_t) \in \mathcal{O}^+\}$. By the previous statement, $1 \leqslant j < n-1$, and so, since $\mathcal{O}$ is an orientation of $\overrightarrow{S}_k$, $(A_{j+1}, B_{j+1}) \in \mathcal{O}^-$. However, $|A_{j+1} \cap B_j| < \omega$, contradicting the assumption that $\mathcal{O}$ is an ω-diblockage of $\overrightarrow{S}_k$.

We will prove the statement by induction on $|\overrightarrow{S}_\mathcal{P}|$. Suppose that $|\overrightarrow{S}_\mathcal{P}| = 0$, in which case $\mathcal{P}$ is a consistent orientation of $\overrightarrow{S}_k$. Suppose that $\mathcal{P}$ is not an ω-diblockage. Since $\mathcal{P}^+_\omega \subseteq \mathcal{P}^+$ and $\mathcal{P}^-_\omega \subseteq \mathcal{P}^-$ there must exist a pair $(A, B) \leqslant (C, D)$ with $(A, B) \in \mathcal{P}^+$, $(C, D) \in \mathcal{P}^-$ and $|B \cap C| < \omega$. Then, we can then form an $(\omega, \mathcal{P})$-admissable $\overrightarrow{S}_k$-path as follows: Let $P = \{v_1, v_2, v_3\}$ and let $\alpha(v_1, v_2) = (A, B)$ and $\alpha(v_2, v_3) = (C, D)$. It is a simple check that (P, α) is a $\mathcal{P}$-admissable $\overrightarrow{S}_k$-tree.

So, let us suppose that $|\overrightarrow{S}_\mathcal{P}| > 0$, and that there is no ω-diblockage $\mathcal{O}$ of $\overrightarrow{S}_k$ extending $\mathcal{P}$. There exists some separation $(A, B) \in \overrightarrow{S}_k \setminus (\mathcal{P}^+ \cup \mathcal{P}^-)$. Let us choose $(C, D) \leqslant (A, B)$ minimal with $(C, D) \in \overrightarrow{S}_k \setminus (\mathcal{P}^+ \cup \mathcal{P}^-)$ and $(A, B) \leqslant (E, F)$ maximal with $(E, F) \in \overrightarrow{S}_k \setminus (\mathcal{P}^+ \cup \mathcal{P}^-)$.

We claim that $\mathcal{P}_1 = (\mathcal{P}^+ \cup (C, D), \mathcal{P}^-)$ and $\mathcal{P}_2 = (\mathcal{P}^+, \mathcal{P}^- \cup (E, F))$ are both consistent partial orientations of $\overrightarrow{S}_k$. Indeed, by minimality of (C, D) every separation $(U, V) < (C, D)$ is in $\mathcal{P}^+ \cup \mathcal{P}^-$ and since $(U, V) < (C, D) \notin \mathcal{P}^+ \cup \mathcal{P}^-$ by the consistency of $\mathcal{P}$ it follows that $(U, V) \in \mathcal{P}^+$. Similarly $(U, V) \in \mathcal{P}^-$ for all $(E, F) < (U, V)$. Hence both $\mathcal{P}_1$ and $\mathcal{P}_2$ are consistent partial orientations of $\overrightarrow{S}_k$. Furthermore $|\overrightarrow{S}_{\mathcal{P}_1}|, |\overrightarrow{S}_{\mathcal{P}_2}| < |\overrightarrow{S}_\mathcal{P}|$. Therefore, we can apply the induction hypothesis to both $\mathcal{P}_1$ and $\mathcal{P}_2$.

Since an ω-diblockage of $\overrightarrow{S}_k$ extending $\mathcal{P}_1$ or $\mathcal{P}_2$ also extends $\mathcal{P}$, we may assume that there exists an $(\omega, \mathcal{P}_1)$-admissable $\overrightarrow{S}_k$-path (P_1, α_1) and an $(\omega, \mathcal{P}_2)$-admissable $\overrightarrow{S}_k$-path (P_2, α_2). Furthermore, we may assume that they are not $(\omega, \mathcal{P})$-admissable, and so the initial leaf separation of (P_1, α_1) is (C, D) and the terminal leaf separation of (P_2, α_2) is (E, F). Let us pick a separation $(C, D) \leqslant (X, Y) \leqslant (E, F)$ such that $|X, Y| = \lambda\big((C, D), (E, F)\big)$. Note that, (X, Y) is up-linked to (C, D) and down-linked to (E, F).

Let (P'_1, α'_1) be the up-shift of (P_1, α_1) onto (X, Y) and let (P'_2, α'_2) be the down-shift of (P_2, α_2) onto (X, Y). Note that the initial leaf separation of (P'_1, α'_1) and the terminal leaf separation of (P'_2, α'_2) are both (X, Y). We form $(\hat{P}, \hat{\alpha})$ by taking $\hat{P}$ to be the path formed by identifying the terminal leaf of P'_2 with the initial leaf of P'_1, with $\hat{\alpha}$ defined to be α'_1 on P'_1 and α'_2 on P'_2. We claim that $(\hat{P}, \hat{\alpha})$ is a $\mathcal{P}$-admissable $\overrightarrow{S}_k$-path. Let us write $V(\hat{P}) = \{\hat{t}_1, \hat{t}_2, \ldots, \hat{t}_{\hat{n}}\}$, with $\hat{\alpha}(\hat{t}_j, \hat{t}_{j+1}) = (\hat{A}_j, \hat{B}_j)$ for each j and $\hat{B}_0 = \hat{A}_{\hat{n}} = V(D)$.

For each $j \neq 1, \hat{n}$, the bag $\hat{A}_j \cap \hat{B}_{j-1}$ is the shift of some non-leaf bag in (P_1, α_1) or (P_2, α_2). Since these were $(\omega, \mathcal{P}_1)$ and $(\omega, \mathcal{P}_2)$-admissable respectively, the size of the bag was less than ω. Therefore, since (X, Y) is up-linked to (C, D) and down-linked to (E, F), by Lemma 5.2.5 $|\hat{A}_j \cap \hat{B}_{j-1}| \leqslant \omega$.

Finally, consider the separations $(\hat{A}_1, \hat{B}_1)$ and $(\hat{A}_{\hat{n}-1}, \hat{B}_{\hat{n}-1})$. If we denote by (U, V) the initial leaf separation in (P_2, α_2) then $(U, V) \in \mathcal{P}^+ \cup \mathcal{P}^+_0$, since (P_2, α_2) is $(\omega, \mathcal{P}_2)$-admissable. Then, $(\hat{A}_1, \hat{B}_1) = (U, V) \wedge (X, Y) \leqslant (U, V)$. We note that, since $\mathcal{P}$ is consistent, $\mathcal{P}^+$ is down-closed, as is $\mathcal{P}^+_0$ by inspection, and so it follows that $(\hat{A}_1, \hat{B}_1) \in \mathcal{P}^+ \cup \mathcal{P}^+_0$. A similar argument shows that $(\hat{A}_{\hat{n}-1}, \hat{B}_{\hat{n}-1}) \in \mathcal{P}^- \cup \mathcal{P}^-_0$.

□

5.4 Finding an arborescence as a butterfly minor

In directed graphs, it is not clear what the best way to generalise the minor operation from undirected graphs. One suggestion (see for example Johnson, Robertson and Seymour [83]), is that of *butterfly minors*. We say an edge $e = (u, v)$ in a digraph D is *contractible* if either $d^-(v) = 1$ or $d^+(u) = 1$ where d^- and d^+ are the in- and out-degree respectively. We say a digraph D' is a butterfly minor of D, which we write $D' \preccurlyeq D$, if D' can be obtained from D by a sequence of vertex deletions, edge deletions and contractions of contractible edges.

We say an $\overrightarrow{S}$-path (P, α) with $V(P) = \{t_1, t_2, \ldots, t_n\}$ and $\alpha(t_i, t_{i+1}) = (A_i, B_i)$ is a *partial* $\overrightarrow{S}$*-path of width* $< k$ if for each $2 \leqslant i \leqslant n-1$, $\omega_i := |A_i \cap B_{i-1}| \leqslant k$ and $\omega_n = |B_{n-1}| \leqslant k$. That is, it is a path-decomposition in which each bag except the first has size $\leqslant k$. Note that a partial $\overrightarrow{S}$-path of width $< k$ is necessarily an $\overrightarrow{S}_{k+1}$-path.

Let $\overrightarrow{S}'_{k+1} \subseteq \overrightarrow{S}$ be the set of (A, B) such that $(A, B) = \alpha(t_1, t_2)$ for some partial $\overrightarrow{S}$-path of width $< k$. Note that, since the second bag in a partial $\overrightarrow{S}$-path of width $< k$ has size at most k, it follows that $|A, B| \leqslant k$ and hence $\overrightarrow{S}'_{k+1} \subseteq \overrightarrow{S}_{k+1}$.

Theorem 5.1.5. *For every forest of arborescences F, every digraph D with directed path-width $\geqslant |V(F)| - 1$ has a butterfly minor isomorphic to F.*

Proof. We may assume without loss of generality that F is in fact an arborescence. Therefore, every vertex in F apart from the root v_0 has exactly one in-neighbour, and there is some ordering of the vertices $V(F) = v_0, v_1, \ldots, v_n$ such that every v_i has no in-neighbours in $\{v_{i+1}, \ldots, v_n\}$. Furthermore, without loss of generality we may assume that the digraph D is weakly connected.

Let us define (C_0, D_0) to be a $\leqslant$-minimal separation of D such that:

- $|C_0, D_0| = 0$;
- $(C_0, D_0) \in \overrightarrow{S}'_{n+1}$.

Note that, since (P, α) with $V(P) = \{t_1, t_2\}$ and $\alpha(t_1, t_2) = (V, \emptyset)$ is a partial $\overrightarrow{S}$-path of width $< n$, at least once such separation exists. Let $x^0_0 \in C_0 \setminus D_0$, which is non-empty as $\mathrm{dpw}(D) \geqslant n$.

We shall construct inductively x^i_i and (C_i, D_i) for $1 \leqslant i \leqslant n$ where (C_i, D_i) is a $\leqslant$-minimal separation satisfying the following properties:

- $(C_i, D_i) \leqslant (C_{i-1}, D_{i-1} \cup \{x^{i-1}_{i-1}\}) \leqslant (C_{i-1}, D_{i-1})$;
- $|C_i, D_i| = i$;
- $(C_i, D_i) \in \overrightarrow{S}'_{n+1}$.

Furthermore, we can label $C_i \cap D_i = \{x^i_0, \ldots, x^i_{i-1}\}$ such that there exists a family of vertex disjoint paths $\{P^i_j : j \leqslant i-1\}$ such that P^i_j is a path from x^{i-1}_j to x^i_j paths P^i_j for each j. Finally, there is some $x^i_i \in C_i \setminus D_i$ such that if $v_i \in F$ has an in-neighbour $v_{k(i)}$ then there is an edge $(x^i_{k(i)}, x^i_i) \in E(D)$.

Suppose we have constructed x^{i-1}_{i-1} and (C_{i-1}, D_{i-1}). Since $(C_{i-1}, D_{i-1}) \in \overrightarrow{S}'_{n+1}$ there is a partial $\overrightarrow{S}$-path (P, α) of width $< n$ with $V(P) = \{t_1, \ldots, t_m\}$ such that $\alpha(t_1, t_2) = (C_{i-1}, D_{i-1})$. We can then form a partial $\overrightarrow{S}$-path (P', α') by letting P' be a path with $V(P') = \{t_0, t_1, \ldots, t_m\}$, $\alpha'(t_0, t_1) = (C_{i-1}, D_{i-1} \cup \{x^{i-1}_{i-1}\})$ and $\alpha' = \alpha$ on P.

Since $i \leqslant n$, it is clear that this is a partial $\overrightarrow{S}$-path of width $< n$, and so $(C_{i-1}, D_{i-1} \cup \{x^{i-1}_{i-1}\}) \in \overrightarrow{S}'_{n+1}$. Therefore, the set of separations satisfying the properties is non-empty, and so there is some $\leqslant$-minimal element (C_i, D_i) which satisfies the three properties.

We claim that $\lambda((C_i, D_i), (C_{i-1}, D_{i-1} \cup \{x^{i-1}_{i-1}\})) = i$. Indeed, suppose for contradiction there exists $(C_i, D_i) < (X, Y) < (C_{i-1}, D_{i-1} \cup \{x^{i-1}_{i-1}\})$ with $|X, Y| = \lambda((C_i, D_i), (C_{i-1}, D_{i-1} \cup \{x^{i-1}_{i-1}\})) < i$. Note that, since $|C_i, D_i| = |C_{i-1}, D_{i-1} \cup \{x^{i-1}_{i-1}\}| = i$, the inequalities are strict.

By assumption, there is a partial $\overrightarrow{S}$-path (P, α) of width $< n$ with initial leaf separation (C_i, D_i). By construction (X, Y) is up-linked to (C_i, D_i), and so by Lemma 5.2.5 the up-shift of (P, α) onto (X, Y) with respect to (C_i, D_i) is a partial $\overrightarrow{S}$-path of width $< n$ with initial leaf separation (X, Y). Hence $(X, Y) \in \overrightarrow{S}'_{n+1}$.

However, $|X, Y| := \ell < i$ and $(X, Y) < (C_{i-1}, D_{i-1}) \leqslant (C_\ell, D_\ell)$, contradicting the minimality of (C_ℓ, D_ℓ). Therefore $\lambda((C_i, D_i), (C_{i-1}, D_{i-1} \cup \{x^{i-1}_{i-1}\})) = i$, and so by Menger's theorem there exists a family of vertex disjoint $C_i \cap D_i$ to $C_{i-1} \cap (D_{i-1} \cup \{x^{i-1}_{i-1}\})$ paths. Let us label the vertices of $C_i \cap D_i = \{x^i_0, \ldots, x^i_{i-1}\}$ such that these paths are from x^{i-1}_j to x^i_j.

We claim that every x^i_j with $j \leqslant i - 1$ has an out-neighbour in $C_i \setminus D_i$. Indeed, suppose x^i_j does not, then $(C_i \setminus \{x^i_j\}, D_i) \in \overrightarrow{S}$. There exists a partial $\overrightarrow{S}$-path (P, α) with $V(P) = \{t_1, \ldots, t_m\}$ of width $< n$ with initial leaf separation (C_i, D_i) and so if we consider (P, α') with $\alpha'(t_1, t_2) = (C_i \setminus \{x^i_j\}, D_i)$ and $\alpha = \alpha'$ on $P[\{t_2, \ldots, t_m\}]$, we see that (P, α') is also a partial $\overrightarrow{S}$-path (P, α) of width $< n$, with initial leaf separation $(C_i \setminus \{x^i_j\}, D_i)$.

Therefore, $(C_i, D_i) > (C_i \setminus \{x^i_j\}, D_i) \in \overrightarrow{S}'_{n+1}$, contradicting the minimality of (C_{i-1}, D_{i-1}). Hence, if $v_{k(i)}$ is the in-neighbour of v_i in F, then we can pick $x^i_i \in C_i \setminus D_i$ such that there is an edge $(x^i_{k(i)}, x^i_i) \in E(D)$.

For each $0 \leqslant j \leqslant n - 1$ let $P_j = \bigcup_{i=j}^n P^i_j$. Then P_j is a path from x^j_j to x^n_j containing x^i_j for each $j \leqslant i \leqslant n - 1$. Consider the subgraph of D given by

$$D' = \bigcup_{j=0}^{n} P_j \cup \{(x^i_{k(i)}, x^i_i) \,:\, 1 \leqslant i \leqslant n\}$$

Note that D' is an arborescence, and so each vertex has at most 1 in-neighbour. Hence every edge is contractible, and by contracting each P_i we obtain F as a butterfly minor. $\square$

One of the strengths of Theorem 5.1.3 is that there exist forests with unbounded path-width, and so the theorem gives a family graphs that must appear as a minor of a graph with sufficiently large path-width, and conversely cannot appear as a minor of a graph with small path-width.

Unfortunately, there do not exist arborescences of unbounded directed path-width. Indeed, since the vertices of an arborescence can be linearly ordered such that no edge goes 'backwards', every arborescence has directed path-width 0, and moreso this observation is even true for all directed acyclic graphs. It would be interesting to know if these methods could be used to prove a theorem similar to Theorem 5.1.5 for a class of graphs whose directed path-width is unbounded.

5.5 Linked directed path-decompositions

A digraph D is *simple* if it is loopless and there is at most one edge (u, v) for every $u, v \in V(D)$. A simple digraph is *semi-complete* if for every pair $u, v \in V(D)$ either $(u, v) \in E(D)$ or $(v, u) \in E(D)$. A semi-complete digraph is a *tournament* if exactly one of (u, v) and (v, u) is an edge. Kim and Seymour [85] considered the following property of a directed path-decomposition $(P, \mathcal{V})$:

$$\text{If } |V_k| \geqslant t \text{ for every } i \leqslant k \leqslant j, \text{ then there exists a collection of } t \text{ vertex-disjoint directed paths from } V_j \text{ to } V_i. \tag{5.5.1}$$

Theorem 5.5.1 (Seymour and Kim). *Let $D = (V, E)$ be a semi-complete digraph. D has a directed path-decomposition $(P, \mathcal{V})$ of width $dpw(D)$ satisfying (5.5.1).*

Seymour and Kim called directed path-decompositions satisfying (5.5.1), as well as two other technical conditions, 'linked'. However, when thinking about directed path-decompositions in terms of separations, perhaps a more natural concept to call 'linked' is the following (See [59]). We say an $\overrightarrow{S}$-path (P, α) with $P = \{t_1, \ldots, t_n\}$ is *linked* if for every $1 \leqslant i < j \leqslant n-1$

$$\min\{|\alpha(t_k, t_{k+1})| : i \leqslant k \leqslant j\} = \lambda\big((A_i, B_i), (A_j, B_j)\big). \tag{5.5.2}$$

Remark. *Given a linked directed path-decompositon $(P, \mathcal{V})$ we can form an $\overrightarrow{S}$-path as in the discussion preceding Lemma 5.2.2. It is easy to check that this $\overrightarrow{S}$-path is linked.*

Conversely, given a linked $\overrightarrow{S}$-path (P, α) we can form a directed path-decomposition $(P, \mathcal{V})$ as in Lemma 5.2.2. However, it is not true that this directed path-decomposition will be linked in the sense of (5.5.1), however it is easy to adapt it to form a linked directed path-decomposition by subdividing each edge (t_i, t_{i+1}) of P by a new vertex s_i and adding a bag at s_i which is the adhesion set of the edge $(t_i, t_{i+1}$. Note that this process does not increase the width or adhesion of the directed path-decomposition.

Given an $\overrightarrow{S}$-path (P, α) and $r \in \mathbb{N}$ let us write P_r for the linear sub-forest of P induced by the edges $e \in E(P)$ such that $|\alpha(e)| \geqslant r$. Let us denote by $e(P_r)$ and $c(P_r)$ the number of edges and components of P_r respectively.

Again we will prove a slightly more general theorem about $\overrightarrow{S}$-paths where we fix independently the size of the adhesion sets. Theorem 5.1.6 will follow from the following theorem, and Remark , if we let $\omega = k$. We note that a stronger result, which would imply Theorem 5.1.6, is claimed in a preprint of Kintali [86]. However we were unable to verify the proof, and include a counterexample to his claim in Section 5.6.

Theorem 5.5.2. *Let $D = (V, E)$ be a digraph and let $k \leqslant \omega \in \mathbb{N}$ be such that there exists an $\overrightarrow{S}_k$-path of width $< \omega - 1$. There exists a linked $\overrightarrow{S}_k$-path of width $< \omega - 1$.*

Remark. *As in Lemma 5.2.2 and the discussion preceding it, it is easy to see that a digraph D has is a linked $\overrightarrow{S}_k$-path of width $< \omega - 1$ if and only if it has a linked path-decomposition of adhesion $< k$ and width $< \omega - 1$.*

Proof. Let us define a partial order on the set of $\overrightarrow{S}_k$-paths of D of width $< \omega - 1$ by letting $(P, \alpha) \leqslant (Q, \beta)$ if there is some r such that,

- for all $r' > r$, $e(P_{r'}) = e(Q_{r'})$ and $c(P_{r'}) = c(Q_{r'})$;
- either $e(P_r) < e(Q_r)$, or $e(P_r) = e(Q_r)$ and $c(P_r) > c(Q_r)$;

Note that, since $c(P_r)$ is at most $e(P_r) + 1$, it is relatively simple to show that there are no ifninite decreasing chains in this partial order. Hence, since there exists at least one directed path-decomposition of D of width $< \omega - 1$, there is some minimal element in this partial order, (P, α) with $V(P) = \{t_1, \ldots, t_n\}$ and $\alpha(t_i, t_{i+1}) = (A_i, B_i)$. We claim that (P, α) is linked.

Suppose for contradiction that (P, α) is not linked. That is, there exists $1 \leqslant i < j \leqslant n-1$ such that

$$\lambda\big((A_i, B_i), (A_j, B_j)\big) < \min\{|A_k, B_k| : i \leqslant k \leqslant j\}.$$

We will construct another directed path-decomposition $(\hat{P}, \hat{\alpha})$ of width $< \omega - 1$ such that $(\hat{P}, \hat{\alpha}) < (P, \alpha)$.

Let us choose a separation $(A_i, B_i) \leqslant (X, Y) \leqslant (A_j, B_j)$ such that $|X, Y| = \lambda\big((A_i, B_i), (A_j, B_j)\big)$. Note that (X, Y) is up-linked to (A_i, B_i) and down-linked to (A_j, B_j)

We form two new $\overrightarrow{S_k}$-paths (P', α') and (P'', α'') by taking the up-shift of (P, α) onto (X, Y) with respect to (A_i, B_i) and the down-shift of (P, α) onto (X, Y) with respect to (A_j, B_j). Let us denote by $(A'_i, B'_i), \ldots, (A'_{n-1}, B'_{n-1})$ and $(A''_1, B''_1), \ldots, (A''_j, B''_j)$ for the images of α' and α''. We note that the initial leaf separation of (P', α') and the terminal leaf separation of (P'', α'') are both (X, Y).

We form a new $\overrightarrow{S}_k$-path $(\hat{P}, \hat{\alpha})$ by letting $\hat{P}$ be the path formed by identifying the initial leaf of (P', α') with the terminal leaf of (P'', α'') and taking $\hat{\alpha}$ to be α' on $E(P')$ and α'' on $E(P'')$.

By Lemma 5.2.5, since (P, α) was of width $< \omega - 1$, so is $(\hat{P}, \hat{\alpha})$. We claim that $(\hat{P}, \hat{\alpha}) < (P, \alpha)$. Given a vertex $t_k \in V(P)$ we will write t'_k and t''_k for the copy of t_k in P' or P'' respectively, and carry these labels over onto $\hat{P}$. Note that, not every vertex will appear in both P' and P''.

Claim 5.5.3. *For every $k \in [n]$ $\max\{|A'_k, B'_k|, |A''_k, B''_k|\} \leqslant |A_k, B_k|$. Furthermore, if $|A_k, B_k| = |A'_k, B'_k|$ then $|A''_k, B''_k| \leqslant |X, Y|$. Similarly if $|A_k, B_k| = |A''_k, B''_k|$ then $|A'_k, B'_k| \leqslant |X, Y|$. (Here we are assuming for convenience that if (A'_k, B'_k) or (A''_k, B''_k) do not exist then their order is 0.)*

Proof of claim. For the first claim we note that the proof of Lemma 5.2.4 in fact shows the stronger statement that if (X, Y) is up/down-linked to (A_i, B_i) then the order of each separation in the up/down-shift of (P, α) onto (X, Y) with respect to (A_i, B_i) does not increase.

For the second claim, it is sufficient to prove it for $i \leqslant k \leqslant j$, since otherwise one of (A'_k, B'_k) or (A''_k, B''_k) has order 0.

Since (P', α') was the up-shift of (P, α) onto (X, Y) with respect to (A_i, B_i), $(A'_k, B'_k) = (A_k, B_k) \vee (X, Y) = (A_k \cup X, B_k \cap Y)$. Similarly $(A''_k, B''_k) = (A_k, B_k) \wedge (X, Y) = (A_k \cap X, B_k \cup Y)$.

Hence, by (5.2.1),

$$|A'_k, B'_k| + |A''_k, B''_k| = |A_k, B_k| + |X, Y|.$$

□

Let us write e_k for the edge $(t_k, t_{k+1}) \in E(P)$, and e'_k, e''_k for the two copies of e_k in $E(\hat{P})$ (when they exist).

Claim 5.5.4. *For every $r > |X, Y|$ and every $k \in [n-1]$ such that $|A_k, B_k| = r$ exactly one of (A'_k, B'_k), (A''_k, B''_k) has order $|A_k, B_k|$, and the other has order $\leqslant |X, Y|$. Furthermore, for each component C of P^r, and $e_p, e_q \in C$, then $|A'_p, B'_p| = |A_p, B_p|$ if and only if $|A'_q, B'_q| = |A_q, B_q|$, and similarly $|A''_p, B''_p| = |A_p, B_p|$ if and only if $|A''_q, B''_q| = |A_q, B_q|$.*

Proof of Claim. We will prove the claim by reverse induction on r, starting with r being the order of the largest separation in (P, α). Note that, since $|A_i, B_i|, |A_j, B_j| > |X, Y|$, it follows that $r > |X, Y|$.

By the first part of Claim 5.5.3 we have that the order of the largest separation in $(\hat{P}, \hat{\alpha})$ is at most r. Hence, since (P, α) was minimal, $e(P_r) \leqslant e(\hat{P}_r)$. However, by Claim 5.5.3, for each (A_k, B_k) with $|A_k, B_k| = r$, at most one of the two separations (A'_k, B'_k) and (A''_k, B''_k) have order r, and if it does, then the other has order $\leqslant |X, Y| < r$. Therefore it follows that $e(\hat{P}_r) \leqslant e(P_r)$, and so $e(P_r) = e(\hat{P}_r)$, and the first part of the claim follows.

By minimality of (P, α) again, it follows that $c(P_r) \geqslant c(\hat{P}_r)$. However, since the edge $e'_i = e''_j$ in $\hat{P}$ is mapped to the separation (X, Y) by $\hat{\alpha}$, it follows from the first half of the claim that $c(\hat{P}_r) \geqslant c(P_r)$, and so $c(P_r) = c(\hat{P}_r)$, and the second part of the claim follows.

Suppose then that the claim holds for all $r' > r$. It follows that $e(P_{r'}) = e(\hat{P}_{r'})$ and $c(P_{r'}) = c(\hat{P}_{r'})$ for all $r' > r$ and so, since (P, α) was minimal, $e(P_r) \leqslant e(\hat{P}_r)$.

However, since by Lemma 5.2.4 the order of each separation does not increase when we shift an $\overrightarrow{S}_k$-tree, the only edges in $\hat{P}_r$ come from copies of edges in $P_{r'}$ with $r' \geqslant r$. If $r' > r$ then, by the induction hypothesis, these copies have order r', or $\leqslant |X, Y|$. If $r' = r$ then, by Claim

5.5.3 at most one of the two copies of the edge has order r, and if it does the other has order $\leqslant |X,Y|$. It follows that $e(\hat{P}_r) \leqslant e(P_r)$, and so $e(P_r) = e(\hat{P}_r)$, and the first part of the claim follows.

By minimality of (P,α) again, it follows that $c(P_r) \geqslant c(\hat{P}_r)$. However, since the edge $e'_i = e''_j$ in $\hat{P}_r$ is mapped to the separation (X,Y) by $\hat{\alpha}$, it follows from the first half of the claim that $c(\hat{P}_r) \geqslant c(P_r)$, and so $c(P_r) = c(\hat{P}_r)$, and the second part of the claim follows. □

Recall that $|A_k, B_k| > |X,Y|$ for all $i \leqslant k \leqslant j$ by assumption. Hence e_i and e_j lie in the same component of $P_{|X,Y|+1}$. However, $(A'_i, B'_i) = (A''_j, B''_j) = (X,Y)$, contradicting the second part of Claim 5.5.4. □

5.6 Counterexample to the existence of lean directed path-decompositions

Kintali [86] defines a directed path-decomposition to be *lean* if it satisfies the following condition:

$$\begin{aligned}&\text{Given } k > 0,\ t_1 \leqslant t_2 \in [n] \text{ and subsets } Z_1 \subseteq V_{t_1} \text{ and } Z_2 \subseteq V_{t_2} \text{ with } |Z_1| = |Z_2| = k\\ &\text{either } G \text{ contains } k \text{ vertex-disjoint directed paths from } Z_2 \text{ to } Z_1 \text{ or there exists}\\ &i \in [t_1, t_2 - 1] \text{ such that } |V_i \cap V_{i+1}| < k\end{aligned} \tag{5.6.1}$$

Note, this is a strengthening of (5.5.1). In particular, (5.6.1) has content in the case $t_1 = t_2$. When Thomas proved his result on the existence of linked tree-decompositions of minimal width [124] he in fact established the existence of tree-decompositions satisfying a stronger condition in the vein of (5.6.1) (which are sometimes called *lean* tree-decompositions in the literature [43]). Kintali claims the following analogous result.

Theorem 5.6.1 ([86] Theorem 7)**.** *Every digraph D has a directed path-decomposition of width $dpw(D)$ satisfying (5.6.1).*

However, we note that this theorem cannot hold. Indeed, consider a perfect binary tree of depth n, with all edges in both directions. Let us write T_n for the undirected tree and $\overrightarrow{T}_n$ for the digraph. It it easy to see that the directed path-width of $\overrightarrow{T}_n$ is equal to the path width of T_n, which is $\frac{n-1}{2}$. Hence, in any directed path-decomposition of $\overrightarrow{T}_n$ there is some bag of size at least $\frac{n+1}{2}$. Suppose that a lean directed path-decomposition exists, let us denote by V_i a bag such that $|V_i| \geqslant \frac{n+1}{2}$.

If we consider V_i as a subset of T_n, then it follows from (5.6.1) that for every $k > 0$ and every $Z_1, Z_2 \subseteq V_i$ with $|Z_1| = |Z_2| = k$, T_n contains k vertex-disjoint paths between Z_1 and Z_2. This property is known in the literature as being *well-linked*, and the size of the largest well-linked set in a graph is linearly related to the tree-width (see for example [77]). Specifically, since T_n contains a well-linked set of size $\geqslant \frac{n+1}{2}$ it follows that $\mathrm{tw}(T_n) \geqslant \frac{n+1}{6}$. However, the tree-width of any tree is one, contradicting the existence of a lean directed path-decomposition of $\overrightarrow{T}_n$ for $n \geqslant 6$.

Chapter 6

A short derivation of the structure theorem for graphs with excluded topological minors

6.1 Introduction

A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges. In a series of 23 papers, published between 1983 and 2012, Robertson and Seymour developed a deep theory of graph minors which culminated in the proof of *Wagner's Conjecture* [115], which asserts that in any infinite set of finite graphs there is one which is a minor of another. One of the landmark results proved along the way, and indeed a fundamental step in resolving Wagner's Conjecture, is a structure theorem for graphs excluding a fixed graph as a minor [111]. It is easy to see that G cannot contain H as a minor if there is a surface into which G can be embedded but H cannot. Loosely speaking, the structure theorem of Robertson and Seymour asserts an approximate converse to this, thereby revealing the deep connection between topological graph theory and the theory of graph minors:

Theorem 6.1.1 ([111] (informal))**.** *For any $n \in \mathbb{N}$, every graph excluding the complete graph K_n as a minor has a tree-decomposition in which every torso is almost embeddable into a surface into which K_n is not embeddable.*

A graph H is a *topological minor* of a graph G if G contains a subdivision of H as a subgraph. It is easy to see that G then also contains H as a minor. The converse is not true, as there exist cubic graphs with arbitrarily large complete minors. For topological minors, we thus have an additional degree-based obstruction, which is fundamentally different from the topological obstruction of surface-embeddings for graph minors. Grohe and Marx [70] proved a result in a similar spirit to Theorem 6.1.1 for graphs excluding a fixed graph as a topological minor:

Theorem 6.1.2 ([70] (informal))**.** *For any $n \in \mathbb{N}$, every graph excluding K_n as a topological minor has a tree-decomposition in which every torso either*

(cl.1) *has a bounded number of vertices of high degree, or*

(cl.2) *is almost embeddable into a surface of bounded genus.*

More recently, Dvořák [56] refined the embeddability condition of this theorem to reflect more closely the topology of embeddings of an arbitrary graph H which is to be excluded as a topological minor.

The proof given in [70], which uses Theorem 6.1.2 as a block-box, is algorithmic and explicitly provides a construction of the desired tree-decomposition, however as a result the proof is quite technical in parts. In this paper, we give a short proof of Theorem 6.1.2 which also provides a good heuristic for the structure of graphs without a large complete topological minor, as well as improving the implicit bounds given in [70] on many of the parameters in their theorem. Our proof is non-constructive, but we note that it can easily be adapted to give an algorithm to find either a subdivision of K_r or an appropriate tree-decomposition. However, the run time of this algorithm will be much slower than that of the algorithm given in [70].

One of the fundamental structures we consider are k-blocks. A *k-block* in a graph G is a set B of at least k vertices which is inclusion-maximal with the property that for every separation (U, W) of order $< k$, we either have $B \subseteq U$ or $B \subseteq W$. The notion of a k-block, which was first studied by Mader [99, 98], has previously been considered in the study of graph decompositions [33, 35, 37].

It is clear that a subdivision of a clique on $k + 1$ vertices yields a k-block. The converse is not true for any $k \geqslant 4$, as there exist planar graphs with arbitrarily large blocks. The second author [127] proved a structure theorem for graphs without a k-block:

Theorem 6.1.3 ([127])**.** *Let G be a graph and $k \geqslant 2$. If G has no $(k + 1)$-block then G has a tree-decomposition in which every torso has at most k vertices of degree at least $2k(k - 1)$.*

Now, since a subdivision of a complete graph gives rise to both a complete minor and a block, there are two obvious obstructions to the existence of a large topological minor, the absence of a large complete minor or the absence of a large block. The upshot of Theorem 6.1.2 is that in a local sense these are the only obstructions, any graph without a large topological minor has a tree-decomposition into parts whose torsos either don't contain a large minor, or don't contain a large block. Furthermore, by Theorem 6.1.1 and Theorem 6.1.3, the converse should also be true: if we can decompose the graph into parts whose torsos either don't contain a large minor or don't contain a large block, then we can refine this tree-decomposition into one satisfying the requirements of Theorem 6.1.2.

The idea of our proof is as follows. Both large minors and large blocks point towards a 'big side' of every separation of low order. A subdivision of a clique simultaneously gives rise to both a complete minor and a block and, what's more, the two are hard to separate in that they choose the same 'big side' for every low-order separation. A qualitative converse to this is already implicit in previous work on graph minors and linkage problems: if a graph contains a large complete minor and a large block which cannot be separated from that minor, then the graph contains a subdivision of a complete graph.

Therefore, if we assume our graph does not contain a subdivision of K_r, then we can separate any large minor from every large block. It then follows from the *tangle tree theorem* of Robertson and Seymour [110] – or rather its extension to *profiles* [80, 48, 35] – that there exists a tree-decomposition which separates the blocks from the minors. Hence each part is either free of large minors or of large blocks.

However, in order to apply Theorems 6.1.1 and 6.1.3, we need to have control over the *torsos*, and not every tree-decomposition will provide that: it might be, for example, that separating some set of blocks created a large minor in one of the torsos. We therefore contract some parts of our tree-decomposition and use the minimality of the remaining separations to prove that this does not happen.

A second nice feature of our proof is that we avoid the difficulty of constructing such a tree-decomposition by choosing initially a tree-decomposition with certain connectivity properties, the proof of whose existence already exists in the literature, and then simply *deducing* that this tree-decomposition has the required properties.

We are going to prove the following:

Theorem 6.1.4. *Let r be a positive integer and let G be a graph containing no subdivision of K_r. Then G has a tree-decomposition of adhesion $<r^2$ such that every torso either*

(cl.1) *has fewer than r^2 vertices of degree at least $2r^4$, or*

(cl.2) *has no K_{2r^2}-minor.*

Combining Theorems 6.1.1 and 6.1.4 then yields Theorem 6.1.2.

Let us briefly compare the bounds we get to the result of Grohe and Marx [70, Theorem 4.1]. It is implicit in their results that if G contains no subdivision of K_r, then G has a tree-decomposition of adhesion $O(r^6)$ such that every torso either has $O(r^6)$ vertices of degree $\Omega(r^7)$, has no $K_{\Omega(r^6)}$ minor or has size at most $O(r^6)$. In this way, Theorem 6.1.4 gives an improvement on the bounds for each of the parameters. Recently Liu and Thomas [94] also proved an extension of the work of Dvořák [56], with the aim to more closely control the bound on the degrees of the vertices in (i). Their results, however, only give this structure 'relative' to some tangle.

6.2 Notation and background material

All graphs considered here are finite and undirected and contain neither loops nor parallel edges. Our notation and terminology mostly follow that of [43].

Given a tree T and $s, t \in V(T)$, we write sTt for the unique s-t-path in T. A *separation* of a graph $G = (V, E)$ is a pair (A, B) with $V = A \cup B$ such that there are no edges between $A \setminus B$ and $B \setminus A$. The *order* of (A, B) is the number of vertices in $A \cap B$. We call the separation (A, B) *tight* if for all $x, y \in A \cap B$, both $G[A]$ and $G[B]$ contain an x-y-path with no internal vertices in $A \cap B$.

The set of all separations of G of order $<k$ will be denoted by $S_k(G)$. An *orientation* of $S_k(G)$ is a subset of $S_k(G)$ containing precisely one element from each pair $\{(A, B), (B, A)\} \subseteq S_k(G)$. The orientation is *consistent* if it does not contain two separations $(A, B), (C, D)$ with $B \subseteq C$ and $D \subseteq A$. A separation *distinguishes* two orientations O_1, O_2 of $S_k(G)$ if precisely one of O_1, O_2 contains it. It does so *efficiently* if it has minimum order among all separations distinguishing them.

Recall that, given an integer k, a set B of at least k vertices of G is a *k-block* if it is inclusion-maximal with the property that for every separation (U, W) of order $<k$, either $B \subseteq U$ or $B \subseteq W$. Observe that B induces a consistent orientation $O_B := \{(U, W) \colon B \subseteq W\}$ of $S_k(G)$.

Given an integer m, a *model of K_m* is a family $\mathcal{X}$ of m pairwise disjoint sets of vertices of G such that $G[X]$ is connected for every $X \in \mathcal{X}$ and G has an edge between X and Y for any two $X, Y \in \mathcal{X}$. The elements of $\mathcal{X}$ are called *branch sets*. Note that, if (U, W) is a separation of order $<m$, then exactly one of $U \setminus W$ and $W \setminus U$ contains some branch set. In this way, $\mathcal{X}$ induces a consistent orientation $O_{\mathcal{X}}$ of $S_k(G)$, where $(U, W) \in O_{\mathcal{X}}$ if and only if some branch set of $\mathcal{X}$ is contained in W.

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

- for every $v \in V(G)$, the set of $t \in V(T)$ with $v \in V_t$ induces a non-empty subtree of T;
- for every edge $vw \in E(G)$ there is a $t \in V(T)$ with $v, w \in V_t$.

If $(T, \mathcal{V})$ is a tree-decomposition of G, then every $st \in E(T)$ induces a separation

$$(U_s, W_t) := \Big(\bigcup_{t \notin uTs} V_u, \bigcup_{s \notin vTt} V_v\Big).$$

Note that $U_s \cap W_t = V_s \cap V_t$. In this way, every edge $e \in E(T)$ has an *order* given by the order of the separation it induces, which we will write as $|e|$. Similarly, an edge of T *(efficiently) distinguishes* two orientations if the separation it induces does. We say that $(T, \mathcal{V})$ *(efficiently) distinguishes* two orientations O and P if some edge of T does. We call $(T, \mathcal{V})$ *tight* if every separation induced by an edge of T is tight.

The *adhesion* of $(T, \mathcal{V})$ is the maximum order of an edge. If the adhesion of $(T, \mathcal{V})$ is less than k and O is an orientation of $S_k(G)$, then O induces an orientation of the edges of T by orienting an edge st towards t if $(U_s, W_t) \in O$. If O is consistent, then all edges will be directed towards some node $t \in V(T)$, which we denote by t_O and call the *home node* of O. When O is induced by a block B or model $\mathcal{X}$, we abbreviate $t_B := t_{O_B}$ and $t_{\mathcal{X}} := t_{O_{\mathcal{X}}}$, respectively. Observe that and edge $e \in E(T)$ distinguishes two orientations O and P if and only if $e \in E(t_O T t_P)$.

Given $t \in V(T)$, the *torso at* t is the graph obtained from $G[V_t]$ by adding, for every neighbor s of t, an edge between any two non-adjacent vertices in $V_s \cap V_t$. More generally, given a subtree $S \subseteq T$, the *torso at* S is the graph obtained from $G\left[\bigcup_{s \in S} V_s\right]$ by adding, for every edge $st \in E(T)$ with $S \cap \{s, t\} = \{s\}$, an edge between any two non-adjacent vertices in $V_s \cap V_t$.

We also define contractions on tree-decompositions: Given $(T, \mathcal{V})$ and an edge $st \in E(T)$, to contract the edge st we form a tree-decomposition $(T', \mathcal{V}')$ where

- T' is obtained by contracting st in T to a new vertex x;
- Let $V'_x := V_s \cup V_t$ and $V'_u := V_u$ for all $u \in V(T) \setminus \{s, t\}$.

It is simple to check that $(T', \mathcal{V}')$ is a tree-decomposition. We note that the separations induced by an edge in $E(T) \setminus \{st\}$ remain the same, as do the torsos of parts V_u for $u \neq s, t$.

We say a tree-decomposition $(T, \mathcal{V})$ is *k-lean* if it has adhesion $< k$ and the following holds for all $p \in [k]$ and $s, t \in T$: If sTt contains no edge of order $< p$, then every separation (A, B) with $|A \cap V_s| \geqslant p$ and $|B \cap V_t| \geqslant p$ has order at least p.

Let $n := |G|$. The *fatness* of $(T, \mathcal{V})$ is the sequence $(a_0, \ldots, a_n)$, where a_i denotes the number of parts of order $n - i$. A tree-decomposition of lexicographically minimum fatness among all tree-decompositions of adhesion smaller than k is called *k-atomic*. These tree-decompositions play a pivotal role in our proof, but we actually only require two properties that follow from this definition. It was observed by Carmesin, Diestel, Hamann and Hundertmark [34] that the short proof of Thomas' Theorem [124] given by Bellenbaum and Diestel in [17] also shows that k-atomic tree-decompositions are k-lean (see also [66]).

Lemma 6.2.1 ([17]). *Every k-atomic tree-decomposition is k-lean.*

It is also not hard to see that k-atomic tree-decompositions are tight. In [127], the second author used k-atomic tree-decompositions to prove a structure theorem for graphs without a k-block. In fact, the proof given there yields the following:

Lemma 6.2.2 ([127]). *Let G be a graph and k a positive integer. Let $(T, \mathcal{V})$ be a k-atomic tree-decomposition of G and $t \in V(T)$ such that V_t contains no k-block of G. Then the torso at t contains fewer than k vertices of degree at least $2k^2$.*

Let G be a graph and $Z \subseteq V(G)$. We denote by G^Z the graph obtained from G by making the vertices of Z pairwise adjacent. A *Z-based model* is a model $\mathcal{X}$ of $K_{|Z|}$ such that $X \cap Z$ consists of a single vertex for every $X \in \mathcal{X}$.

The following lemma of Robertson and Seymour [114] is crucial to our proof.

Lemma 6.2.3 ([114]). *Let G be a graph, $Z \subseteq V(G)$ and $p := |Z|$. Let $q \geqslant 2p - 1$ and let $\mathcal{X}$ be a model of K_q in G^Z. If $\mathcal{X}$ and Z induce the same orientation of $S_p(G^Z)$, then G has a Z-based model.*

6.3 The proof

Let us fix throughout this section a graph G with no subdivision of K_r, let $k := r(r-1)$, $m := 2k$, and let $(T, \mathcal{V})$ be a k-atomic tree-decomposition of G.

First, we will show that $(T, \mathcal{V})$ efficiently distinguishes every k-block from every model of K_m in G. This allows us to split T into two types of sub-trees, those containing a k-block and those containing a model of K_m. Lemma 6.2.2 allows us to bound the number of high degree degree vertices in the torsos in the latter components. We will then show that if we choose these sub-trees in a sensible way then we can also bound the order of a complete minor contained in the torsos of the former. Hence, by contracting each of these sub-trees in $(T, \mathcal{V})$ we will have our desired tree-decomposition.

To show that $(T, \mathcal{V})$ distinguishes every k-block from every model of K_m in G, we must first show that they are distinguishable, that is, no k-block and K_m induce the same orientation. The following lemma, as well as its proof, is similar to Lemma 6.11 in [70].

Lemma 6.3.1. *Let B be a k-block and $\mathcal{X}$ a model of K_m in G. If B and $\mathcal{X}$ induce the same orientation of S_k, then G contains a subdivision of K_r with arbitrarily prescribed branch vertices in B.*

Proof. Suppose B and $\mathcal{X}$ induce the same orientation and let B_0 be an arbitrary subset of B of size r. Let H be the graph obtained from G by replacing every $b \in B_0$ by an independent set J_b of order $(r-1)$, where every vertex of J_b is adjacent to every neighbor of b in G and to every vertex of J_c if b, c are adjacent. Let $J := \bigcup_b J_b$ and note that $|J| = k$. We regard G as a subgraph of H by identifying each $b \in B$ with one arbitrary vertex in J_b. In this way we can regard $\mathcal{X}$ as a model of K_m in H.

Assume for a contradiction that there was a separation (U, W) of H such that $|U \cap W| < |J|$, $J \subseteq U$ and $X \subseteq W \setminus U$ for some $X \in \mathcal{X}$. We may assume without loss of generality that for every $b \in B_0$, either $J_b \subseteq U \cap W$ or $J_b \cap (U \cap W) = \emptyset$. Indeed, if there is a $z \in J_b \setminus (U \cap W)$, then $z \in U \setminus W$, and we can delete any $z' \in J_b \cap W$ from W and maintain a separation (because $N(z) = N(z')$) with the desired properties. In particular, for every $b \in B_0$ we find $b \in W$ if and only if $J_b \subseteq W$. Since $|U \cap W| < |J|$, it follows that there is at least one $b_0 \in B_0$ with $J_{b_0} \subseteq (U \setminus W)$. Let $(U', W') := (U \cap V(G), W \cap V(G))$ be the induced separation of G. Then $X \subseteq W' \setminus U'$ and $b_0 \in U' \setminus W'$. Since $|U' \cap W'| \leqslant |U \cap W| < k$ and B is a k-block, we have $B \subseteq U'$. But then (U', W') distinguishes B and $\mathcal{X}$, which is a contradiction to our initial assumption.

We can now apply Lemma 6.2.3 to H and find a J-based model $\mathcal{Y} = (Y_j)_{j \in J}$ in H. For each $b \in B_0$, label the vertices of J_b as $(v^b_c)_{c \in B_0 \setminus \{b\}}$. For $b \neq c$, H has a path $P'_{b,c} \subseteq Y_{v^b_c} \cup Y_{v^c_b}$ and the paths obtained like this are pairwise disjoint, because the Y_j are, and $P'_{b,c} \cap J = \{v^b_c, v^c_b\}$. For each such path $P'_{b,c}$, obtain $P_{b,c} \subseteq G$ by replacing v^b_c by b and v^c_b by c. The collection of these paths $(P_{b,c})_{b,c \in B_0}$ gives a subdivision of K_r with branch vertices in B_0. □

Now we can show that $(T, \mathcal{V})$ efficiently distinguishes every k-block from every model of K_m in G.

Lemma 6.3.2. *$(T, \mathcal{V})$ efficiently distinguishes all orientations of $S_k(G)$ induced by k-blocks or models of K_m.*

Proof. Let us call a consistent orientation O of $S_k(G)$ *anchored* if for every $(U, W) \in O$, there are at least k vertices in $W \cap V_{t_O}$.

Note that every orientation $O = O_B$ induced by a k-block B is trivially anchored, since $B \subseteq V_{t_B}$. But the same is true for the orientation $O = O_{\mathcal{X}}$ induced by a model $\mathcal{X}$ of K_m. Indeed, let $(U, W) \in O_{\mathcal{X}}$. Then every set in $\mathcal{X}$ meets $V_{t_{\mathcal{X}}}$. At least k branch sets of $\mathcal{X}$ are disjoint from $U \cap W$, say $X_1, \ldots, X_k$, and they all lie in $W \setminus U$. For $1 \leqslant i \leqslant k$, let $x_i \in X_i \cap V_{t_{\mathcal{X}}}$ and note that $R := \{x_1, \ldots, x_k\} \subseteq W \cap V_{t_{\mathcal{X}}}$.

We now show that $(T, \mathcal{V})$ efficiently distinguishes all anchored orientations of $S_k(G)$. Let O_1, O_2 be anchored orientations of $S_k(G)$ and let their home nodes be t_1 and t_2 respectively. If $t_1 \neq t_2$, let p be the minimum order of an edge along $t_1 T t_2$, and put $p := k$ otherwise. Choose some $(U, W) \in O_2 \setminus O_1$ of minimum order. Since O_1 and O_2 are anchored, we have $|U \cap V_{t_1}| \geqslant k$ and $|W \cap V_{t_2}| \geqslant k$. As $(T, \mathcal{V})$ is k-lean, it follows that $|U \cap W| \geqslant p$. Hence $t_1 \neq t_2$ and $(T, \mathcal{V})$ efficiently distinguishes O_1 and O_2. □

Let us call a node $t \in V(T)$ a *block-node* if it is the home node of some k-block and *model-node* if it is the home node of a model of K_m.

Let $F \subseteq E(T)$ be inclusion-minimal such that every k-block is efficiently distinguished from every model of K_m by some separation induced by an edge in F. We now define a red/blue colouring $c : V(T) \to \{r, b\}$ by letting $c(t) = b$ if the component of $T - F$ containing t contains a block-node and letting $c(t) = r$ if it contains a model-node. Let us first show that this is in fact a colouring of $V(T)$.

Lemma 6.3.3. *Every node receives exactly one colour.*

Proof. Suppose first that $t \in V(T)$ is such that the component of $T - F$ containing t contains both a block node and a model node. Then there is a k-block B and a K_m-minor $\mathcal{X}$ such that $t_B T t$ and $t_{\mathcal{X}} T t$ both contain no edges of F. But then B and $\mathcal{X}$ are not separated by the separations induced by F, a contradiction.

Suppose now that $t \in V(T)$ is such that the component S of $T - F$ containing t contains neither a block nor a minor. Let $f_1, \ldots, f_n$ be the edges of T between S and $T \setminus S$, ordered such that $|f_1| \geqslant |f_i|$ for all $i \leqslant n$. By minimality of F, there is a block-node t_B and a model-node $t_{\mathcal{X}}$ such that f_1 is the only edge of F that efficiently distinguishes B and $\mathcal{X}$. Since $t_B, t_{\mathcal{X}} \notin S$, there is a $j \geqslant 2$ such that $f_j \in E(t_B T t_{\mathcal{X}})$, and so f_j distinguishes B and $\mathcal{X}$ as well, and since $|f_1| \geqslant |f_j|$, it does so efficiently, contradicting our choice of B and $\mathcal{X}$ □

Lemma 6.3.4. *Let $st \in E(T)$ and suppose s is blue and t is red. Then $G[W_t]$ has a $(V_s \cap V_t)$-based model.*

Proof. Let $Q := V_s \cap V_t$. Let t_B be a block-node in the same component of $T - F$ as s and let $t_{\mathcal{X}}$ be a model-node in the same component as t. Since the separations induced by F efficiently distinguish B and $\mathcal{X}$, it must be that $st \in F$ and (U_s, W_t) efficiently distinguishes B and $\mathcal{X}$.

Let $\mathcal{Y} := (X \cap W_t)_{X \in \mathcal{X}}$. Since $(U_s, W_t) \in O_{\mathcal{X}}$, $\mathcal{Y}$ is a model of K_m in $G[W_t]^Q$. We wish to apply Lemma 6.2.3 to Q and $\mathcal{Y}$ in the graph $G[W_t]$. Suppose Q and $\mathcal{Y}$ do not induce the same orientation of $S_{|Q|}(G[W_t]^Q)$. That is, there is a separation (U, W) of $G[W_t]^Q$ with $|U \cap W| < |Q|$ and $Q \subseteq U$ such that $Y \cap U = \emptyset$ for some $Y \in \mathcal{Y}$. There is an $X \in \mathcal{X}$ so that $Y = X \cap G[W_t]$. Note that $X \cap U$ is empty as well. Now $(U', W') := (U \cup U_s, W)$ is a separation of G. Note that

$$X \cap U' = X \cap U_s = \emptyset,$$

because X is connected, meets W_t and does not meet Q. Therefore $X \subseteq W' \setminus U'$ and $B \subseteq U_s \subseteq U'$. But $|U' \cap W'| = |U \cap W| < |Q|$, which contradicts the fact that (U_s, W_t) efficiently distinguishes B and $\mathcal{X}$. Therefore, by Lemma 6.2.3, $G[W_t]$ has a Q-based model. □

Using the above we can bound the size of a complete minor in the torso of a blue component. The next lemma plays a similar role to Lemma 6.9 in [70].

Lemma 6.3.5. *Let $S \subseteq T$ be a maximal subtree consisting of blue nodes. Then the torso of S has no K_m-minor.*

Proof. Let $F_S := \{(s,t)\colon st \in E(T), s \in S, t \notin S\}$. For every $(s,t) \in F_S$, the node s is blue and t is red. By Lemma 6.3.4, G_t has a $(V_s \cap V_t)$-based complete minor $\mathcal{Y}^{s,t}$. Contract each of its branch sets onto the single vertex of $V_s \cap V_t$ that it contains. Do this for every $(s,t) \in F_S$. After deleting any vertices outside of $V_S := \bigcup_{s \in S} V_s$, we obtain the torso of S as a minor of the graph G.

Suppose the torso of S contained a K_m-minor. Then G has a K_m-minor $\mathcal{X}$ such that every $X \in \mathcal{X}$ meets V_S. Therefore $\mathcal{X}$ orients every edge $st \in E(T)$ with $(s,t) \in F_S$ towards s. But then $t_{\mathcal{X}} \in S$, contradicting the assumption that S contains no red nodes. $\square$

We can now finish the proof. Let $(T', \mathcal{V}')$ be obtained from $(T, \mathcal{V})$ by contracting every maximal subtree consisting of blue nodes and let the vertices of T' inherit the colouring from $V(T)$. We claim that $(T', \mathcal{V}')$ satisfies the conditions of Theorem 6.1.4.

Indeed, firstly, the adhesion of $(T', \mathcal{V}')$ is at most that of $(T, \mathcal{V})$, and hence is at most k. Secondly, the torso of every red node in $(T', \mathcal{V}')$ is the torso of some red node in $(T, \mathcal{V})$, which by Lemma 6.2.2 has fewer than k vertices of degree at least $2k^2$. Finally, by Lemma 6.3.5 the torso of every blue node in $(T', \mathcal{V}')$ has no K_m minor. Since $k = r(r-1)$ and $m = 2k$, the theorem follows.

As claimed in the introduction, it is not hard to turn this proof into an algorithm to find either a subdivision of K_r or an appropriate tree-decomposition. Indeed, the proof of Lemma 6.2.1 can easily be adapted to give an algorithm to find a tight k-lean tree-decomposition. Similarly, in order to colour the vertices of the tree red or blue we must check for the existence of a K_m minor or a k-block having this vertex as a home node, both of which can be done algorithmically (See [114] and [34]). However, we note that the running time of such an algorithm, or at least a naive implementation of one, would have run time $\sim |V(G)|^{f(r)}$ for some function of the size of the topological minor K_r we are excluding, whereas the algorithm of Grohe and Marx has run time $g(r)|V(G)|^{O(1)}$, which should be much better for large values of r.

Part II

Infinitary combinatorics

Chapter 7

A counterexample to the reconstruction conjecture for locally finite trees

7.1 Introduction

We say that two graphs G and H are *(vertex-)*hypomorphic if there exists a bijection φ between the vertices of G and H such that the induced subgraphs $G - v$ and $H - \varphi(v)$ are isomorphic for each vertex v of G. Any such bijection is called a *hypomorphism*. We say that a graph G is *reconstructible* if $H \cong G$ for every H hypomorphic to G. The following conjecture, attributed to Kelly and Ulam, is perhaps one of the most famous unsolved problems in the theory of graphs.

Conjecture 7.1.1 (The Reconstruction Conjecture)**.** *Every finite graph with at least three vertices is reconstructible.*

For an overview of results towards the Reconstruction Conjecture for finite graphs see the survey of Bondy and Hemminger [23]. Harary [75] proposed the Reconstruction Conjecture for infinite graphs, however Fisher [62] found a counterexample, which was simplified to the following counterexample by Fisher, Graham and Harary [63]: consider the infinite tree G in which every vertex has countably infinite degree, and the graph H formed by taking two disjoint copies of G, which we will write as $G \sqcup G$. For each vertex v of G, the induced subgraph $G - v$ is isomorphic to $G \sqcup G \sqcup \cdots$, a disjoint union of countably many copies of G, and similarly for each vertex w of H, the induced subgraph $H - w$ is isomorphic to $G \sqcup G \sqcup \cdots$ as well. Therefore, any bijection from $V(G)$ to $V(H)$ is a hypomorphism, but G and H are clearly not isomorphic. Hence, the tree G is not reconstructible.

These examples, however, contain vertices of infinite degree. Regarding locally finite graphs, Harary, Schwenk and Scott [76] showed that there exists a non-reconstructible locally finite forest. However, they conjectured that the Reconstruction Conjecture should hold for locally finite trees.

Conjecture 7.1.2 (The Harary-Schwenk-Scott Conjecture)**.** *Every locally finite tree is reconstructible.*

This conjecture has been verified in a number of special cases. Kelly [84] showed that finite trees on at least three vertices are reconstructible. Bondy and Hemminger [22] showed that every tree with at least two but a finite number of ends is reconstructible, and Thomassen [125] showed that this also holds for one-ended trees. Andreae [10] proved that also every tree with countably many ends is reconstructible.

A survey of Nash-Williams [103] on the subject of reconstruction problems in infinite graphs gave the following three main open problems in this area, which have remained open until now.

Problem 1 (Nash-Williams)**.** *Is every locally finite connected infinite graph reconstructible?*

Problem 2 (Nash-Williams)**.** *If two infinite trees are hypomorphic, are they also isomorphic?*

Problem 3 (Halin)**.** *If G and H are hypomorphic, do there exist embeddings $G \hookrightarrow H$ and $H \hookrightarrow G$?*

Problem 2 has been emphasized in Andreae's [12], which contains partial affirmative results on Problem 2. A positive answer to Problem 1 or 2 would verify the Harary-Schwenk-Scott Conjecture. In this paper we construct a pair of trees which are not only a counterexample to the Harary-Schwenk-Scott Conjecture, but also answer the three questions of Nash-Williams and Halin in the negative. Our counterexample will in fact have bounded degree.

Theorem 7.1.3. *There are two (vertex)-hypomorphic infinite trees T and S with maximum degree three such that there is no embedding $T \hookrightarrow S$ or $S \hookrightarrow T$.*

Our example also provides a strong answer to a question by Andreae [11] about edge-reconstructibility. Two graphs G and H are *edge-hypomorphic* if there exists a bijection $\varphi\colon E(G) \to E(H)$ such that $G - e \cong H - \varphi(e)$ for each $e \in E(G)$. A graph G is *edge-reconstructible* if $H \cong G$ for all H edge-hypomorphic to G. In [11] Andreae constructed countable forests which are not edge-reconstructible, but conjectured that no locally finite such examples can exist.

Problem 4 (Andreae)**.** *Is every locally finite graph with infinitely many edges edge-reconstructible?*

Our example answers Problem 4 in the negative: the trees T and S we construct for Theorem 7.1.3 will also be edge-hypomorphic. Besides answering Problem 4, this appears to be the first known example of two non-isomorphic graphs that are *simultaneously* vertex- and edge-hypomorphic.

The Reconstruction Conjecture has also been considered for general locally finite graphs. Nash-Williams [102] showed that any locally finite graph with at least three, but a finite number of ends is reconstructible, and in [104], he established the same result for two-ended graphs. The following problems, also from [103], remain open:

Problem 5 (Nash-Williams)**.** *Is every locally finite graph with exactly one end reconstructible?*

Problem 6 (Nash-Williams)**.** *Is every locally finite graph with countably many ends reconstructible?*

In a paper in preparation [29], we will extend the methods developed in the present paper to also construct counterexamples to Problems 5 and 6.

This paper is organised as follows. In the next section we will give a short, high-level overview of our counterexample to the Harary-Schwenk-Scott Conjecture. In Section 7.3, we will develop the technical tools necessary for our construction, and in Section 7.4, we will prove Theorem 7.1.3.

For standard graph theoretical concepts we follow the notation in [43].

7.2 Sketch of the construction

In this section we sketch the main ideas of the construction. For the sake of simplicity we only indicate how to ensure that the trees T and S are vertex-hypomorphic and non-isomorphic, but not that they are edge-hypomorphic as well, nor that neither embeds into the other.

Our plan is to build the trees T and S recursively, where at each step of the construction we ensure for some vertex v already chosen for T that there is a corresponding vertex w of S with $T - v \cong S - w$, or vice versa. This will ensure that by the end of the construction, the trees we have built are hypomorphic.

More precisely, at step n we will construct subtrees T_n and S_n of our eventual trees, where some of the leaves of these subtrees have been coloured in two colours, say red and blue. We will only further extend the trees from these coloured leaves, and we will extend from leaves of the same colour in the same way.

That is, the plan is that there should be two further rooted trees R and B such that T can be obtained from T_n by attaching copies of R at all red leaves and copies of B at all blue leaves, and S can be obtained from S_n in the same way. At step n, however, we do not yet know what these trees R and B will eventually be.

Nevertheless, we can ensure that the induced subgraphs, $T - v$ and $S - w$, of the vertices we have dealt with so far really will match up. More precisely, by step n we have vertices $x_1, \ldots, x_n$ of T_n and $y_1, \ldots, y_n$ of S_n for which we intend that $T - x_j$ should be isomorphic to $S - y_j$ for each j. We ensure this by arranging that for each j there is an isomorphism from $T_n - x_j$ to $S_n - y_j$ which preserves the colours of the leaves.

The T_n will be nested, and we will take T to be the union of all of them; similarly the S_n will be nested and we take S to be the union of all of them.

There is a trick to ensure that T and S do not end up being isomorphic. First we ensure, for each n, that there is no isomorphism from T_n to S_n. We also ensure that the part of T or S beyond any coloured leaf of T_n or S_n begins with a long non-branching path (called a *bare* path), longer than any such path appearing in T_n or S_n. Call the length of these long paths k_{n+1}.

Suppose now for a contradiction that there is an isomorphism from T to S. Then there must exist some large n such that the isomorphism sends some vertex t of T_n to a vertex s of S_n. However, T_n is the component of T containing t after all bare paths of length k_{n+1} have been removed[1], and so it must map isomorphically onto the component of S containing s after all bare paths of length k_{n+1} have been removed, namely onto S_n. However, there is no isomorphism from T_n onto S_n, so we have the desired contradiction.

Figure 7.1: A first approximation of T_{n+1} on the left, and S_{n+1} on the right. All dotted lines are non-branching paths of length k_{n+1}.

Suppose now that we have already constructed T_n and S_n and wish to construct T_{n+1} and S_{n+1}. Suppose further that we are given a vertex v of T_n for which we wish to find a partner w in S_{n+1} so that $T - v$ and $S - w$ are isomorphic. We begin by building a tree $\hat{T}_n \ncong T_n$ which

[1] Here and throughout this section we will omit minor technical details for brevity.

has some vertex w such that $T_n - v \cong \hat{T}_n - w$. This can be done by taking the components of $T_n - v$ and arranging them suitably around the new vertex w.

We will take S_{n+1} to include S_n and $\hat{T}_n$, with the copies of red and blue leaves in $\hat{T}_n$ also coloured red and blue respectively. As indicated on the right in Figure 7.1, we add paths of length k_{n+1} to some blue leaf b of S_n and to some red leaf $\hat{r}$ of $\hat{T}_n$ and join these paths at their other endpoints by some edge e_n. We also join two new leaves y and g to the endvertices of e_n. We colour the leaf y yellow and the leaf g green (to avoid confusion with the red and blue leaves from step n, we take the two colours applied to the leaves in step $n+1$ to be yellow and green).

To ensure that $T_{n+1} - v \cong S_{n+1} - w$, we take T_{n+1} to include T_n together with a copy $\hat{S}_n$ of S_n, coloured appropriately and joined up in the same way, as indicated on the left in Figure 7.1.

The only problem up to this point is that we have not been faithful to our intention of extending in the same way at each red or blue leaf of T_n and S_n. Thus, we now copy the same subgraph appearing beyond r in Fig. 7.1, including its coloured leaves, onto all the other red leaves of S_n and T_n. Similarly we copy the subgraph appearing beyond the blue leaf b of S_n onto all other blue leaves of S_n and T_n.

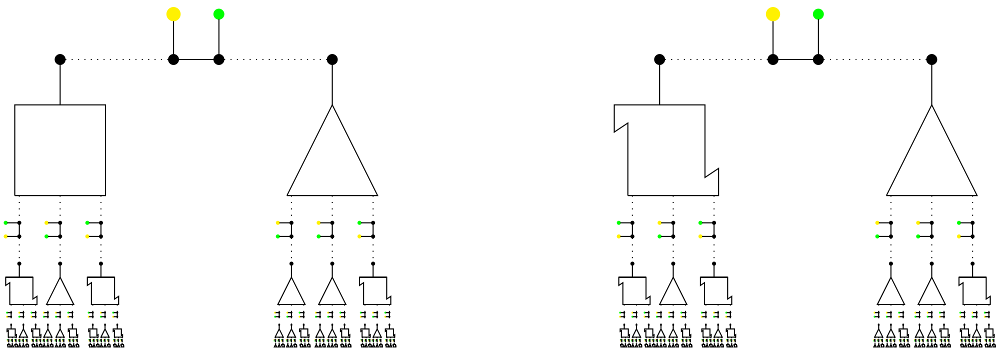

Figure 7.2: A sketch of T_{n+1} and S_{n+1} after countably many steps.

At this point, we would have kept our promise of adding the same thing behind every red and blue leaf of T_n and S_n, and hence would have achieved $T_{n+1} - x_j \cong S_{n+1} - y_j$ for all $j \leqslant n$. However, by gluing the additional copies to blue and red leaves of T_n and S_n, we now have ruined the isomorphism between $T_{n+1} - v$ and $S_{n+1} - w$. In order to repair this, we also have to copy the graphs appearing beyond r and b in Fig. 7.1 respectively onto all red and blue leaves of $\hat{S}_n$ and $\hat{T}_n$. This repairs $T_{n+1} - v \cong S_{n+1} - w$, but again violates our initial promises. In this way, we keep adding, step by step, further copies of the graphs appearing beyond r and b in Fig. 7.1 respectively onto all red and blue leaves of everything we have constructed so far.

At every step we preserved the colours of leaves in all newly added copies, so we get new red leaves and blue leaves, and we continue the process of copying onto those new leaves as well. After countably many steps we have dealt with all red or blue leaves. We take these new trees to be S_{n+1} and T_{n+1}. They are non-isomorphic, since after removing all long bare paths, T_{n+1} contains T_n as a component, whereas S_{n+1} does not.

Figure 7.2 shows how T_{n+1} and S_{n+1} might appear. We have now fulfilled our intention of sticking the same thing onto all red leaves and the same thing onto all blue leaves, but we have also ensured that $T_{n+1} - v \cong S_{n+1} - w$, as desired.

7.3 Closure with respect to promises

In this section, we formalise the ideas set forth in the proof sketch of how to extend a graph so that it looks the same beyond certain sets of leaves.

Given a directed edge $\vec{e} = \vec{xy}$ in some forest $G = (V, E)$, we denote by $G(\vec{e})$ the unique component of $G - e$ containing the vertex y. We think of $G(\vec{e})$ as a rooted tree with root y. As indicated in the previous section, in order to make T and S hypomorphic at the end, we will often have to guarantee $S(\vec{e}) \cong T(\vec{f})$ for certain pairs of edges $\vec{e}$ and $\vec{f}$.

Definition 7.3.1 (Promise structure)**.** *A promise structure* $\mathcal{P} = \left(G, \vec{P}, \mathcal{L}\right)$ *consists of:*

- *a forest* G*,*
- $\vec{P} = \{\vec{p_i} \colon i \in I\}$ *a set of directed edges* $\vec{P} \subseteq \vec{E}(G)$*, and*
- $\mathcal{L} = \{L_i \colon i \in I\}$ *a set of pairwise disjoint sets of leaves of* G*.*

Often, when the context is clear, we will not make a distinction between $\mathcal{L}$ and the set $\bigcup_i L_i$, for notational convenience.

We will call an edge $\vec{p_i} \in \vec{P}$ a *promise edge*, and leaves $\ell \in L_i$ *promise leaves.* A promise edge $\vec{p_i} \in \vec{P}$ is called a *placeholder-promise* if the component $G(\vec{p_i})$ consists of a single leaf $c_i \in L_i$, then called a *placeholder-leaf.* We write

$$\mathcal{L}_p = \{L_i \colon i \in I,\ \vec{p_i} \text{ a placeholder-promise}\} \text{ and } \mathcal{L}_q = \mathcal{L} \setminus \mathcal{L}_p.$$

Given a leaf ℓ in G, there is a unique edge $q_\ell \in E(G)$ incident with ℓ, and this edge has a natural orientation $\vec{q_\ell}$ towards ℓ. Informally, we think of the 'promise' $\ell \in L_i$ as saying that if we extend G to a graph $H \supset G$, we will do so in such a way that $H(\vec{q_\ell}) \cong H(\vec{p_i})$. Given a promise structure $\mathcal{P} = \left(G, \vec{P}, \mathcal{L}\right)$, we would like to construct a graph $H \supset G$ which satisfies all the promises in $\mathcal{P}$. This will be done by the following kind of extension.

Definition 7.3.2 (Leaf extension)**.** *Given an inclusion* $H \supseteq G$ *of forests and a set* L *of leaves of* G*,* H *is called a* leaf extension*, or more specifically an* L-extension, of G*, if:*

- *every component of* H *contains precisely one component of* G*, and*
- *for every vertex* $h \in H \setminus G$ *and every vertex* $g \in G$ *in the same component as* h*, the unique* $g - h$ *path in* H *meets* L*.*

In the remainder of this section we describe a construction of a forest $\mathrm{cl}(G)$ which has the following properties.

Proposition 7.3.3. *Let* G *be a forest and let* $\left(G, \vec{P}, \mathcal{L}\right)$ *be a promise structure. Then there is a forest* $\mathrm{cl}(G)$ *such that:*

(cl.1) $\mathrm{cl}(G)$ *is an* $\mathcal{L}_q$*-extension of* G*, and*

(cl.2) *for every* $\vec{p_i} \in \vec{P}$ *and all* $\ell \in L_i$*,*

$$\mathrm{cl}(G)(\vec{p_i}) \cong \mathrm{cl}(G)(\vec{q_\ell})$$

are isomorphic as rooted trees.

We first describe the construction of $\mathrm{cl}(G)$, and then verify the properties asserted in Proposition 7.3.3. Let us define a sequence of promise structures $\left(H^{(i)}, \vec{P}, \mathcal{L}^{(i)}\right)$ as follows. We set $\left(H^{(0)}, \vec{P}, \mathcal{L}^{(0)}\right) = \left(G, \vec{P}, \mathcal{L}\right)$. We construct a sequence of graphs

$$G = H^{(0)} \subseteq H^{(1)} \subseteq H^{(2)} \subseteq \cdots,$$

and each $H^{(n)}$ will get a promise structure whose set of promise edges is equal to $\vec{P}$ again, yet whose set of promise leaves depends on n as follows: given $\left(H^{(n)}, \vec{P}, \mathcal{L}^{(n)}\right)$, we construct $H^{(n+1)}$ by gluing, for each i, at every promise leaf $\ell \in L_i^{(n)}$ a rooted copy of $G(\vec{p_i})$. As promise leaves for $H^{(n+1)}$ we take all promise leaves from the newly added copies of $G(\vec{p_i})$. That is, if a leaf $\ell \in G(\vec{p_i})$ was such that $\ell \in L_j$, then every copy of that leaf will be in $L_j^{(n+1)}$.

Formally, suppose that $\left(G, \vec{P}, \mathcal{L}\right)$ is a promise structure. For each $\vec{p_i} \in \vec{P}$ let $C_i = G(\vec{p_i})$ and let c_i be the root of this tree. If U is a set and H is a graph, then we denote by $U \times H$ the graph whose vertices are pairs (u, v) with $u \in U$ and v a vertex of H, and with an edge from (u, v) to (u, w) whenever vw is an edge of H. Let $\left(H^{(0)}, \vec{P}, \mathcal{L}^{(0)}\right) = \left(G, \vec{P}, \mathcal{L}\right)$ and given $\left(H^{(n)}, \vec{P}, \mathcal{L}^{(n)}\right)$ let us define:

- $H^{(n+1)}$ to be the quotient of $H^{(n)} \sqcup \bigsqcup_{i \in I}(L_i^{(n)} \times C_i)$ w.r.t. the relation
$$l \sim (l, c_i) \text{ for } l \in L_i^{(n)} \in \mathcal{L}^{(n)}.$$
- $\mathcal{L}^{(n+1)} = \left\{L_i^{(n+1)} : i \in I\right\}$ with $L_i^{(n+1)} = \bigcup_{j \in I} L_j^{(n)} \times (C_j \cap L_i)$.

There is a sequence of natural inclusions $G = H^{(0)} \subseteq H^{(1)} \subseteq \cdots$ and we define $\mathrm{cl}(G)$ to be the direct limit of this sequence.

Definition 7.3.4 (Promise-respecting map)**.** *Let G be a forest, $F^{(1)}$ and $F^{(2)}$ be leaf extensions of G, and $\mathcal{P}^{(1)} = \left(F^{(1)}, \vec{P}, \mathcal{L}^{(1)}\right)$ and $\mathcal{P}^{(2)} = \left(F^{(2)}, \vec{P}, \mathcal{L}^{(2)}\right)$ be promise structures with $\vec{P} \subseteq \vec{E}(G)$. Suppose $X^{(1)} \subseteq V(F^{(1)})$ and $X^{(2)} \subseteq V(F^{(2)})$.*

A bijection $\varphi\colon X^{(1)} \to X^{(2)}$ is $\vec{P}$-respecting *(with respect to $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$) if the image of $L_i^{(1)} \cap X^{(1)}$ under φ is $L_i^{(2)} \cap X^{(2)}$ for all i.*

Since both promise structures $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$ refer to the same edge set $\vec{P}$, we can think of them as defining a $|\vec{P}|$-colouring on some sets of leaves. Then a mapping is $\vec{P}$-respecting if it preserves leaf colours.

Lemma 7.3.5. *Let $\left(G, \vec{P}, \mathcal{L}\right)$ be a promise structure and let $G = H^{(0)} \subseteq H^{(1)} \subseteq \cdots$ be as defined above. Then the following statements hold:*

- *$H^{(n)}$ is an $\mathcal{L}_q$-extension of G for all n,*
- *$\Delta(H^{(n+1)}) = \Delta(H^{(n)})$ for all n, and*
- *For each $\ell \in L_i \in \mathcal{L}$ there exists a sequence of $\vec{P}$-respecting rooted isomorphisms $\varphi_{\ell,n}\colon H^{(n)}(\vec{p_i}) \to H^{(n+1)}(\vec{q_\ell})$ such that $\varphi_{\ell,n+1}$ extends $\varphi_{\ell,n}$ for all $n \in \mathbb{N}$.*

Proof. The first two statements are clear. We will prove the third by induction on n. To construct $H^{(1)}$ from G, we glued a rooted copy of $G(\vec{p_i})$ to each $\ell \in L_i$, keeping all copies of promise leaves. Hence, for any given $\ell \in L_i$, the natural isomorphism $\varphi_{\ell,0}\colon G(\vec{p_i}) \to H^{(1)}(\vec{q_\ell})$ is $\vec{P}$-respecting as desired.

Now suppose that $\varphi_{\ell,n}$ exists for all $\ell \in \mathcal{L}$. To form $H^{(n+1)}(\vec{p_i})$, we glued on a copy of $G(\vec{p_i})$ to each $\ell \in L_i^{(n)} \cap H^{(n)}(\vec{p_i})$, and to construct $H^{(n+2)}(\vec{q_\ell})$, we glued on a copy of $G(\vec{p_i})$ to each $\ell \in L_i^{(n+1)} \cap H^{(n+1)}(\vec{q_\ell})$, in both cases keeping all copies of promise leaves.

Therefore, since $\varphi_{\ell,n}$ was a $\vec{P}$-respecting rooted isomorphism from $H^{(n)}(\vec{p_i})$ to $H^{(n+1)}(\vec{q_\ell})$, we can combine the individual isomorphisms between the newly added copies of $G(\vec{p_i})$ with $\varphi_{\ell,n}$ to form $\varphi_{\ell,n+1}$. □

We can now complete the proof of Proposition 7.3.3.

Proof of Proposition 7.3.3. First, we note that $G \subseteq \mathrm{cl}(G)$, and since each $H^{(n)}$ is an $\mathcal{L}_q$-extension of G for all n, so is $\mathrm{cl}(G)$. Also, since each $H^{(n)}$ is a forest it follows that $\mathrm{cl}(G)$ is a forest.

Let us show that $\mathrm{cl}(G)$ satisfies property (cl.2). Since we have the sequence of inclusions $G = H^{(0)} \subseteq H^{(1)} \subseteq \ldots$, it follows that $\mathrm{cl}(G)(\vec{q_\ell})$ is the direct limit of the sequence $H^{(0)}(\vec{q_\ell}) \subseteq H^{(1)}(\vec{q_\ell}) \subseteq \cdots$ and also $\mathrm{cl}(G)(\vec{p_i})$ is the direct limit of the sequence $H^{(0)}(\vec{p_i}) \subseteq H^{(1)}(\vec{p_i}) \subseteq \cdots$. By Lemma 7.3.5 there is a sequence of rooted isomorphisms $\varphi_{\ell,n} \colon H^{(n)}(\vec{p_i}) \to H^{(n+1)}(\vec{q_\ell})$ such that $\varphi_{\ell,n+1}$ extends $\varphi_{\ell,n}$, so $\varphi_\ell = \bigcup_n \varphi_{\ell,n}$ is the required isomorphism. □

We remark that it is possible to show that $\mathrm{cl}(G)$ is in fact determined, uniquely up to isomorphism, by the properties (cl.1) and (cl.2). Also we note that since each $H^{(n)}$ has the same maximum degree as G, it follows that $\Delta(\mathrm{cl}(G)) = \Delta(G)$.

There is a natural promise structure on $\mathrm{cl}(G)$ given by the placeholder promises in $\vec{P}$ and their corresponding promise leaves. In the construction sketch from Section 7.2, these leaves corresponded to the yellow and green leaves. We now show how to keep track of the placeholder promises when taking the closure of a promise structure.

Note that if $\vec{p_i}$ is a placeholder promise, then for each $\left(H^{(n)}, \mathcal{P}, \mathcal{L}^{(n)}\right)$ we have $L_i^{(n)} \supseteq L_i^{(n-1)}$. Indeed, for each leaf in $L_i^{(n-1)}$ we glue a copy of the component c_i together with the associated promises on the leaves in this component. However, c_i is just a single vertex, with a promise corresponding to $\vec{p_i}$, and hence $L_i^{(n)} \supseteq L_i^{(n-1)}$. For every placeholder promise $\vec{p_i} \in \vec{P}$ we define $\mathrm{cl}(L_i) = \bigcup_n L_i^{(n)}$.

Definition 7.3.6 (Closure of a promise structure)**.** *The* closure *of the promise structure* $(G, \mathcal{P}, \mathcal{L})$ *is the promise structure* $\mathrm{cl}(\mathcal{P}) = \left(\mathrm{cl}(G), \mathrm{cl}(\vec{P}), \mathrm{cl}(\mathcal{L})\right)$, *where:*

- $\mathrm{cl}(\vec{P}) = \left\{\vec{p_i} \colon \vec{p_i} \in \vec{P} \text{ is a placeholder-promise}\right\}$, *and*
- $\mathrm{cl}(\mathcal{L}) = \{\mathrm{cl}(L_i) \colon \vec{p_i} \in \vec{P} \text{ is a placeholder-promise}\}$.

We note that, since each isomorphism $\varphi_{\ell,n}$ from Lemma 7.3.5 was $\vec{P}$-respecting, it is possible to strengthen Proposition 7.3.3 in the following way.

Proposition 7.3.7. *Let G be a forest and let $\left(G, \vec{P}, \mathcal{L}\right)$ be a promise structure. Then the forest* $\mathrm{cl}(G)$ *satisfies:*

(cl.3) *for every $\vec{p_i} \in \vec{P}$ and every $\ell \in L_i$,*

$$\mathrm{cl}(G)(\vec{p_i}) \cong \mathrm{cl}(G)(\vec{q_\ell})$$

are isomorphic as rooted trees, and this isomorphism is $\mathrm{cl}(\vec{P})$*-respecting with respect to* $\mathrm{cl}(\mathcal{P})$.

Proof. Since each isomorphism $\varphi_{\ell,n} \colon H^{(n)}(\vec{p_i}) \to H^{(n+1)}(\vec{q_\ell})$ in Proposition 7.3.5 is $\vec{P}$-respecting, we have

$$\varphi_{\ell,n}\left(L_i^{(n)} \cap H^{(n)}(\vec{p_i})\right) = L_i^{(n+1)} \cap H^{(n+1)}(\vec{q_\ell}).$$

For each placeholder promise we have that $\mathrm{cl}(L_i) = \bigcup_n L_i^{(n)}$, and so it follows that

$$\mathrm{cl}(L_i) \cap \mathrm{cl}(G)(\vec{q_\ell}) = \bigcup_n \left(L_i^{(n)} \cap H^{(n)}(\vec{q_\ell})\right)$$

and

$$\mathrm{cl}(L_i) \cap \mathrm{cl}(G)(\vec{p_i}) = \bigcup_n \left(L_i^{(n)} \cap H^{(n)}(\vec{p_i})\right).$$

From this it follows that $\varphi_\ell = \bigcup_n \varphi_{l,n}$ is a $\mathrm{cl}(\vec{P})$-respecting isomorphism between $\mathrm{cl}(G)(\vec{p_i})$ and $\mathrm{cl}(G)(\vec{q_\ell})$ as rooted trees. □

It is precisely this property (cl.3) of the promise closure that will allow us, in Claim 7.4.14 below, to maintain partial hypomorphisms during our recursive construction.

7.4 The construction

In this section we construct two hypomorphic locally finite trees neither of which embed into the other, establishing our main theorem announced in the introduction.

7.4.1 Preliminary definitions

Definition 7.4.1 (Bare path)**.** *A path $P = v_0, v_1, \dots, v_n$ in a graph G is called a* bare path *if $\deg_G(v_i) = 2$ for all internal vertices v_i for $0 < i < n$. The path P is a* maximal bare path *(or* maximally bare*) if in addition $\deg_G(v_0) \neq 2 \neq \deg_G(v_n)$. An infinite path $P = v_0, v_1, v_2, \dots$ is* maximally bare *if $\deg_G(v_0) \neq 2$ and $\deg_G(v_i) = 2$ for all $i \geqslant 1$.*

Lemma 7.4.2. *Let T be a tree and $e \in E(T)$. If every maximal bare path in T has length at most $k \in \mathbb{N}$, then every maximal bare path in $T - e$ has length at most $2k$.*

Proof. We first note that every maximal bare path in $T - e$ has finite length, since any infinite bare path in $T_n - e$ would contain a subpath which is an infinite bare path in T. If $P = \{x_0, x_1, \dots, x_n\}$ is a maximal bare path in $T - e$ which is not a subpath of any maximal bare path in T, then there is at least one $1 \leqslant i \leqslant n-1$ such that e is adjacent to x_i, and since T was a tree, x_i is unique. Therefore, both $\{x_0, x_1, \dots, x_i\}$ and $\{x_i, x_{i+1}, \dots, x_n\}$ are maximal bare paths in T. By assumption both i and $n - i$ are at most k, and so the length of P is at most $2k$, as claimed. □

Definition 7.4.3 (Bare extension)**.** *Given a forest G, a subset B of leaves of G, and a component T of G, we say that a tree $\hat{T} \supset T$ is a* bare extension *of T at B to length k if $\hat{T}$ can be obtained from T by adjoining, at each vertex $l \in B \cap V(T)$, a new path of length k starting at l and a new leaf whose only neighbour is l.*

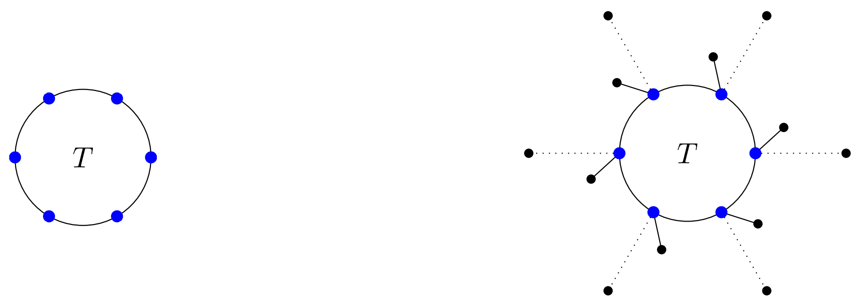

A tree T with designated leaf set B. A bare extension of T at B.

Figure 7.3: Building a bare extension of a tree T at B to length k. All dotted lines are maximal bare paths of length k.

Note that the new leaves attached to each $l \in B$ ensure that the paths of length k are indeed maximal bare paths.

Definition 7.4.4 (k-ball)**.** *For G a subgraph of H, the k-ball $\mathrm{Ball}_H(G, k)$ is the induced subgraph of H on the set of vertices within distance k of some vertex of G.*

Definition 7.4.5 (Binary tree)**.** *For $k \geqslant 1$, the* binary tree of height k *is the unique rooted tree on $2^k - 1 = 1 + 2 + \cdots + 2^{k-1}$ vertices such that the root has degree 2, there are 2^{k-1} leaves, and all other vertices have degree 3. By a* binary tree *we mean a binary tree of height k for some $k \in \mathbb{N}$.*

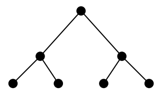

Figure 7.4: The binary tree of height 3.

7.4.2 The back-and-forth construction

We prove the following theorem.

Theorem 7.4.6. *There are two (vertex-)hypomorphic infinite trees T and S with maximum degree 3 such that there is no embedding $T \hookrightarrow S$ or $S \hookrightarrow T$.*

To do this we shall recursively construct, for each $n \in \mathbb{N}$,

- disjoint (possibly infinite) rooted trees T_n and S_n,
- disjoint (possibly infinite) sets R_n and B_n of leaves of the forest $T_n \sqcup S_n$,
- finite sets $X_n \subset V(T_n)$ and $Y_n \subset V(S_n)$, and bijections $\varphi_n \colon X_n \to Y_n$,
- a family of isomorphisms $\mathcal{H}_n = \{h_{n,x} \colon T_n - x \to S_n - \varphi_n(x) \colon x \in X_n\}$,
- strictly increasing sequences of integers $k_n \geqslant 2$ and $b_n \geqslant 3$,

such that (letting all objects indexed by -1 be the empty set) for all $n \in \mathbb{N}$:

(†1) $T_{n-1} \subset T_n$ and $S_{n-1} \subset S_n$ as induced subgraphs,

(†2) the vertices of T_n and S_n all have degree at most 3,

(†3) the root of T_n is in R_n and the root of S_n is in B_n,

(†4) all binary trees appearing as subgraphs of $T_n \sqcup S_n$ are finite and have height at most b_n,

(†5) all bare paths in $T_n \sqcup S_n$ are finite and have length at most k_n,

(†6) $\mathrm{Ball}_{T_n}(T_{n-1}, k_{n-1} + 1)$ is a bare extension of T_{n-1} at $R_{n-1} \cup B_{n-1}$ to length $k_{n-1} + 1$ and does not meet $R_n \cup B_n$,

(†7) $\mathrm{Ball}_{S_n}(S_{n-1}, k_{n-1} + 1)$ is a bare extension of S_{n-1} at $R_{n-1} \cup B_{n-1}$ to length $k_{n-1} + 1$ and does not meet $R_n \cup B_n$,

(†8) there is no embedding from T_n into any bare extension of S_n at $R_n \cup B_n$ to any length, nor from S_n into any bare extension of T_n at $R_n \cup B_n$ to any length,

(†9) any embedding of T_n into a bare extension of T_n at $R_n \cup B_n$ to any length fixes the root of T_n and has image T_n,

(†10) any embedding of S_n into a bare extension of S_n at $R_n \cup B_n$ to any length fixes the root of S_n and has image S_n,

(†11) there are enumerations $V(T_n) = \{t_j : j \in J_n\}$ and $V(S_n) = \{s_j : j \in J_n\}$ such that

- $J_{n-1} \subset J_n \subset \mathbb{N}$,
- $\{t_j : j \in J_n\}$ extends the enumeration $\{t_j : j \in J_{n-1}\}$ of $V(T_{n-1})$, and similarly for $\{s_j : j \in J_n\}$,
- $|\mathbb{N} \setminus J_n| = \infty$,
- $\{0, 1, \dots, n\} \subset J_n$,

(†12) $\{t_j, s_j : j \leqslant n\} \cap (R_n \cup B_n) = \emptyset$,

(†13) the finite sets of vertices X_n and Y_n satisfy $|X_n| = n = |Y_n|$, and

- $X_{n-1} \subset X_n$ and $Y_{n-1} \subset Y_n$,
- $\varphi_n \restriction X_{n-1} = \varphi_{n-1}$,
- $\{t_j : j \leqslant n\} \subset X_{2n+1}$ and $\{s_j : j \leqslant n\} \subset Y_{2(n+1)}$,
- $(X_n \cup Y_n) \cap (R_n \cup B_n) = \emptyset$,

(†14) the families of isomorphisms $\mathcal{H}_n$ satisfy

- $h_{n,x} \restriction (T_{n-1} - x) = h_{n-1,x}$ for all $x \in X_{n-1}$,
- the image of $R_n \cap V(T_n)$ under $h_{n,x}$ is $R_n \cap V(S_n)$, and
- the image of $B_n \cap V(T_n)$ under $h_{n,x}$ is $B_n \cap V(S_n)$ for all $x \in X_n$.

7.4.3 The construction yields the desired non-reconstructible trees.

By property (†1), we have $T_0 \subset T_1 \subset T_2 \subset \cdots$ and $S_0 \subset S_1 \subset S_2 \subset \cdots$. Let T and S be the union of the respective chains. It is clear that T and S are trees, and that as a consequence of (†2), both trees have maximum degree 3.

We claim that the map $\varphi = \bigcup_n \varphi_n$ is a hypomorphism between T and S. Indeed, it follows from (†11) and (†13) that φ is a well-defined bijection from $V(T)$ to $V(S)$. To see that φ is a hypomorphism, consider any vertex x of T. This vertex appears as some t_j in our enumeration of $V(T)$, so by (†14) the map

$$h_x := \bigcup_{n > 2j} h_{n,x} \colon T - x \to S - \varphi(x)$$

is an isomorphism between $T - x$ and $S - \varphi(x)$.

Now suppose for a contradiction that $f \colon T \hookrightarrow S$ is an embedding of T into S. Then $f(t_0)$ is mapped into S_n for some $n \in \mathbb{N}$. Properties (†5) and (†6) imply that after deleting all maximal bare paths in T of length $> k_n$, the connected component of t_0 is a bare extension of T_n to length 0. Further, by (†7), $\mathrm{Ball}_S(S_n, k_n + 1)$ is a bare extension of S_n at $R_n \cup B_n$ to length $k_n + 1$. But combining the fact that $f(T_n) \cap S_n \neq \emptyset$ and the fact that T_n does not contain long maximal bare paths, it is easily seen that $f(T_n) \subset \mathrm{Ball}_S(S_n, k_n + 1)$, contradicting (†8).[2]

The case $S \hookrightarrow T$ yields a contradiction in a symmetric fashion, completing the proof.

[2]To get the non-embedding property, we have used (†5)–(†8) at every step n. While at the first glance, properties (†4), (†9) and (†10) do not seem to be needed at this point, they are crucial during the construction to establish (†8) at step $n + 1$. See Claim 7.4.11 below for details.

7.4.4 The base case: there are finite rooted trees T_0 and S_0 satisfying requirements (†1)–(†14).

Choose a pair of non-isomorphic, equally sized trees T_0 and S_0 of maximum degree 3, and pick a leaf each as roots $\mathrm{r}(T_0)$ and $\mathrm{r}(S_0)$ for T_0 and S_0, subject to conditions (†8)–(†10) with $R_0 = \{\mathrm{r}(T_0)\}$ and $B_0 = \{\mathrm{r}(S_0)\}$. A possible choice is given in Fig. 7.5. Here, (†8) is satisfied, because any embedding of T_0 into a bare extension of S_0 has to map the binary tree of height 3 in T_0 to the binary tree in S_0, making it impossible to embed the middle leaf. Properties (†9) and (†10) are similar.

Figure 7.5: A possible choice for finite rooted trees T_0 and S_0.

Let $J_0 = \{0, 1, \ldots, |T_0| - 1\}$ and choose enumerations $V(T_0) = \{t_j : j \in J_0\}$ and $V(S_0) = \{s_j : j \in J_0\}$ with $t_0 \neq \mathrm{r}(T_0)$ and $s_0 \neq \mathrm{r}(S_0)$. This takes care of (†11) and (†12). Finally, (†13) and (†14) are satisfied for $X_0 = Y_0 = \mathcal{H}_0 = \varphi_0 = \emptyset$. Set $k_0 = 2$ and $b_0 = 3$.

7.4.5 The inductive step: set-up

Now, assume that we have constructed trees T_k and S_k for all $k \leqslant n$ such that (†1)–(†14) are satisfied up to n. If $n = 2m$ is even, then we have $\{t_j : j \leqslant m-1\} \subset X_n$, so in order to satisfy (†13) we have to construct T_{n+1} and S_{n+1} such that the vertex t_m is taken care of in our partial hypomorphism. Similarly, if $n = 2m+1$ is odd, then we have $\{s_j : j \leqslant m-1\} \subset Y_n$ and we have to construct T_{n+1} and S_{n+1} such that the vertex s_m is taken care of in our partial hypomorphism. Both cases are symmetric, so let us assume in the following that $n = 2m$ is even.

Now let v be the vertex with the least index in the set $\{t_j : j \in J_n\} \setminus X_n$, i.e.

$$v = t_i \quad \text{for} \quad i = \min\{\ell : t_\ell \in V(T_n) \setminus X_n\}. \tag{7.4.1}$$

Then by assumption (†13), v will be t_m, unless t_m was already in X_n anyway. In any case, since $|X_n| = |Y_n| = n$, it follows from (†11) that $i \leqslant n$, so by (†12), v does not lie in our leaf sets $R_n \cup B_n$, i.e.

$$v \notin R_n \cup B_n. \tag{7.4.2}$$

In the next sections, we will demonstrate how to to obtain trees $T_{n+1} \supset T_n$ and $S_{n+1} \supset S_n$ with $X_{n+1} = X_n \cup \{v\}$ and $Y_{n+1} = Y_n \cup \{\varphi_{n+1}(v)\}$ satisfying (†1)—(†10) and (†13)–(†14).

After we have completed this step, since $|\mathbb{N} \setminus J_n| = \infty$, it is clear that we can extend our enumerations of T_n and S_n to enumerations of T_{n+1} and S_{n+1} as required, making sure to first list some new elements that do not lie in $R_{n+1} \cup B_{n+1}$. This takes care of (†11) and (†12) and completes the recursion step $n \mapsto n+1$.

7.4.6 The inductive step: construction

Given the two trees T_n and S_n, we extend each of them through their roots as indicated in Figure 7.6 to trees $\tilde{T}_n$ and $\tilde{S}_n$ respectively. The trees T_{n+1} and S_{n+1} will be obtained as components of the promise closure of the forest $G_n = \tilde{T}_n \sqcup \tilde{S}_n$ with respect to the coloured promise edges.

Since v is not the root of T_n, there is a first edge e on the unique path in T_n from v to the root.

$$\text{This edge we also call } e(v). \tag{7.4.3}$$

Then $T_n - e$ has two connected components: one that contains the root of T_n which we name $T_n(r)$, and one that contains v which we name $T_n(v)$.

Since every maximal bare path in T_n has length at most k_n by (†5), it follows from Lemma 7.4.2 that all maximal bare paths in $T_n - e$, and so all bare paths in $T_n(r)$ and $T_n(v)$, have bounded length. Let $k = \tilde{k}_n$ be twice the maximum of the length of bare paths in T_n, S_n, $T_n(r)$ and $T_n(v)$, which exists by (†5).

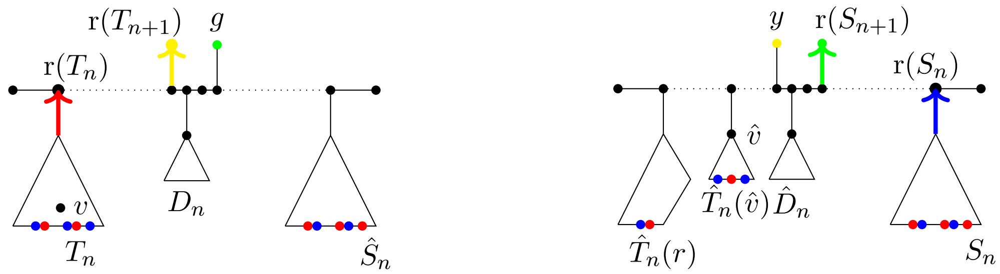

(a) tree $\tilde{T}_n$ The tree $\tilde{S}_n$.

Figure 7.6: All dotted lines are maximal bare paths of length at least $k = \tilde{k}_n$. The trees D_n are binary trees of height $b_n + 3$, hence $D_n \not\leftrightarrow T_n$ and $D_n \not\leftrightarrow S_n$ by ((†4)).

To obtain $\tilde{T}_n$, we extend T_n through its root $\mathrm{r}(T_n) \in R_n$ by a path

$$\mathrm{r}(T_n) = u_0, u_1, \dots, u_{p-1}, u_p = \mathrm{r}\left(\hat{S}_n\right)$$

of length $p = 4(\tilde{k}_n + 1) + 3$, where at its last vertex u_p we glue a rooted copy $\hat{S}_n$ of S_n (via an isomorphism $\hat{w} \leftrightarrow w$), identifying u_p with the root of $\hat{S}_n$.

Next, we add two additional leaves at u_0 and u_p, so that $\deg(\mathrm{r}(T_n)) = 3 = \deg\left(\mathrm{r}\left(\hat{S}_n\right)\right)$. Further, we add a leaf $\mathrm{r}(T_{n+1})$ at u_{2k+2}, which will be our new root for the next tree T_{n+1}; and another leaf g at u_{2k+5}. Finally, we take a copy D_n of a rooted binary tree of height $b_n + 3$ and connect its root via an edge to u_{2k+3}. This completes the construction of $\tilde{T}_n$.

The construction of $\tilde{S}_n$ is similar, but with a twist. For its construction, we extend S_n through its root $\mathrm{r}(S_n) \in B_n$ by a path

$$\mathrm{r}(S_n) = v_p, v_{p-1}, \dots, v_1, v_0 = \mathrm{r}\left(\hat{T}_n(r)\right)$$

of length p, where at its last vertex v_0 we glue a copy $\hat{T}_n(r)$ of $T_n(r)$, identifying v_0 with the root of $\hat{T}_n(r)$. Then, we take a copy $\hat{T}_n(\hat{v})$ of $T_n(v)$ and connect $\hat{v}$ via an edge to v_{k+1}.

$$\text{This edge we call } e(\hat{v}). \tag{7.4.4}$$

Finally, as before, we add two leaves at v_0 and v_p so that $\deg\left(\mathrm{r}\left(\hat{T}_n(r)\right)\right) = 3 = \deg\left(\mathrm{r}(S_n)\right)$. Next, we add a leaf $\mathrm{r}(S_{n+1})$ to v_{2k+5}, which will be our new root for the next tree S_{n+1}; and another leaf y to v_{2k+2}. Finally, we take another copy $\hat{D}_n$ of a rooted binary tree of height $b_n + 3$ and connect its root via an edge to v_{2k+3}. This completes the construction of $\tilde{S}_n$.

By the induction hypothesis, certain leaves of T_n have been coloured with one of the two colours $R_n \cup B_n$, and also some leaves of S_n have been coloured with one of the two colours $R_n \cup B_n$. In the above construction, we colour leaves of $\hat{S}_n$, $\hat{T}_n(r)$ and $\hat{T}_n(\hat{v})$ accordingly:

$$\begin{aligned} \tilde{R}_n &= \Big(R_n \cup \Big\{\hat{w} \in \hat{S}_n \cup \hat{T}_n(r) \cup \hat{T}_n(\hat{v}) \colon w \in R_n\Big\}\Big) \setminus \Big\{\mathrm{r}(T_n), \mathrm{r}\Big(\hat{T}_n(r)\Big)\Big\}, \\ \tilde{B}_n &= \Big(B_n \cup \Big\{\hat{w} \in \hat{S}_n \cup \hat{T}_n(r) \cup \hat{T}_n(\hat{v}) \colon w \in B_n\Big\}\Big) \setminus \Big\{\mathrm{r}(S_n), \mathrm{r}\Big(\hat{S}_n\Big)\Big\}. \end{aligned} \tag{7.4.5}$$

Now put $G_n := \tilde{T}_n \sqcup \tilde{S}_n$ and consider the following promise structure $\mathcal{P} = \Big(G_n, \vec{P}, \mathcal{L}\Big)$ on G_n, consisting of four promise edges $\vec{P} = \{\vec{p}_1, \vec{p}_2, \vec{p}_3, \vec{p}_4\}$ and corresponding leaf sets $\mathcal{L} = \{L_1, L_2, L_3, L_4\}$, as follows:

$$\begin{aligned} &\bullet\ \vec{p}_1 \text{ pointing in } T_n \text{ towards the root } \mathrm{r}(T_n), \text{ with } L_1 = \tilde{R}_n, \\ &\bullet\ \vec{p}_2 \text{ pointing in } S_n \text{ towards the root } \mathrm{r}(S_n), \text{ with } L_2 = \tilde{B}_n, \\ &\bullet\ \vec{p}_3 \text{ pointing in } \tilde{T}_n \text{ towards the root } \mathrm{r}(T_{n+1}), \text{ with } L_3 = \{\mathrm{r}(T_{n+1}), y\}, \\ &\bullet\ \vec{p}_4 \text{ pointing in } \tilde{S}_n \text{ towards the root } \mathrm{r}(S_{n+1}), \text{ with } L_4 = \{\mathrm{r}(S_{n+1}), g\}. \end{aligned} \tag{7.4.6}$$

Note that our construction so far has been tailored to provide us with a $\vec{P}$-respecting isomorphism

$$h \colon \tilde{T}_n - v \to \tilde{S}_n - \hat{v}. \tag{7.4.7}$$

Consider the closure $\mathrm{cl}(G_n)$ with respect to the promise structure $\mathcal{P}$ defined above. Since $\mathrm{cl}(G_n)$ is a leaf-extension of G_n, it has two connected components, just as G_n. We now define

$$\begin{aligned} T_{n+1} &= \text{ the component containing } T_n \text{ in } \mathrm{cl}(G_n), \text{ and} \\ S_{n+1} &= \text{ the component containing } S_n \text{ in } \mathrm{cl}(G_n). \end{aligned} \tag{7.4.8}$$

It follows that $\mathrm{cl}(G_n) = T_{n+1} \sqcup S_{n+1}$ and $\hat{v} \in V(S_{n+1})$. Further, since $\vec{p}_3$ and $\vec{p}_4$ are placeholder promises, $\mathrm{cl}(G)$ carries a corresponding promise structure, see Definition 7.3.6. We define

$$R_{n+1} = \mathrm{cl}(L_3) \text{ and } B_{n+1} = \mathrm{cl}(L_4). \tag{7.4.9}$$

Lastly, we set

$$\begin{aligned} X_{n+1} &= X_n \cup \{v\}, \\ Y_{n+1} &= Y_n \cup \{\hat{v}\}, \text{ and} \\ \varphi_{n+1} &= \varphi_n \cup \{(v, \hat{v})\}, \end{aligned} \tag{7.4.10}$$

and put

$$k_{n+1} = 2\tilde{k}_n + 3 \text{ and } b_{n+1} = b_n + 3 \tag{7.4.11}$$

The construction of trees T_{n+1} and S_{n+1}, coloured leaf sets R_{n+1} and B_{n+1}, the bijection $\varphi_{n+1} \colon X_{n+1} \to Y_{n+1}$, and integers k_{n+1} and b_{n+1} is now complete. In the following, we verify that (†1)–(†14) are indeed satisfied for the $(n+1)^{\text{th}}$ instance.

7.4.7 The inductive step: verification

Claim 7.4.7. *T_{n+1} and S_{n+1} extend T_n and S_n. Moreover, they are rooted trees of maximum degree 3 such that their respective roots are contained in R_{n+1} and B_{n+1}. Hence, (†1)–(†3) are satisfied.*

Proof. Property (†1) follows from (cl.1), i.e. that $\mathrm{cl}(G_n)$ is a leaf-extension of G_n. Thus, T_{n+1} is a leaf extension of $\tilde{T}_n$, which in turn is a leaf extension of T_n, and similar for S_n. This shows (†1).

As noted after the proof of Proposition 7.3.3, taking the closure does not affect the maximum degree, i.e. $\Delta(\mathrm{cl}(G_n)) = \Delta(G_n) = 3$. This shows (†2).

Finally, (7.4.9) implies (†3), as $\mathrm{r}(T_{n+1}) \in R_{n+1}$ and $\mathrm{r}(S_{n+1}) \in B_{n+1}$. □

Claim 7.4.8. *All binary trees appearing as subgraphs of $T_{n+1} \sqcup S_{n+1}$ have height at most b_{n+1}, and every such tree of height b_{n+1} is some copy D_n or $\hat{D}_n$. Hence, T_{n+1} and S_{n+1} satisfy (†4).*

Proof. We first claim that all binary trees appearing as subgraphs of $\tilde{T}_n \sqcup \tilde{S}_n$ which are not contained in D_n or $\hat{D}_n$ have height at most $b_n + 1$. Indeed, note that any binary tree appearing as a subgraph of T_n, $\hat{T}_n(r)$, $\hat{T}_n(v)$, $\hat{S}_n$ or S_n has height at most b_n by the inductive hypothesis. Since the paths we added to the roots of T_n and $\hat{S}_n$ to form $\tilde{T}_n$ were sufficiently long, any binary tree appearing as a subgraph of $\tilde{T}_n$ can only meet one of T_n, $\hat{S}_n$ or D_n. Since the roots of T_n and $\hat{S}_n$ are adjacent to two new vertices in $\tilde{T}_n$, one of degree 1, any such tree meeting T_n or $\hat{S}_n$ must have height at most $b_n + 1$. By Figure 7.6 we see that any binary tree in $\tilde{T}_n$ which meets D_n but whose root lies outside of D_n has height at most $3 \leqslant b_n + 1$. Consider then a binary tree whose root lies inside D_n, but that is not contained in D_n. Again, by Figure 7.6 we see that the root of D_n must lie in one of the bottom three layers of this binary tree. Hence, if the root of this tree lies on the kth level of D_n, then the tree can have height at most $\min\{b_n + 3 - k, k + 2\}$, and hence the tree has height at most $b_n/2 + 2 \leqslant b_n + 1$. Any other binary tree meeting D_n is then contained in D_n. It follows that the only binary tree of height $b_n + 3$ appearing as a subgraph of $\tilde{T}_n$ is D_n, and a similar argument holds for $\tilde{S}_n$ and $\hat{D}_n$.

Recall that T_{n+1} and S_{n+1} are the components of $\mathrm{cl}(\tilde{T}_n \sqcup \tilde{S}_n)$ containing $\tilde{T}_n$ and $\tilde{S}_n$ respectively. If we refer back to Section 7.3 we see that T_{n+1} can be formed from $\tilde{T}_n$ by repeatedly gluing components isomorphic to $\tilde{T}_n(\vec{p_1})$ or $\tilde{S}_n(\vec{p_2})$ to leaves. Consider a binary tree appearing as a subgraph of T_{n+1} which is contained in $\tilde{T}_n$ or one of the copies of $\tilde{T}_n(\vec{p_1})$ or $\tilde{S}_n(\vec{p_2})$. By the previous paragraph, this tree has height at most $b_n + 3$, and if it has height $b_n + 3$ it is a copy D_n or $\hat{D}_n$. Suppose then that there is a binary tree, of height b, whose root is in $\tilde{T}_n$, but is not contained in $\tilde{T}_n$. Such a tree must contain some vertex $\ell \in \tilde{T}_n$ which is adjacent to a vertex not in $\tilde{T}_n$. Hence, ℓ must have been a leaf in $\tilde{T}_n$ at which a copy of $\tilde{T}_n(\vec{p_1})$ or $\tilde{S}_n(\vec{p_2})$ was glued on. However, the roots of each of these components are adjacent to just two vertices, one of degree 1, and hence this leaf ℓ must either be in the bottom, or second to bottom layer of the binary tree. Therefore, $b \leqslant b_n + 2$. A similar argument holds when the root lies in some copy of $\tilde{T}_n(\vec{p_1})$ or $\tilde{S}_n(\vec{p_2})$, and also for S_{n+1}.

Therefore, all binary trees appearing as subgraphs of $T_{n+1} \sqcup S_{n+1}$ have height at most $b_n + 3$, and every such tree is some copy D_n or $\hat{D}_n$. Hence, since $b_{n+1} = b_n + 3$, it follows that $b_{n+1} \geqslant b_n$ and T_{n+1} and S_{n+1} satisfy (†4). □

Claim 7.4.9. *Every maximal bare path in $T_{n+1} \sqcup S_{n+1}$ has length at most k_{n+1}. Hence, T_{n+1} and S_{n+1} satisfy (†5).*

Proof. We first claim that all maximal bare paths in $\tilde{T}_n \sqcup \tilde{S}_n$ have length at most $2\tilde{k}_n + 3$. Firstly, we note that any maximal bare path which is contained in T_n or $\hat{S}_n$ has length at most $k_n \leqslant \tilde{k}_n$

by the induction hypothesis. Also, since the roots of T_n and $\hat{S}_n$ have degree 3 in $\tilde{T}_n$, any maximal bare path is either contained in T_n or $\hat{S}_n$, or does not contain any interior vertices from T_n or $\hat{S}_n$. However, it is clear from the construction that any maximal bare path in $\tilde{T}_n$ that does not contain any interior vertices from T_n or $\hat{S}_n$ has length at most $2\tilde{k}_n + 3$. Similarly, any maximal bare path which is contained in $\hat{T}_n(r)$, $\hat{T}_n(v)$, or S_n has length at most $\tilde{k}_n$ by definition. By the same reasoning as above, any maximal bare path in $\tilde{S}_n$ not contained in $\hat{T}_n(r)$, $\hat{T}_n(v)$, or S_n has length at most $2\tilde{k}_n + 3$.

Again, recall that T_{n+1} can be formed from $\tilde{T}_n$ by repeatedly gluing components isomorphic to $\tilde{T}_n(\vec{p_1})$ or $\tilde{S}_n(\vec{p_2})$ to leaves. Any maximal bare path in T_{n+1} which is contained in $\tilde{T}_n$ or one of the copies of $\tilde{T}_n(\vec{p_1})$ or $\tilde{S}_n(\vec{p_2})$ has length at most $2\tilde{k}_n + 3$ by the previous paragraph. However, since every interior vertex in a maximal bare path has degree two, and the vertices in T_{n+1} at which we, at some point in the construction, stuck on copies of $\tilde{T}_n(\vec{p_1})$ or $\tilde{S}_n(\vec{p_2})$ have degree 3, any maximal bare path in T_{n+1} must be contained in $\tilde{T}_n$ or one of the copies of $\tilde{T}_n(\vec{p_1})$ or $\tilde{S}_n(\vec{p_2})$. Again, a similar argument holds for S_{n+1}. Hence, all maximal bare paths in $T_{n+1} \sqcup S_{n+1}$ have length at most $2\tilde{k}_n + 3$. Therefore, since $k_{n+1} = 2\tilde{k}_n + 3$, it follows that $k_{n+1} \geqslant k_n$ and T_{n+1} and S_{n+1} satisfy (†5). □

Claim 7.4.10. *$\mathrm{Ball}_{T_{n+1}}(T_n, k_n + 1)$ is a bare extension of T_n at $R_n \cup B_n$ to length $k_n + 1$ and does not meet $R_{n+1} \cup B_{n+1}$ and similarly for S_{n+1}. Hence, T_{n+1} and S_{n+1} satisfy* (†6) *and* (†7) *respectively.*

Proof. We will show that T_{n+1} satisfies (†6), the proof that S_{n+1} satisfies (†7) is analogous. By Proposition 7.3.3, the tree T_{n+1} is an $\Big((\tilde{R}_n \cup \tilde{B}_n) \cap V(\tilde{T}_n)\Big)$-extension of $\tilde{T}_n$. Hence T_{n+1} is an

$$\Big(\Big((\tilde{R}_n \cup \tilde{B}_n) \cap V(T_n)\Big) \cup r(T_n)\Big) = \big((R_n \cup B_n) \cap V(T_n)\big)\text{-extension of } T_n. \tag{7.4.12}$$

By looking at the construction of $\mathrm{cl}(G)$ from Section 7.3, we see that T_{n+1} is also an L'-extension of the supertree $T' \supseteq T_n$ formed by gluing a copy of $\tilde{T}_n(\vec{p_1})$ to every leaf in $R_n \cap V(T_n)$ and a copy of $\tilde{S}_n(\vec{p_2})$ to every leaf in $B_n \cap V(T_n)$, where the leaves in L' are the inherited promise leaves from the copies of $\tilde{T}_n(\vec{p_1})$ and $\tilde{S}_n(\vec{p_2})$.

However, we note that every promise leaf in $\tilde{T}_n(\vec{p_1})$ and $\tilde{S}_n(\vec{p_2})$ is at distance at least $\tilde{k}_n + 1$ from the respective root, and so $\mathrm{Ball}_{T_{n+1}}(T_n, \tilde{k}_n) = \mathrm{Ball}_{T'}(T_n, \tilde{k}_n)$. However, $\mathrm{Ball}_{T'}(T_n, \tilde{k}_n)$ can be seen immediately to be a bare extension of T_n at $R_n \cup B_n$ to length $\tilde{k}_n$, and since $\tilde{k}_n \geqslant k_n + 1$ it follows that $\mathrm{Ball}_{T_{n+1}}(T_n, k_n + 1)$ is a bare extension of T_n at $R_n \cup B_n$ to length $k_n + 1$ as claimed.

Finally, we note that $R_{n+1} \cup B_{n+1}$ is the set of promise leaves $\mathrm{cl}(\mathcal{L}_n)$. By the same reasoning as before, $\mathrm{Ball}_{T_{n+1}}(T_n, k_n + 1)$ contains no promise leaf in $\mathrm{cl}(\mathcal{L}_n)$, and so does not meet $R_{n+1} \cup B_{n+1}$ as claimed. □

Claim 7.4.11. *Let U_{n+1} be a bare extension of $\mathrm{cl}(G_n) = T_{n+1} \sqcup S_{n+1}$ at $R_{n+1} \cup B_{n+1}$ to any length. Then any embedding of T_{n+1} or S_{n+1} into U_{n+1} fixes the respective root. Hence, T_{n+1} and S_{n+1} satisfy* (†8).

Proof. Recall that the promise closure was constructed by recursively adding copies of rooted trees C_i and identifying their roots with promise leaves. For the promise structure $\mathcal{P} = \Big(G_n, \vec{P}, \mathcal{L}\Big)$ on G_n we have $C_1 = \tilde{T}_n(\vec{p_1})$ and $C_2 = \tilde{S}_n(\vec{p_2})$.

Note that by (†5), the image of any embedding $T_n \hookrightarrow U_{n+1}$ cannot contain a bare path of length $k_n + 1$. Also, by construction, every copy of T_n, S_n, $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$ in T_{n+1} has the property that its $(k_n + 1)$-ball in T_{n+1} is a bare extension to length $k_n + 1$ of this copy. Hence, if the root of T_n embeds into some copy of T_n, S_n, $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$, then the whole tree embeds into a bare extension of this copy. The same is true for S_n.

By (†8), there are no embeddings of T_n into a bare extension of S_n, or of S_n into a bare extension of T_n. Moreover, since both $\hat{T}_n(r)$ and $\hat{T}_n(\hat{v})$ are subtrees of T_n, there is no embedding of T_n or S_n into bare extensions of them by (†8) and (†9).

Thus, only the following embeddings are possible:

- T_n embeds into a bare extension of a copy of T_n, or S_n embeds into a bare extension of a copy of S_n. In both cases, the root must be preserved, as otherwise we contradict (†9) or (†10).

Let $f\colon T_{n+1} \hookrightarrow U_{n+1}$ be an embedding. By Claim 7.4.8, U_{n+1} contains no binary trees of height $b_n + 3$ apart from D_n, $\hat{D}_n$, and the copies of those two trees that were created by adding copies of C_1 and C_2. Consequently f maps D_n to one of these copies, mapping the root to the root. The neighbours of $\mathrm{r}(T_{n+1})$ and g must map to vertices of degree 3 at distance two and three from the image of the root of D_n respectively, which forces $f(\mathrm{r}(T_{n+1})) \in R_{n+1}$. If $f(\mathrm{r}(T_{n+1})) = \mathrm{r}(T_{n+1})$ then we are done.

Otherwise there are two possibilities for $f(\mathrm{r}(T_{n+1}))$. If $f(\mathrm{r}(T_{n+1}))$ is contained in a copy of C_1, then $\mathrm{r}(T_n)$ maps to a promise leaf other than the root in a copy of T_n, S_n, $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$. If $f(\mathrm{r}(T_{n+1})) = y$ or $f(\mathrm{r}(T_{n+1}))$ is contained in a copy of C_2, then $\mathrm{r}(T_n)$ maps to a copy of $\mathrm{r}\left(\hat{T}_n(r)\right)$ or some vertex of $\hat{T}_n(\hat{v})$. In both cases the root of T_n does not map to the root of a copy of T_n, which is impossible by the first bullet point.

Finally, let $f\colon S_{n+1} \hookrightarrow U_{n+1}$ be an embedding. By the same arguments as above $f(\mathrm{r}(S_{n+1})) \in B_{n+1}$. If f fixes $\mathrm{r}(S_{n+1})$, we are done.

Otherwise we have again two cases. If $f(\mathrm{r}(S_{n+1})) = g$, or $f(\mathrm{r}(S_{n+1}))$ is contained in a copy of C_1, then v_{k+1} (the neighbour of $\hat{v}$ on the long path) would have to map to a vertex of degree 2, giving an immediate contradiction. If $f(\mathrm{r}(S_{n+1}))$ is contained in a copy of C_2, then $\mathrm{r}(S_n)$ maps to a promise leaf other than the root in a copy of T_n, S_n, $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$ which is also impossible by the observations in the bullet points. □

Claim 7.4.12. *Let U_{n+1} be as in Claim 7.4.11. Then there is no embedding of T_{n+1} or S_{n+1} into U_{n+1} whose image contains vertices outside of $\mathrm{cl}(G_n)$, i.e. vertices that have been added to form the bare extension.*

Since a root-preserving embedding of a locally finite tree into itself must be an automorphism, this together with the previous claim implies (†9)and (†10).

Proof. We prove this claim for T_{n+1}, the proof for S_{n+1} is similar. Assume for a contradiction that there is a vertex w of T_{n+1} and an embedding $f\colon T_{n+1} \hookrightarrow U_{n+1}$ such that $f(w) \notin \mathrm{cl}(G_n)$. By definition of bare extension, removing $f(w)$ from U_{n+1} splits the component of $f(w)$ into at most two components, one of which is a path.

Note first that w does not lie in a copy of D_n or $\hat{D}_n$, because these must map to binary trees of the same height by Claim 7.4.8. Furthermore, all vertices in $R_{n+1} \cup B_{n+1}$ have a neighbour of degree 3 whose neighbours all have degree $\geqslant 2$, thus $w \notin R_{n+1} \cup B_{n+1}$. Finally, only one component of $T_{n+1} - w$ can contain vertices of degree 3. Consequently, w must lie in a copy C of T_n, S_n, $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$.

All maximal bare paths in the image $f(C)$ have length at most $k = \tilde{k}_n$, so $f(C)$ cannot intersect any copies of T_n, S_n, $\hat{T}_n(r)$, or $(\hat{T}_n(\hat{v}) + v_{k+1})$. Let r be the root of C (where $r = \hat{v}$ in the last case). Now $f(r)$ must have the following properties: it is a vertex of degree 3, and the root of a nearest binary tree of height b_{n+1} not containing $f(r)$ lies at distance d from $f(r)$, where $5 \leqslant d \leqslant 2k + 4$.

But the only vertices with these properties are contained in copies of T_n, $\hat{S}_n$, $\hat{T}_n(r)$, or $(\hat{T}_n(\hat{v}) + v_{k+1})$. This contradicts the fact that $f(C)$ does not intersect any of these copies. □

Claim 7.4.13. *The function φ_{n+1} is a well-defined bijection extending φ_n, such that its domain and range do not intersect $R_{n+1} \cup B_{n+1}$. Hence, property (†13) holds for $\varphi_{n+1}\colon X_{n+1} \to Y_{n+1}$.*

Proof. By the choice of x in (7.4.1) and the definition of $\varphi_{n+1}\colon X_{n+1} \to Y_{n+1}$ in (7.4.10), the first three items of property (†13) hold.

Since v does not lie in $R_n \cup B_n$ by (7.4.2), it follows by our construction of the promise structure $\mathcal{P} = \left(G_n, \vec{P}, \mathcal{L}\right)$ in (7.4.5) and (7.4.6) that neither v nor $\hat{v} = \varphi_{n+1}(v)$ appear as promise leaves in $\mathcal{L}$. Furthermore, by the induction hypothesis, $(X_n \cup Y_n) \cap (R_n \cup B_n) = \emptyset$, so no vertex in $(X_n \cup Y_n)$ appears as a promise leaf in $\mathcal{L}$ either. Thus, in formulas,

$$(X_{n+1} \cup Y_{n+1}) \cap \bigcup_{L \in \mathcal{L}} L = \emptyset. \tag{7.4.13}$$

In particular, since

$$(R_{n+1} \cup B_{n+1}) \cap G_n = (\mathrm{cl}(L_3) \cup \mathrm{cl}(L_4)) \cap G_n = L_3 \cup L_4,$$

and $X_{n+1} \cup Y_{n+1} \subset G_n$, we get $(X_{n+1} \cup Y_{n+1}) \cap (R_{n+1} \cup B_{n+1}) = \emptyset$. Thus, also the last item of (†13) is verified. □

Claim 7.4.14. *There is a family of isomorphisms $\mathcal{H}_{n+1} = \{h_{n+1,x}\colon x \in X_{n+1}\}$ witnessing that $T_{n+1} - x$ and $S_{n+1} - \varphi_{n+1}(x)$ are isomorphic for all $x \in X_{n+1}$, such that $h_{n+1,x}$ extends $h_{n,x}$ for all $x \in X_n$. Hence, property (†14) holds.*

Proof. There are four things to be verified for this claim. Firstly, we need an isomorphism $h_{n+1,v}$ witnessing that $T_{n+1} - v$ and $S_{n+1} - \hat{v}$ are isomorphic. Secondly, we need to *extend* all previous isomorphisms $h_{n,x}$ between $T_n - x$ and $S_n - \varphi_n(x)$ to $T_{n+1} - x$ and $S_{n+1} - \varphi_n(x)$. This will take care of the first item of (†14). To also comply with the remaining two items, we need to make sure that each isomorphism in

$$\mathcal{H}_{n+1} = \{h_{n+1,x}\colon x \in X_{n+1}\}$$

maps leaves in $R_{n+1} \cap V(T_{n+1})$ bijectively to leaves in $R_{n+1} \cap V(S_{n+1})$, and similarly for B_{n+1}.

To find the first isomorphism, note that by construction of the promise structure $\mathcal{P} = \left(G_n, \vec{P}, \mathcal{L}\right)$ on G_n in (7.4.5), and properties (cl.1) and (cl.3) of the promise closure, the trees T_{n+1} and S_{n+1} are obtained from $\tilde{T}_n$ and $\tilde{S}_n$ by attaching at every leaf $r \in \tilde{R}_n$ a copy of the rooted tree $\mathrm{cl}(G_n)(\vec{p}_1)$, and by attaching at every leaf $b \in \tilde{B}_n$ a copy of the rooted tree $\mathrm{cl}(G_n)(\vec{p}_2)$.

By (7.4.13), neither v nor $\varphi_{n+1}(v)$ are mentioned in $\mathcal{L}$. As observed in (7.4.7), there is a $\vec{P}$-respecting isomorphism

$$h\colon \tilde{T}_n - v \to \tilde{S}_n - \varphi_{n+1}(v).$$

In other words, h maps promise leaves in $L_i \cap V(\tilde{T}_n)$ bijectively to the promise leaves in $L_i \cap V(\tilde{S}_n)$ for all $i = 1, 2, 3, 4$. Our plan is to extend h to an isomorphism between $T_{n+1} - v$ and $S_{n+1} - \varphi_n(v)$ by mapping the corresponding copies of $\mathrm{cl}(G_n)(\vec{p}_1)$ and $\mathrm{cl}(G_n)(\vec{p}_2)$ attached to the various red and blue leaves to each other.

Formally, by (cl.3) there is for each $\ell \in \left(\tilde{R}_n \cup \tilde{B}_n\right) \cap V(T)$ a $\mathrm{cl}(\vec{P})$-respecting isomorphism of rooted trees

$$\mathrm{cl}(G_n)(\vec{q}_\ell) \cong \mathrm{cl}(G_n)(\vec{q}_{h(\ell)}).$$

Therefore, by combining the isomorphism h between $\tilde{T}_n - v$ and $\tilde{S}_n - \varphi_{n+1}(v)$ with these isomorphisms between each $\mathrm{cl}(G_n)(\vec{q}_\ell)$ and $\mathrm{cl}(G_n)(\vec{q}_{h(\ell)})$ we get a $\mathrm{cl}(\vec{P})$-respecting isomorphism

$$h_{n+1,v}\colon T_{n+1} - v \to S_{n+1} - \varphi_{n+1}(v).$$

And since R_{n+1} and B_{n+1} have been defined in (7.4.9) to be the promise leaf sets of $\mathrm{cl}(\mathcal{P})$, by definition of $\mathrm{cl}(\vec{P})$-respecting (Def. 7.3.4), the image of $R_{n+1} \cap V(T_{n+1})$ under $h_{n+1,v}$ is $R_{n+1} \cap V(S_{n+1})$, and similarly for B_{n+1}.

It remains to extend the old isomorphisms in $\mathcal{H}_n$. As argued in (7.4.12), both trees T_{n+1} and S_{n+1} are leaf extensions of T_n and S_n at $R_n \cup B_n$ respectively. By property (cl.3), these leaf extensions are obtained by attaching at every leaf $r \in R_n$ a copy of the rooted tree $\mathrm{cl}(G_n)(\vec{p}_1)$, and similarly by attaching at every leaf $b \in B_n$ a copy of the rooted tree $\mathrm{cl}(G_n)(\vec{p}_2)$.

By induction assumption (†14), for each $x \in X_n$ the isomorphism

$$h_{n,x} \colon T_n - x \to S_n - \varphi_n(x)$$

maps the red leaves of T_n bijectively to the red leaves of S_n, and the blue leaves of T_n bijectively to the blue leaves of S_n. Thus, by property (cl.3), there are $\mathrm{cl}(\vec{P})$-respecting isomorphisms of rooted trees

$$\mathrm{cl}(G_n)(\vec{q}_\ell) \cong \mathrm{cl}(G_n)(\vec{q}_{h_{n,x}(\ell)})$$

for all $\ell \in (R_n \cup B_n) \cap V(T_n)$. By combining the isomorphism $h_{n,x}$ between $T_n - x$ and $S_n - \varphi_n(x)$ with these isomorphisms between each $\mathrm{cl}(G_n)(\vec{q}_\ell)$ and $\mathrm{cl}(G_n)(\vec{q}_{h_{n,x}(l)})$, we obtain a $\mathrm{cl}(\vec{P})$-respecting extension

$$h_{n+1,x} \colon T_{n+1} - x \to S_{n+1} - \varphi_n(x).$$

As before, by definition of $\mathrm{cl}(\vec{P})$-respecting, the image of $R_{n+1} \cap V(T_{n+1})$ under $h_{n+1,x}$ is $R_{n+1} \cap V(S_{n+1})$, and similarly for B_{n+1}.

Finally, by construction we have $h_{n+1,x} \restriction (T_n - x) = h_{n,x}$ for all $x \in X_n$ as desired. The proof is complete. □

7.5 The trees are also edge-hypomorphic

In this final section, we briefly indicate why the trees T and S yielded by our strategy above are automatically edge-hypomorphic: we claim the correspondence

$$\psi \colon e(x) \mapsto e(\varphi(x))$$

as introduced in (7.4.3) and (7.4.4) is an edge-hypomorphism between T and S. For this, we need to verify that

(a) ψ is a bijection between $E(T)$ and $E(S)$, and that

(b) the maps $h_x \cup \{\langle x, \varphi(x)\rangle\} \colon G - e(x) \to H - e(\varphi(x))$ are isomorphisms.

Regarding (b), observe that the map h as defined in (7.4.7) yields, by construction, also a $\vec{P}$-respecting isomorphism

$$h \cup \{(v, \hat{v})\} \colon \tilde{T}_n - e(v) \to \tilde{S}_n - e(\hat{v}),$$

and from there, the arguments are entirely the same as in the previous section.

For (a), we use the canonical bijection between the edge set of a rooted tree, and its vertices other than the root; namely the bijection mapping every such vertex to the first edge on its unique path to the root. Thus, given the enumeration of $V(T_n)$ and $V(S_n)$ in (†11), we obtain corresponding enumerations of $E(T_n)$ and $E(S_n)$, and since the rooted trees T_n and S_n are order-preserving subtrees of the rooted trees T_{n+1} and S_{n+1} (cf. Figure 7.6), it follows that also our enumerations of $E(T_n)$ and $E(S_n)$ extend the enumerations of $E(T_{n-1})$ and $E(S_{n-1})$ respectively. But now it follows from (†13) and the definition of ψ that by step $2(n+1)$ we have dealt with the first n edges in our enumerations of $E(T)$ and $E(S)$ respectively.

Chapter 8

Non-reconstructible locally finite graphs

8.1 Introduction

Two graphs G and H are *hypomorphic* if there exists a bijection φ between their vertex sets such that the induced subgraphs $G - v$ and $H - \varphi(v)$ are isomorphic for each vertex v of G. We say that a graph G is *reconstructible* if $H \cong G$ for every H hypomorphic to G. The *Reconstruction Conjecture*, a famous unsolved problem attributed to Kelly and Ulam, suggests that every finite graph with at least three vertices is reconstructible.

For an overview of results towards the Reconstruction Conjecture for finite graphs see the survey of Bondy and Hemminger [23]. The corresponding reconstruction problem for infinite graphs is false: the countable regular tree T_∞, and two disjoint copies of it (written as $T_\infty \cup T_\infty$) are easily seen to be non-homeomorphic reconstructions of each other. This example, however, contains vertices of infinite degree. Regarding locally finite graphs, Harary, Schwenk and Scott [76] showed that there exists a non-reconstructible locally finite forest. However, they conjectured that the Reconstruction Conjecture should hold for locally finite trees. This conjecture has been verified for locally finite trees with at most countably many ends in a series of paper [10, 22, 125]. However, very recently, the present authors have constructed a counterexample to the conjecture of Harary, Schwenk and Scott.

Theorem 8.1.1 (Bowler, Erde, Heinig, Lehner, Pitz [29]). *There exists a non-recon-structible tree of maximum degree three.*

The Reconstruction Conjecture has also been considered for general locally finite graphs. Nash-Williams [102] showed that if $p \geqslant 3$ is an integer, then any locally finite graph with exactly p ends is reconstructible; and in [104] he showed the same is true for $p = 2$. The case $p = 2$ is significantly more difficult. Broadly speaking this is because every graph with $p \geqslant 3$ ends has some identifiable finite 'centre', from which the ends can be thought of as branching out. A two-ended graph however can be structured like a double ray, without an identifiable 'centre'.

The case of 1-ended graphs is even harder, and the following problems from a survey of Nash-Williams [103], which would generalise the corresponding results established for trees, have remained open.

Problem 7 (Nash-Williams). *Is every locally finite graph with exactly one end reconstructible?*

Problem 8 (Nash-Williams). *Is every locally finite graph with countably many ends reconstructible?*

In this paper, we extend our methods from [29] to construct examples showing that both of Nash-Williams' questions have negative answers. Our examples will not only be locally finite, but in fact have bounded degree.

Theorem 8.1.2. *There is a connected one-ended non-reconstructible graph with bounded maximum degree.*

Theorem 8.1.3. *There is a connected countably-ended non-reconstructible graph with bounded maximum degree.*

Since every locally finite graph has either finitely many, countably many or continuum many ends, Theorems 8.1.1, 8.1.2 and 8.1.3 together with the results of Nash-Williams provide a complete picture about what can be said about number of ends versus reconstruction:

- A locally finite tree with at most countably many ends is reconstructible; but there are non-reconstructible locally finite trees with continuum many ends.
- A locally finite graph with at least two, but a finite number of ends is reconstructible; but there are non-reconstructible locally finite graphs with one, countably many, and continuum many ends respectively.

This paper is organised as follows: In the next section we give a short, high-level overview of our constructions which answer Nash-Williams' problems. In Sections 8.3 and 8.4, we develop the technical tools necessary for our construction, and in Sections 8.5 and 8.6, we prove Theorems 8.1.2 and 8.1.3.

For standard graph theoretical concepts we follow the notation in [43].

8.2 Sketch of the construction

In this section we sketch the main ideas of the construction in three steps. First, we quickly recall our construction of two hypomorphic, non-isomorphic locally finite trees from [29]. We will then outline how to adapt the construction to obtain a one-ended-, and a countably-ended counterexample respectively.

8.2.1 The tree case

This section contains a very brief summary of the much more detailed sketch from [29]. The strategy is to build trees T and S recursively, where at each step of the construction we ensure for some new vertex v already chosen for T that there is a corresponding vertex w of S with $T - v \cong S - w$, or vice versa. This will ensure that by the end of the construction, the trees we have built are hypomorphic.

More precisely, at step n we will construct subtrees T_n and S_n of our eventual trees, where some of the leaves of these subtrees have been coloured in two colours, say red and blue. We will only further extend the trees from these coloured leaves, and we will extend from leaves of the same colour in the same way. We also make sure that earlier partial isomorphisms between $T_n - v_i \cong S_n - w_i$ preserve leaf colours. Together, these requirements guarantee that earlier partial isomorphisms always extend to the next step.

The T_n will be nested, and we will take T to be the union of all of them; similarly the S_n will be nested and we take S to be the union of all of them. To ensure that T and S do not end up being isomorphic, we first ensure, for each n, that there is no isomorphism from T_n to S_n. Our second requirement is that T or S beyond any coloured leaf of T_n or S_n begins with a long

Figure 8.1: A first approximation of T_{n+1} on the left, and S_{n+1} on the right. All dotted lines are long non-branching paths.

non-branching path, longer than any such path appearing in T_n or S_n. Together, this implies that T and S are not isomorphic.

Algorithm Stage One: Suppose now that we have already constructed T_n and S_n and wish to construct T_{n+1} and S_{n+1}. Suppose further that we are given a vertex v of T_n for which we wish to find a partner w in S_{n+1} so that $T - v$ and $S - w$ are isomorphic. We begin by building a tree $\hat{T}_n \not\cong T_n$ which has some vertex w such that $T_n - v \cong \hat{T}_n - w$. This can be done by taking the components of $T_n - v$ and arranging them suitably around the new vertex w.

We will take S_{n+1} to include S_n and $\hat{T}_n$, with the copies of red and blue leaves in $\hat{T}_n$ also coloured red and blue respectively. As indicated on the right in Figure 8.1, we add long non-branching paths to some blue leaf b of S_n and to some red leaf r of $\hat{T}_n$ and join these paths at their other endpoints by some edge e_n. We also join two new leaves y and g to the endvertices of e_n. We colour the leaf y yellow and the leaf g green. To ensure that $T_{n+1} - v \cong S_{n+1} - w$, we take T_{n+1} to include T_n together with a copy $\hat{S}_n$ of S_n, with its leaves coloured appropriately, and joined up in the same way, as indicated on the left in Figure 8.1. Note that, whilst $\hat{S}_n$ and S_n are isomorphic as graphs, we make a distinction as we want to lift the partial isomorphisms between $T_n - v_i \cong S_n - w_i$ to these new graphs, and our notation aims to emphasize the natural inclusions $T_n \subset T_{n+1}$ and $S_n \subset S_{n+1}$.

Algorithm Stage Two: We now have committed ourselves to two targets which are seemingly irreconcilable: first, we promised to extend in the same way at each red or blue leaf of T_n and S_n, but we also need that $T_{n+1} - v \cong S_{n+1} - w$. The solution is to copy the same subgraph appearing beyond r in Fig. 8.1, including its coloured leaves, onto all the other red leaves of S_n and T_n. Similarly we copy the subgraph appearing beyond the blue leaf b of S_n onto all other blue leaves of S_n and T_n. In doing so, we create new red and blue leaves, and we will keep adding, step by step, further copies of the graphs appearing beyond r and b in Fig. 8.1 respectively onto all red and blue leaves of everything we have constructed so far.

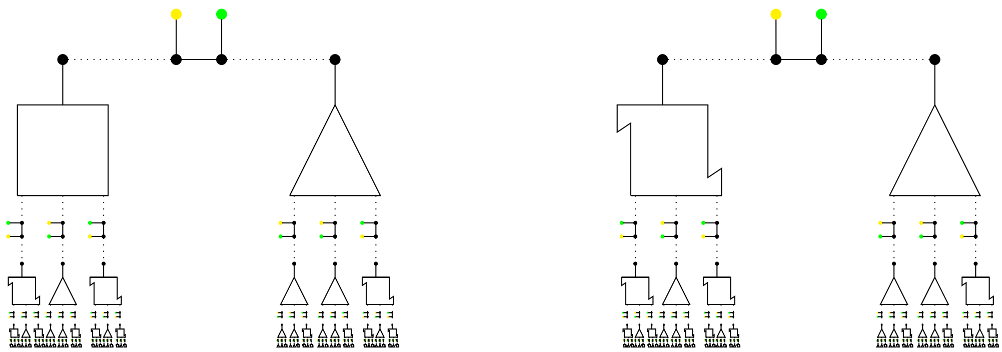

Figure 8.2: A sketch of T_{n+1} and S_{n+1} after countably many steps.

After countably many steps we have dealt with all red and blue leaves, and it can be checked that both our targets are achieved. We take these new trees to be S_{n+1} and T_{n+1}. They are non-

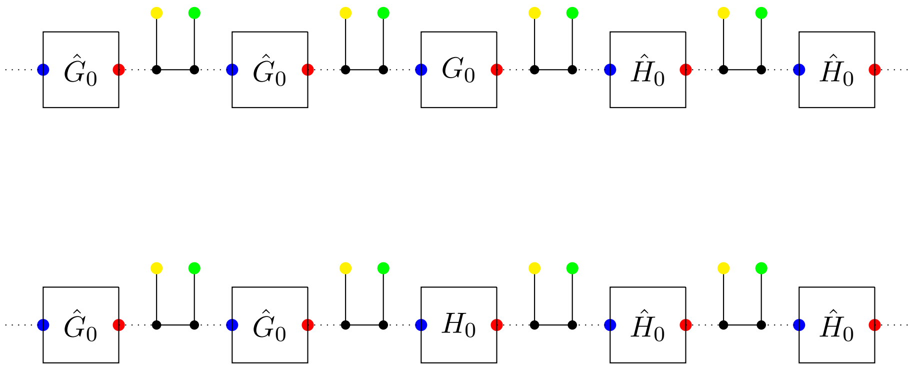

Figure 8.3: A sketch of G_1 (above) and H_1 (below).

isomorphic, as after removing all long non-branching paths, T_{n+1} contains T_n as a component, whereas S_{n+1} does not.

8.2.2 The one-ended case

To construct a one-ended non-reconstructible graph, we initially follow the same strategy as in the tree case and build locally finite graphs G_n and H_n and some partial hypomorphisms between them. Simultaneously, however, we will also build one-ended locally finite graphs of a grid-like form $F_n \times \mathbb{N}$ (the Cartesian product of a locally finite tree F_n with a ray) which share certain symmetries with G_n and H_n. These will allow us to glue $F_n \times \mathbb{N}$ onto both G_n and H_n, in order to make them one-ended, without spoiling the partial hypomorphisms. Let us illustrate this idea by explicitly describing the first few steps of the construction.

We start with two non-isomorphic graphs G_0 and H_0, such that G_0 and H_0 each have exactly one red and one blue leaf. After stage one of our algorithm, our approximations to G_1 and H_1 as in Figure 8.1 contain, in each of G_0, $\hat{H}_0$, $\hat{G}_0$ and H_0, one coloured leaf. In stage two, we add copies of these graphs recursively. It follows that the resulting graphs G_1' and H_1' have the global structure of a double ray, along which parts corresponding to copies of G_0, $\hat{H}_0$, $\hat{G}_0$ and H_0 appear in a repeating pattern. Crucially, however, each graph G_1' and H_1' has infinitely many yellow and green leaves, which appear in an alternating pattern extending to infinity in both directions along the double ray.

Consider the minor F_1 of G_1' obtained by collapsing every subgraph corresponding to G_0, $\hat{H}_0$, $\hat{G}_0$ and H_0 to a single point. Write $\psi_G \colon G_1' \to F_1$ for the quotient map. Then F_1 is a double ray with alternating coloured leaves hanging off it. Note that we could have started with H_1' and obtained the same F_1. In other words, F_1 approximates the global structures of both G_1' and H_1'. Consider the one-ended grid-like graph $F_1 \times \mathbb{N}$, where we let $F_1 \times \{0\}$ inherit the colours from F_1. We now form G_1 and H_1 by gluing $F_1 \times \mathbb{N}$ onto G_1', by identifying corresponding coloured vertices y and $\psi_G(y)$, and similarly for H_2'. [1] Since the coloured leaves contained both ends of our graphs in their closure, the graphs G_1 and H_1 are now one-ended.

It remains to check that our partial isomorphism $h_1 \colon G_1' - v_1 \to H_1' - w_1$ guaranteed by step two can be extended to $G_1 - v_1 \to H_1 - w_1$. This can be done essentially because of the following property: let us write $\mathcal{L}(\cdot)$ for the set of coloured leaves. It can be checked that there is an automorphism $\pi_1 \colon F_1 \to F_1$ such that the diagram

[1] For technical reasons, in the actual construction we identify $\psi_G(y)$ with the corresponding base vertex of the leaf y in G_1'. In this way the coloured leaves of G_1' remain leaves, and we can continue our recursive construction.

$$\begin{tikzcd} \mathcal{L}(G_1') \arrow[r, "h_1 \restriction \mathcal{L}(G_1)"] \arrow[d, "\psi_G"'] & \mathcal{L}(H_1') \arrow[d, "\psi_H"] \\ \mathcal{L}(F_1) \arrow[r, "\pi_1 \restriction \mathcal{L}(F_1)"] & \mathcal{L}(F_1) \end{tikzcd}$$

is colour-preserving and commutes. Hence, $\pi_1 \times \mathrm{id}$ is an automorphism of $F_1 \times \mathbb{N}$ which is compatible with our gluing procedure, so it can be combined with h_1 to give us the desired isomorphism.

We are now ready to describe the general step. Instead of describing F_n as a minor of G_n, which no longer works naïvely at later steps, we will directly build F_n by recursion, so that it satisfies the properties of the above diagram.

Suppose at step n we have constructed locally finite graphs G_n and H_n, and also a locally finite tree F_n where some leaves are coloured in one of two colours. Furthermore, suppose we have a family of isomorphisms

$$\mathcal{H}_n = \{h_x \colon G_n - x \to H_n - \varphi(x) \colon x \in X_n\},$$

for some subset $X_n \subset V(G_n)$, a family of isomorphisms $\Pi_n = \{\pi_x \colon F_n \to F_n \colon x \in X_n\}$, and colour-preserving bijections $\psi_{G_n} \colon \mathcal{L}(G_n) \to \mathcal{L}(F_n)$ and $\psi_{H_n} \colon \mathcal{L}(H_n) \to \mathcal{L}(F_n)$ such that the corresponding commutative diagram from above holds for each x. We construct G'_{n+1} and H'_{n+1} according to stages one and two of the previous algorithm. As before our isomorphisms h_x will lift to isomorphisms between $G'_{n+1} - x$ and $H'_{n+1} - \varphi(x)$.

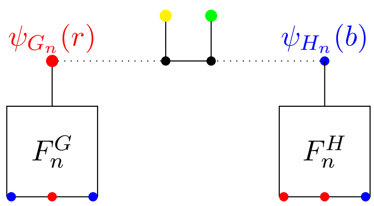

Figure 8.4: The auxiliary graph $\tilde{F}_n$.

Algorithm Stage Three. As indicated in Figure 8.4, we take two copies F_n^G and G_n^H of F_n, and glue them together mimicking stage one of the algorithm, i.e. connect $\psi_{G_n}(r)$ in F_n^G by a path of length three to $\psi_{H_n}(b)$ in F_n^H, and attach two new leaves coloured yellow and green in the middle of the path. Call the resulting graph $\tilde{F}_n$. We then apply stage two of the algorithm to this graph, gluing again and again onto every blue vertex a copy of the graph of $\tilde{F}_n$ behind $\psi_{H_n}(b)$, and similarly for every red leaf, to obtain a tree F_{n+1}. Since this procedure is, in structural terms, so similar to the construction of G'_{n+1} and H'_{n+1}, it can be shown that we do obtain a colour-preserving commuting diagram of the form

$$\begin{tikzcd} \mathcal{L}(G'_{n+1}) \arrow[r, "h_x \restriction \mathcal{L}(G'_{n+1})"] \arrow[d, "\psi_{G_{n+1}}"'] & \mathcal{L}(H'_{n+1}) \arrow[d, "\psi_{H_{n+1}}"] \\ \mathcal{L}(F_{n+1}) \arrow[r, "\pi_x \restriction \mathcal{L}(F_{n+1})"] & \mathcal{L}(F_{n+1}) \end{tikzcd}$$

As before, this means that we can indeed glue together G'_{n+1} and $F_{n+1} \times \mathbb{N}$, and H'_{n+1} and $F_{n+1} \times \mathbb{N}$ to obtain one-ended graphs G_{n+1} and H_{n+1} as desired.

At the end of our construction, after countably many steps, we have built two graphs G and H which are hypomorphic, and for the same reasons as in the tree case the two graphs will not be isomorphic. Further, since all G_n and H_n are one-ended, so will be G and H.

8.2.3 The countably-ended case

In order to produce hypomorphic graphs with countably many ends we follow the same procedure as for the one-ended case, except that we start with one-ended (non-isomorphic) graphs G_0 and H_0.

After the first and second stage of our algorithm, the resulting graphs G_1' and H_1' will again consist of infinitely many copies of G_0 and H_0 glued together along a double ray. After gluing $F_1 \times \mathbb{N}$ to these graphs as before, we obtain graphs with one thick end, with many coloured leaves tending to that end, as well as infinitely many thin ends, coming from the copies of G_0 and H_0, each of which contained a ray. These thin ends will eventually be rays, and so have no coloured leaves tending towards them. This guarantees that in the next step, when we glue $F_2 \times \mathbb{N}$ onto G_2' and H_2', the thin ends will not be affected, and that all the other ends in the graph will be amalgamated into one thick end.

Then, in each stage of the construction, the graphs G_n and H_n will have exactly one thick end, again with many coloured leaves tending towards it, and infinitely many thin ends each of which is eventually a ray. This property lifts to the graphs G and H constructed in the limit: they will have one thick end and infinitely many ends which are eventually rays. However, since G and H are countable, there can only be countably many of these rays. Hence the two graphs G and H have countably many ends in total, and as before they will be hypomorphic but not isomorphic.

8.3 Closure with respect to promises

A *bridge* in a graph G is an edge $e = \{x, y\}$ such that x and y lie in different components of $G - e$. Given a directed bridge $\vec{e} = \vec{xy}$ in some graph $G = (V, E)$, we denote by $G(\vec{e})$ the unique component of $G - e$ containing the vertex y. We think of $G(\vec{e})$ as a rooted graph with root y.

Definition 8.3.1 (Promise structure)**.** *A* promise structure $\mathcal{P} = \left(G, \vec{P}, \mathcal{L}\right)$ *is a triple consisting of:*

- *a graph G,*
- *$\vec{P} = \{\vec{p_i} : i \in I\}$ a set of directed bridges $\vec{P} \subset \vec{E}(G)$, and*
- *$\mathcal{L} = \{L_i : i \in I\}$ a set of pairwise disjoint sets of leaves of G.*

We insist further that, if the component $G(\vec{p_i})$ consists of a single leaf $c \in L_j$, then $i = j$.

Often, when the context is clear, we will not make a distinction between $\mathcal{L}$ and the set $\bigcup_i L_i$, for notational convenience.

We call an edge $\vec{p_i} \in \vec{P}$ a *promise edge*, and leaves $\ell \in L_i$ *promise leaves.* A promise edge $\vec{p_i} \in \vec{P}$ is called a *placeholder-promise* if the component $G(\vec{p_i})$ consists of a single leaf $c \in L_i$, which we call a *placeholder-leaf.* We write

$$\mathcal{L}_p = \{L_i : \vec{p_i} \text{ a placeholder-promise}\} \text{ and } \mathcal{L}_q = \mathcal{L} \setminus \mathcal{L}_p.$$

Given a leaf ℓ in G, there is a unique edge $q_\ell \in E(G)$ incident with ℓ, and this edge has a natural orientation $\vec{q_\ell}$ towards ℓ. Informally, we think of $\ell \in L_i$ as the 'promise' that if we extend G to a graph $H \supset G$, we will do so in such a way that $H(\vec{q_\ell}) \cong H(\vec{p_i})$.

Definition 8.3.2 (Leaf extension)**.** *Given an inclusion $H \supseteq G$ of graphs and a set L of leaves of G, H is called a* leaf extension*, or more specifically an* L-extension, of G*, if:*

- *every component of H contains precisely one component of G, and*
- *every component of $H - G$ is adjacent to a unique vertex l of G, and we have $l \in L$.*

In [29], given a promise structure $\mathcal{P} = \left(G, \vec{P}, \mathcal{L}\right)$, it is shown how to construct a graph $\mathrm{cl}(G) \supset G$ which has the following properties.

Proposition 8.3.3 (Closure w.r.t a promise structure, cf. [29, Proposition 3.3])**.** *Let G be a graph and let $\left(G, \vec{P}, \mathcal{L}\right)$ be a promise structure. Then there is a graph* $\mathrm{cl}(G)$*, called the* closure of G with respect to $\mathcal{P}$*, such that:*

(cl.1) $\mathrm{cl}(G)$ *is an $\mathcal{L}_q$-extension of G,*

(cl.2) *for every $\vec{p}_i \in \vec{P}$ and all $\ell \in L_i$,*

$$\mathrm{cl}(G)(\vec{p}_i) \cong \mathrm{cl}(G)(\vec{q}_\ell)$$

are isomorphic as rooted graphs.

Since the existence of $\mathrm{cl}(G)$ is crucial to our proof, we briefly remind the reader how to construct such a graph. As a first approximation, in order to try to achieve ((cl.2)), we glue a copy of the component $G(\vec{p}_i)$ onto each leaf $\ell \in L_i$, for each $i \in I$. We call this the *1-step extension* $G^{(1)}$ of G. If there were no promise leaves in the component $G(\vec{p}_i)$, then the promises in L_i would be satisfied. However, if there are, then we have grown $G(\vec{p}_i)$ by adding copies of various $G(\vec{p}_j)$s behind promise leaves appearing in $G(\vec{p}_i)$.

However, remembering all promise leaves inside the newly added copies of $G(\vec{p}_i)$ we glued behind each $\ell \in L_i$, we continue this process indefinitely, growing the graph one step at a time by gluing copies of (the original) $G(\vec{p}_i)$ to promise leaves ℓ' which have appeared most recently as copies of $\ell \in L_i$. After a countable number of steps the resulting graph $\mathrm{cl}(G)$ satisfies Proposition 8.3.3. We note also that the maximum degree of $\mathrm{cl}(G)$ equals that of G.

Definition 8.3.4 (Promise-respecting map)**.** *Let G be a graph, $\mathcal{P} = \left(G, \vec{P}, \mathcal{L}\right)$ be a promise structure on G, and let T_1 and T_2 be two components of G.*

Given $x \in T_1$ and $y \in T_2$, a bijection $\varphi \colon T_1 - x \to T_2 - y$ is $\vec{P}$-respecting *(with respect to $\mathcal{P}$) if the image of $L_i \cap T_1$ under φ is $L_i \cap T_2$ for all i.*

We can think of $\mathcal{P}$ as defining a $|\vec{P}|$-colouring on some sets of leaves. Then a mapping is $\vec{P}$-respecting if it preserves leaf colours.

Suppose that $\vec{p}_i$ is a placeholder promise, and $G = H^{(0)} \subseteq H^{(1)} \subseteq \cdots$ is the sequence of 1-step extensions whose direct limit is $\mathrm{cl}(G)$. Then, if we denote by $L_i^{(n)}$ the set of promise leaves associated with $\vec{p}_i$ in $H^{(n)}$, it follows that $L_i^{(n)} \supseteq L_i^{(n-1)}$ since $G(\vec{p}_i)$ is just a single vertex $c_i \in L_i$. For every placeholder promise $\vec{p}_i \in \vec{P}$, we define $\mathrm{cl}(L_i) = \bigcup_n L_i^{(n)}$.

Definition 8.3.5 (Closure of a promise structure)**.** *The* closure *of the promise structure $\left(G, \vec{P}, \mathcal{L}\right)$ is the promise structure* $\mathrm{cl}(\mathcal{P}) = \left(\mathrm{cl}(G), \mathrm{cl}(\vec{P}), \mathrm{cl}(\mathcal{L})\right)$*, where:*

- $\mathrm{cl}(\vec{P}) = \left\{\vec{p}_i \colon \vec{p}_i \in \vec{P} \text{ is a placeholder-promise}\right\}$,
- $\mathrm{cl}(\mathcal{L}) = \{\mathrm{cl}(L_i) \colon \vec{p}_i \in \vec{P} \text{ is a placeholder-promise}\}$.

Proposition 8.3.6 ([29, Proposition 3.3])**.** *Let G be a graph and let $\left(G, \vec{P}, \mathcal{L}\right)$ be a promise structure. Then* $\mathrm{cl}(G)$ *satisfies:*

(cl.3) *for every* $\vec{p_i} \in \vec{P}$ *and every* $\ell \in L_i$,

$$\mathrm{cl}(G)(\vec{p_i}) \cong \mathrm{cl}(G)(\vec{q_\ell})$$

are isomorphic as rooted graphs, and this isomorphism is $\mathrm{cl}(\vec{P})$*-respecting with respect to* $\mathrm{cl}(\mathcal{P})$.

It is precisely this property (cl.3) of the promise closure that will allow us to maintain partial hypomorphisms during our recursive construction.

The last two results of this section serve as preparation for growing G_{n+1}, H_{n+1} and F_{n+1} 'in parallel', as outlined in the third stage of the algorithm in Section 8.2.2. If $\mathcal{L} = \{L_i : i \in I\}$ and $\mathcal{L}' = \{L'_i : i \in I\}$, we say a map $\psi : \bigcup \mathcal{L} \to \bigcup \mathcal{L}'$ is *colour-preserving* if $\psi(L_i) \subseteq L'_i$ for every i.

Lemma 8.3.7. *Let* $\left(G, \vec{P}, \mathcal{L}\right)$ *and* $\left(G', \vec{P'}, \mathcal{L}'\right)$ *be promise structures, and let* $G = H^{(0)} \subseteq H^{(1)} \subseteq \cdots$ *and* $G' = H'^{(0)} \subseteq H'^{(1)} \subseteq \cdots$ *be 1-step extensions approximating their respective closures.*

Assume that $\vec{P} = \{\vec{p}_1, \ldots, \vec{p}_k\}$ *and* $\vec{P'} = \{\vec{r}_1, \ldots, \vec{r}_k\}$*, and that there is a colour-preserving bijection*

$$\psi\colon \bigcup \mathcal{L} \to \bigcup \mathcal{L}'$$

such that (recall that $\mathcal{L}(\cdot)$ *is the set of leaves of a graph that are in* $\mathcal{L}$*)*

$$\psi \restriction G(\vec{p_i})\colon \mathcal{L}(G(\vec{p_i})) \to \mathcal{L}'(G'(\vec{r_i}))$$

is still a colour-preserving bijection for all $\vec{p_i} \in \vec{P}$.

Then for each $i \leqslant k$ *there is a sequence of colour-preserving bijections*

$$\alpha^i_n\colon \mathcal{L}\left(H^{(n)}(\vec{p_i})\right) \to \mathcal{L}'\left(H'^{(n)}(\vec{r_i})\right)$$

such that α^i_{n+1} *extends* α^i_n.

Proof. Fix i. We proceed by induction on n. Put $\alpha^i_0 := \psi \restriction G(\vec{p_i})$.

Now suppose that α^i_n exists. To form $H^{(n+1)}(\vec{p_i})$, we glued a copy of $G(\vec{p_j})$ to each $\ell \in L^{(n)}_j \cap H^{(n)}(\vec{p_i})$ for all $j \leqslant k$, and to construct $H'^{(n+1)}(\vec{r_i})$, we glued a copy of $G'(\vec{r_j})$ to each $\ell' \in L'^{(n)}_j \cap H'^{(n)}(\vec{r_i})$ for all $j \leqslant k$, in both cases keeping all copies of promise leaves.

By assumption, the second part can be phrased equivalently as: we glued on a copy of $G'(\vec{r_j})$ to each $\alpha^i_n(\ell)$ for $\ell \in L^{(n)}_j \cap H^{(n)}(\vec{r_i})$. Thus, we can now combine the bijections $\alpha^i_n(\ell)$ with all the individual bijections ψ between all newly added $G(\vec{p_j})$ and $G'(\vec{r_j})$ to obtain a bijection α^i_{n+1} as desired. □

Corollary 8.3.8. *In the above situation, for each* i *there is a colour-preserving bijection* α^i *between* $\mathcal{L}(\mathrm{cl}(G)(\vec{p_i}))$ *and* $\mathcal{L}'(\mathrm{cl}(G')(\vec{r_i}))$ *with respect to the promise closures* $\mathrm{cl}(\mathcal{P})$ *and* $\mathrm{cl}(\mathcal{P}')$.

Proof. Put $\alpha^i = \bigcup_n \alpha^i_n$. Because all α^i_n respected all colours, they respect in particular the placeholder promises which make up $\mathrm{cl}(\mathcal{P})$ and $\mathrm{cl}(\mathcal{P}')$. □

8.4 Thickening the graph

In this section, we lay the groundwork for the third stage of our algorithm, as outlined in Section 8.2.2. Our aim is to clarify how gluing a one-ended graph F onto a graph G affects automorphisms and the end-space of the resulting graph.

Definition 8.4.1 (Gluing sum)**.** *Given two graphs G and F, and a bijection ψ with* $\mathrm{dom}(\psi) \subseteq V(G)$ *and* $\mathrm{ran}(\psi) \subseteq V(F)$*, the* gluing sum of G and F along ψ*, denoted by $G \oplus_\psi F$, is the quotient graph $(G \cup F)/\sim$ where $v \sim \psi(v)$ for all $v \in$* $\mathrm{dom}(\psi)$*.*

Our first lemma of this section explains how a partial isomorphism from $G_n - x$ to $H_n - \phi(x)$ in our construction can be lifted to the gluing sum of G_n and H_n with a graph F respectively.

Lemma 8.4.2. *Let G, H and F be graphs, and consider two gluing sums $G \oplus_{\psi_G} F$ and $H \oplus_{\psi_H} F$ along partial bijections ψ_G and ψ_H. Suppose there exists an isomorphism $h\colon G - x \to H - y$ that restricts to a bijection between* $\mathrm{dom}(\psi_G)$ *and* $\mathrm{dom}(\psi_H)$*.*

Then h extends to an isomorphism $(G \oplus_{\psi_G} F) - x \to (H \oplus_{\psi_H} F) - y$ provided there is an automorphism π of F such that $\pi \circ \psi_G(v) = \psi_H \circ h(v)$ for all $v \in$ $\mathrm{dom}(\psi_G)$*.*

Proof. We verify that the map

$$\hat{h}\colon (G \oplus_{\psi_G} F) - x \to (H \oplus_{\psi_H} F) - y, \quad v \mapsto \begin{cases} h(v) & \text{if } v \in G - x, \text{ and} \\ \pi(v) & \text{if } v \in F \end{cases}$$

is a well-defined isomorphism. It is well-defined, since if $v \sim \psi_G(v)$ in $G \oplus_{\psi_G} F$, then $\hat{h}(v) \sim \hat{h}(\psi_G(v))$ in $H \oplus_{\psi_H} F$ by assumption on π. Moreover, since h and π are isomorphisms, it follows that $\hat{h}$ is an isomorphism, too. □

For the remainder of this section, all graphs are assumed to be locally finite. A *ray* in a graph G is a one-way infinite path. Given a ray R, then for any finite vertex set $S \subset V(G)$ there is a unique component $C(R, S)$ of $G - S$ containing a tail of R. An *end* in a graph is an equivalence class of rays under the relation

$$R \sim R' \Leftrightarrow \text{for every finite vertex set } S \subset V(G) \text{ we have } C(R, S) = C(R', S).$$

We denote by $\Omega(G)$ the set of ends in the graph G, and write $C(\omega, S) := C(R, S)$ with $R \in \omega$. Let $\Omega(\omega, S) = \{\omega' \colon C(\omega', S) = C(\omega, S)\}$. The singletons $\{v\}$ for $v \in V(G)$ and sets of the form $C(\omega, S) \cup \Omega(\omega, S)$ generate a compact metrizable topology on the set $V(G) \cup \Omega(G)$, which is known in the literature as $|G|$.[2] This topology allows us to talk about the closure of a set of vertices $X \subset V(G)$, denoted by $\overline{X}$. Write $\partial(X) = \overline{X} \setminus X = \overline{X} \cap \Omega(X)$ for the boundary of X: the collection of all ends in the closure of X. Then an end $\omega \in \Omega(G)$ lies in $\partial(X)$ if and only if for every finite vertex set $S \subset V(G)$, we have $|X \cap C(\omega, S)| = \infty$. Therefore $\Omega(G) = \partial(X)$ if and only if for every finite vertex set $S \subset V(G)$, every infinite component of $G - S$ meets X infinitely often. In this case we say that X is *dense* for $\Omega(G)$.

Finally, an end $\omega \in \Omega(G)$ is *free* if for some S, the set $\Omega(\omega, S) = \{\omega\}$. Then $\Omega'(G)$ denotes the *non-free* (or limit-)ends. Note that $\Omega'(G)$ is a closed subset of $\Omega(G)$.

Lemma 8.4.3. *For locally finite connected graphs G and F, consider the gluing sum $G \oplus_\psi F$ for a partial bijection ψ. If F is one-ended and* $\mathrm{dom}(\psi)$ *is infinite, then $\Omega(G \oplus_\psi F) \cong \Omega(G)/\partial(\mathrm{dom}(\psi))$.*

Proof. Note first that for locally finite graphs G and F, also $G \oplus_\psi F$ is locally finite. Observe further that all rays of the unique end of F are still equivalent in $G \oplus_\psi F$, and so $G \oplus_\psi F$ has an end $\hat{\omega}$ containing the single end of F.

We are going to define a continuous surjection $f \colon \Omega(G) \to \Omega(G \oplus_\psi F)$ with the property that f has precisely one non-trivial fibre, namely $f^{-1}(\hat{\omega}) = \partial(\mathrm{dom}(\psi))$. It then follows from

[2] Normally $|G|$ is defined on the 1-complex of G together with its ends, but for our purposes it will be enough to just consider the subspace $V(G) \cup \Omega(G)$. See the survey paper of Diestel [42] for further details.

definition of the quotient topology that f induces a continuous bijection from the compact space $\Omega(G)/\partial(\mathrm{dom}(\psi))$ to the Hausdorff space $\Omega(G \oplus_\psi F)$, which, as such, is necessarily a homeomorphism.

The mapping f is defined as follows. Given an end $\omega \in \Omega(G) \setminus \partial(\mathrm{dom}(\psi))$, there is a finite $S \subset V(G)$ such that $C(\omega, S) \cap \mathrm{dom}(\psi) = \emptyset$, and so $C = C(\omega, S)$ is also a component of $(G \oplus_\psi F) - S$, which is disjoint from F. Define f to be the identity between $\Omega(G) \cap \overline{C}$ and $\Omega(G \oplus_\psi F) \cap \overline{C}$, while for all remaining ends $\omega \in \Omega(G) \cap \overline{\mathrm{dom}(\psi)}$, we put $f(\omega) = \hat{\omega}$.

To see that this assignment is continuous at $\omega \in \Omega(G) \cap \overline{\mathrm{dom}(\psi)}$, it suffices to show that $C := C(\omega, S) \subset G - S$ is a subset of $C' := C(\hat{\omega}, S) \subset (G \oplus_\psi F) - S$ for any finite set $S \subset G \oplus_\psi F$. To see this inclusion, note that by choice of ω, we have $|\mathrm{dom}(\psi) \cap C| = \infty$. At the same time, since F is both one-ended and locally finite, $F - S$ has precisely one infinite component D and $F - D$ is finite, so as ψ is a bijection, there is $v \in \mathrm{dom}(\psi) \cap C$ with $\psi(v) \in D$ (in fact, there are infinitely many such v). Since v and $\psi(v)$ get identified in $G \oplus_\psi F$, we conclude that $C \cup D$ is connected in $(G \oplus_\psi F) - S$, and hence that $C \cup D \subset C'$ as desired.

Finally, to see that f is indeed surjective, note first that the fact that $\mathrm{dom}(\psi)$ is infinite implies that $\overline{\mathrm{dom}(\psi)} \cap \Omega(G) \neq \emptyset$, and so $\hat{\omega} \in \mathrm{ran}(f)$. Next, consider an end $\omega \in \Omega(G \oplus_\psi F)$ different from $\hat{\omega}$. Find a finite separator $S \subset V(G \oplus_\psi F)$ such that $C(\omega, S) \neq C(\hat{\omega}, S)$. It follows that $\mathrm{dom}(\psi) \cap C(\omega, S)$ is finite. So there is a finite $S' \supseteq S$ such that $C := C(\omega, S') \neq C(\hat{\omega}, S')$ and $\mathrm{dom}(\psi) \cap C = \emptyset$. So by definition, f is a bijection between $\Omega(G) \cap \overline{C}$ and $\Omega(G \oplus_\psi F) \cap \overline{C}$, so $\omega \in \mathrm{ran}(f)$. $\square$

Corollary 8.4.4. *Under the above assumptions, if* $\mathrm{dom}(\psi)$ *is dense for* $\Omega(G)$*, then* $G \oplus_\psi F$ *is one-ended.*

Corollary 8.4.5. *Under the above assumptions, if* $\overline{\mathrm{dom}(\psi)} \cap \Omega(G) = \Omega'(G)$*, then* $G \oplus_\psi F$ *has at most one non-free end.*

We remark that more direct proofs for Corollaries 8.4.4 and 8.4.5 can be given that do not need the full power of Lemma 8.4.3.

8.5 The construction

8.5.1 Preliminary definitions

In the precise statement of our construction in Section 8.5.2, we are going to employ the following notation.

Definition 8.5.1 (Mii-path)**.** *A path* $P = v_0, v_1, \ldots, v_n$ *in a graph* G *is called* internally isolated *if* $\deg_G(v_i) = 2$ *for all internal vertices* v_i *for* $0 < i < n$*. The path* P *is* maximal internally isolated *(or* mii *for short) if in addition* $\deg_G(v_0) \neq 2 \neq \deg_G(v_n)$*. An infinite path* $P = v_0, v_1, v_2, \ldots$ *is* mii *if* $\deg_G(v_0) \neq 2$ *and* $\deg_G(v_i) = 2$ *for all* $i \geqslant 1$*.*

Definition 8.5.2 (Mii-spectrum)**.** *The* mii-spectrum *of* G *is*

$$\Sigma(G) := \{k \in \mathbb{N} \colon G \text{ contains an mii-path of length } k\}.$$

If $\Sigma(G)$ *is finite, we let* $\sigma_0(G) = \max \Sigma(G)$ *and* $\sigma_1(G) = \max\left(\Sigma(G) \setminus \{\sigma_0(G)\}\right)$*.*

Lemma 8.5.3. *Let* e *be an edge of a locally finite graph* G*. If* $\Sigma(G)$ *is finite, then* $\Sigma(G - e)$ *is finite.*

Proof. Observe first that every vertex of degree $\leqslant 2$ in any graph can lie on at most one mii-path.

We now claim that for an edge $e = xy$, there are at most two finite mii-paths in $G - e$ which are not subpaths of finite mii-paths of G.

Indeed, if $\deg x = 3$ in G, then x can now be the interior vertex of one new finite mii-path in $G - e$. And if $\deg x = 2$ in G, then x can now be end-vertex of one new finite mii-path in $G - e$ (this is relevant if x lies on an infinite mii-path of G). The argument is for y is the same, so the claim follows. □

Definition 8.5.4 (Spectrally distinguishable)**.** *Given two graphs G and H, we say that G and H are* spectrally distinguishable *if there is some $k \geqslant 3$ such that $k \in \Sigma(G) \triangle \Sigma(H) = \Sigma(G) \setminus \Sigma(H) \cup \Sigma(H) \setminus \Sigma(G)$.*

Note that being spectrally distinguishable is a strong certificate for being non-isomorphic.

Definition 8.5.5 (k-ball)**.** *For G a subgraph of H, and $k > 0$, the k-ball $\mathrm{Ball}_H(G, k)$ is the induced subgraph of H on the set of vertices at distance at most k of some vertex of G.*

Definition 8.5.6 (proper Mii-extension; infinite growth)**.** *Let G be a graph, B a subset of leaves of G, and H a component of G.*

- *A graph $\hat{G} \supset H$ is an* mii-extension *of H at B to length k if $\mathrm{Ball}_{\hat{G}}(H, k)$ can be obtained from H by adjoining, at each vertex $l \in B \cap V(H)$, a new path of length k starting at l, and a new leaf whose only neighbour is l.*[3]

- *A leaf l in a graph G is* proper *if the unique neighbour of l in G has degree $\geqslant 3$. An mii-extension is called* proper *if every leaf in B is proper.*

- *An mii-extension $\hat{G}$ of G is* of infinite growth *if every component of $\hat{G} - G$ is infinite.*

8.5.2 The back-and-forth construction

Our aim in this section is to prove our main theorem announced in the introduction.

Theorem 12.1.4. *There are two (vertex)-hypomorphic infinite trees T and S with maximum degree three such that there is no embedding $T \hookrightarrow S$ or $S \hookrightarrow T$.*

To do this we shall recursively construct, for each $n \in \mathbb{N}$,

- disjoint rooted connected graphs G_n and H_n,
- disjoint sets R_n and B_n of proper leaves of the graph $G_n \cup H_n$,
- trees F_n,
- disjoint sets R'_n and B'_n of leaves of F_n,
- bijections $\psi_{G_n} \colon V(G_n) \cap (R_n \cup B_n) \to R'_n \cup B'_n$ and $\psi_{H_n} \colon V(H_n) \cap (R_n \cup B_n) \to R'_n \cup B'_n$,
- finite sets $X_n \subset V(G_n)$ and $Y_n \subset V(H_n)$, and bijections $\varphi_n \colon X_n \to Y_n$,
- a family of isomorphisms $\mathcal{H}_n = \{h_{n,x} \colon G_n - x \to H_n - \varphi_n(x) \colon x \in X_n\}$,
- a family of automorphisms $\Pi_n = \{\pi_{n,x} \colon F_n \to F_n \colon x \in X_n\}$,

[3] We note that this is a slightly different definition of an mii-extension to that in [29].

- a strictly increasing sequence of integers $k_n \geqslant 2$,

such that for all $n \in \mathbb{N}$:[4]

(†1) $G_{n-1} \subset G_n$ and $H_{n-1} \subset H_n$ as induced subgraphs,

(†2) the vertices of G_n and H_n all have degree at most 5,

(†3) the vertices of F_n all have degree at most 3,

(†4) the root of G_n is in R_n and the root of H_n is in B_n,

(†5) $\sigma_0(G_n) = \sigma_0(H_n) = k_n$,

(†6) G_n and H_n are spectrally distinguishable,

(†7) G_n and H_n have at most one end,

(†8) $\Omega(G_n \cup H_n) \subset \overline{R_n \cup B_n}$,

(†9) (a) G_n is a (proper) mii-extension of infinite growth of G_{n-1} at $R_{n-1} \cup B_{n-1}$ to length $k_{n-1}+1$, and
 (b) $\mathrm{Ball}_{G_n}(G_{n-1}, k_{n-1}+1)$ does not meet $R_n \cup B_n$,

(†10) (a) H_n is a (proper) mii-extension of infinite growth of H_{n-1} at $R_{n-1} \cup B_{n-1}$ to length $k_{n-1}+1$, and
 (b) $\mathrm{Ball}_{H_n}(H_{n-1}, k_{n-1}+1)$ does not meet $R_n \cup B_n$,

(†11) there are enumerations $V(G_n) = \{t_j : j \in J_n\}$ and $V(H_n) = \{s_j : j \in J_n\}$ such that

- $J_{n-1} \subset J_n \subset \mathbb{N}$,
- $\{t_j : j \in J_n\}$ extends the enumeration $\{t_j : j \in J_{n-1}\}$ of $V(G_{n-1})$, and similarly for $\{s_j : j \in J_n\}$,
- $|\mathbb{N} \setminus J_n| = \infty$,
- $\{0, 1, \ldots, n\} \subset J_n$,

(†12) $\{t_j, s_j : j \leqslant n\} \cap (R_n \cup B_n) = \emptyset$,

(†13) the finite sets of vertices X_n and Y_n satisfy $|X_n| = n = |Y_n|$, and

- $X_{n-1} \subset X_n$ and $Y_{n-1} \subset Y_n$,
- $\varphi_n \restriction X_{n-1} = \varphi_{n-1}$,
- $\{t_j : j \leqslant \lfloor (n-1)/2 \rfloor\} \subset X_n$ and $\{s_j : j \leqslant \lfloor n/2 \rfloor - 1\} \subset Y_n$,
- $(X_n \cup Y_n) \cap (R_n \cup B_n) = \emptyset$,

(†14) the families of isomorphisms $\mathcal{H}_n$ satisfy

- $h_{n,x} \restriction (G_{n-1} - x) = h_{n-1,x}$ for all $x \in X_{n-1}$,
- the image of $R_n \cap V(G_n)$ under $h_{n,x}$ is $R_n \cap V(H_n)$,
- the image of $B_n \cap V(G_n)$ under $h_{n,x}$ is $B_n \cap V(H_n)$ for all $x \in X_n$.

(†15) the families of automorphisms Π_n satisfy

[4] If the statement involves an object indexed by $n-1$ we only require that it holds for $n \geqslant 1$.

- $\pi_{n,x} \restriction R'_n$ is a permutation of R'_n for each $x \in X_n$,
- $\pi_{n,x} \restriction B'_n$ is a permutation of B'_n for each $x \in X_n$,
- for each $x \in X_n$, the following diagram commutes:

$$\begin{tikzcd}
\mathcal{L}(G_n) \arrow[r, "h_{n,x} \restriction \mathcal{L}(G_n)"] \arrow[d, "\psi_{G_n}"'] & \mathcal{L}(H_n) \arrow[d, "\psi_{H_n}"] \\
\mathcal{L}(F_n) \arrow[r, "\pi_{n,x} \restriction \mathcal{L}(F_n)"] & \mathcal{L}(F_n)
\end{tikzcd}$$

I.e. for every $\ell \in \mathcal{L}(G_n) := V(G_n) \cap (R_n \cup B_n)$ we have $\pi_{n,x}(\psi_{G_n}(\ell)) = \psi_{H_n}(h_{n,x}(\ell))$.

8.5.3 The construction yields the desired non-reconstructible one-ended graphs.

By property (†1), we have $G_0 \subset G_1 \subset G_2 \subset \cdots$ and $H_0 \subset H_1 \subset H_2 \subset \cdots$. Let G and H be the union of the respective sequences. Then both G and H are connected, and as a consequence of (†2), both graphs have maximum degree 5.

We claim that the map $\varphi = \bigcup_n \varphi_n$ is a hypomorphism between G and H. Indeed, it follows from (†11) and (†13) that φ is a well-defined bijection from $V(G)$ to $V(H)$. To see that φ is a hypomorphism, consider any vertex x of G. This vertex appears as some t_j in our enumeration of $V(G)$, so the map

$$h_x = \bigcup_{n>2j} h_{n,x} \colon G - x \to H - \varphi(x),$$

is a well-defined isomorphism by (†14) between $G - x$ and $H - \varphi(x)$.

Now suppose for a contradiction that there exists an isomorphism $f \colon G \to H$. Then $f(t_0)$ is mapped into H_n for some $n \in \mathbb{N}$. Properties (†5) and (†9) imply that after deleting all mii-paths in G of length $> k_n$, the connected component C of t_0 is a leaf extension of G_n adding one further leaf to every vertex in $V(G_n) \cap (R_n \cup B_n)$. Similarly, properties (†5) and (†10) imply that after deleting all mii-paths in H of length $> k_n$, the connected component D of $f(t_0)$ is a leaf-extension of H_n adding one further leaf to every vertex in $V(H_n) \cap (R_n \cup B_n)$. Note that f restricts to an isomorphism between C and D. However, since C and D are proper extensions, we have $\Sigma(C) \triangle \Sigma(G_n) \subseteq \{1, 2\}$ and $\Sigma(D) \triangle \Sigma(H_n) \subseteq \{1, 2\}$. Hence, since G_n and H_n are spectrally distinguishable by (†6), so are C and D, a contradiction. We have established that G and H are non-isomorphic reconstructions of each other.

Finally, for G being one-ended, we now show that for every finite vertex separator $S \subset V(G)$, the graph $G - S$ has only one infinite component (the argument for H is similar). Suppose for a contradiction $G - S$ has two infinite components C_1 and C_2. Consider n large enough such that $S \subset V(G_n)$. Since G_k is one-ended for all k by (†7), we may assume that $C_1 \cap G_k$ falls apart into finite components for all $k \geqslant n$. Since C_1 is infinite and connected, it follows from (†9)(b) that C_1 intersects $G_{n+1} - G_n$. But since G_{n+1} is an mii-extension of G_n of infinite growth by (†9)(a), we see that that $C_1 \cap (G_{n+1} - G_n)$ contains an infinite component, a contradiction.

8.5.4 The base case: there are finite rooted graphs G_0 and H_0 satisfying requirements (†1)–(†15).

Choose a pair of spectrally distinguishable, equally sized graphs G_0 and H_0 with maximum degree $\leqslant 5$ and $\sigma_0(G_0) = \sigma_0(H_0) = k_0$. Pick a proper leaf each as roots $\mathrm{r}(G_0)$ and $\mathrm{r}(H_0)$ for G_0 and H_0, and further proper leaves $\ell_b \in G_0$ and $\ell_r \in H_0$.

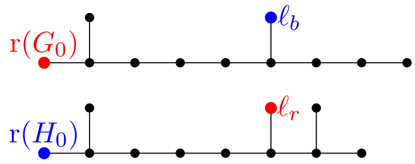

Figure 8.5: A possible choice for the finite rooted graphs G_0 and H_0.

Define $R_0 = \{\mathrm{r}(G_0), \ell_r\}$ and $B_0 = \{\mathrm{r}(H_0), \ell_b\}$. We take F_0 to be two vertices x and y joined by an edge, with $R'_0 = \{x\}$ and $B'_0 = \{y\}$ and take ψ_{G_0} to be the unique bijection sending $R_0 \cap G_0$ to R'_0 and $B_0 \cap G_0$ to B'_0, and similarly for ψ_{H_0}.

Figure 8.6: F_0.

Let $J_0 = \{0, 1, \dots, |G_0| - 1\}$ and choose enumerations $V(G_0) = \{t_j : j \in J_0\}$ and $V(H_0) = \{s_j : j \in J_0\}$ with $t_0 \neq \mathrm{r}(G_0)$ and $s_0 \neq \mathrm{r}(H_0)$. Finally we let $X_0 = Y_0 = \mathcal{H}_0 = \emptyset$. It is a simple check that conditions (†1)–(†15) are satisfied.

8.5.5 The inductive step: set-up

Now, assume that we have constructed graphs G_k and H_k for all $k \leqslant n$ such that (†1)–(†15) are satisfied up to n. If $n = 2m$ is even, then we have $\{t_j : j \leqslant m-1\} \subset X_n$ and in order to satisfy (†13) we have to construct G_{n+1} and H_{n+1} such that the vertex t_m is taken care of in our partial hypomorphism. Similarly, if $n = 2m+1$ is odd, then we have $\{s_j : j \leqslant m-1\} \subset Y_n$ and we have to construct G_{n+1} and H_{n+1} such that the vertex s_m is taken care of in our partial hypomorphism. Both cases are symmetric, so let us assume in the following that $n = 2m$ is even.

Now let v be the vertex with the least index in the set $\{t_j : j \in J_n\} \setminus X_n$, i.e.

$$v = t_i \quad \text{for} \quad i = \min\{j : t_j \in V(G_n) \setminus X_n\}. \tag{8.5.1}$$

Then by assumption (†13), v will be t_m, unless t_m was already in X_n anyway. In any case, since $|X_n| = |Y_n| = n$, it follows from (†11) that $i \leqslant n$, so by (†12), v does not lie in our leaf sets $R_n \cup B_n$, i.e.

$$v \notin R_n \cup B_n. \tag{8.5.2}$$

In the next sections, we will demonstrate how to obtain graphs $G_{n+1} \supset G_n$, $H_{n+1} \supset H_n$ and F_{n+1} with $X_{n+1} = X_n \cup \{v\}$ and $Y_{n+1} = Y_n \cup \{\varphi_{n+1}(v)\}$ satisfying (†1)—(†10) and (†13)–(†15).

After we have completed this step, since $|\mathbb{N} \setminus J_n| = \infty$, it is clear that we can extend our enumerations of G_n and H_n to enumerations of G_{n+1} and H_{n+1} as required, making sure to first list some new elements that do not lie in $R_{n+1} \cup B_{n+1}$. This takes care of (†11) and (†12) and completes the step $n \mapsto n+1$.

8.5.6 The inductive step: construction

We will construct the graphs G_{n+1} and H_{n+1} in three steps. First, in Section 8.5.6 we construct graphs $G'_{n+1} \supset G_n$ and $H'_{n+1} \supset H_n$ such that there is a vertex $\phi_{n+1}(v) \in H'_{n+1}$ with $G'_{n+1} - v \cong H'_{n+1} - \phi_{n+1}(v)$. This first step essentially follows the argument from [29, Section 4.6]. We will also construct a graph F_{n+1} via a parallel process.

Secondly, in Section 8.5.6 we will show that there are well-behaved maps from the coloured leaves of G'_{n+1} and H'_{n+1} to $F_{n+1} \times \mathbb{N}$, such that analogues of (†14) and (†15) hold for G'_{n+1}, H'_{n+1} and F_{n+1}, giving us control over the corresponding gluing sum.

Lastly, in Section 8.5.6, we do the actual gluing process and define all objects needed for step $n+1$ of our inductive construction.

Building the auxiliary graphs

Given the two graphs G_n and H_n, we extend each of them through their roots as indicated in Figure 8.7 to graphs $\tilde{G}_n$ and $\tilde{H}_n$ respectively.

Since v is not the root of G_n, there is a unique component of $G_n - v$ containing the root, which we call $G_n(r)$. Let $G_n(v)$ be the induced subgraph of G_n on the remaining vertices, including v. We remark that if v is not a cutvertex of G_n, then $G_n(v)$ is just a single vertex v. Since $\sigma_0(G_n) = k_n$ by (†5) and $\deg(v) \leqslant 5$ by (†2), it follows from an iterative application of Lemma 8.5.3 that $\Sigma\left(G_n(r)\right)$ and $\Sigma\left(G_n(v)\right)$ are finite. Let $k = \tilde{k}_n = \max\{\sigma_0(G_n), \sigma_0\left(G_n(r)\right), \sigma_0\left(G_n(v)\right), \sigma_0(H_n)\} + 1$.

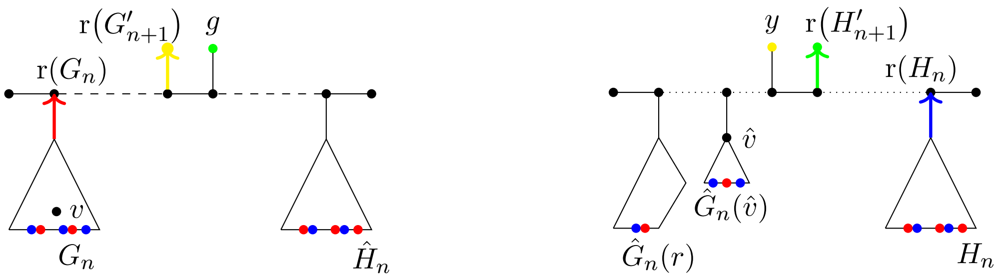

The graph $\tilde{G}_n$. The graph $\tilde{H}_n$.

Figure 8.7: All dotted lines are mii-paths of length at least $k + 1 = \tilde{k}_n + 1$.

To obtain $\tilde{G}_n$, we extend G_n through its root $\mathrm{r}(G_n) \in R_n$ by a path

$$\mathrm{r}(G_n) = u_0, u_1, \ldots, u_{p-1}, u_p = \mathrm{r}\left(\hat{H}_n\right)$$

of length $p = 4(\tilde{k}_n + 1) + 1$, where at its last vertex u_p we glue a rooted copy $\hat{H}_n$ of H_n (via an isomorphism $\hat{z} \leftrightarrow z$), identifying u_p with the root of $\hat{H}_n$.

Next, we add two additional leaves at u_0 and u_p, so that $\deg(\mathrm{r}(G_n)) = 3 = \deg\left(\mathrm{r}\left(\hat{H}_n\right)\right)$. Further, we add a leaf $\mathrm{r}\left(G'_{n+1}\right)$ at u_{2k+2}, which will be our new root for the next tree G'_{n+1}; and another leaf g at u_{2k+3}. This completes the construction of $\tilde{G}_n$.

The construction of $\tilde{H}_n$ is similar, but not entirely symmetric. For its construction, we extend H_n through its root $\mathrm{r}(H_n) \in B_n$ by a path

$$\mathrm{r}(H_n) = v_p, v_{p-1}, \ldots, v_1, v_0 = \mathrm{r}\left(\hat{G}_n(r)\right)$$

of length p, where at its last vertex v_0 we glue a copy $\hat{G}_n(r)$ of $G_n(r)$, identifying v_0 with the root of $\hat{G}_n(r)$. Then, we take a copy $\hat{G}_n(\hat{v})$ of $G_n(v)$ and connect $\hat{v}$ via an edge to v_{k+1}.

Finally, as before, we add two leaves at v_0 and v_p so that $\deg\left(\mathrm{r}\left(\hat{G}_n(r)\right)\right) = 3 = \deg\left(\mathrm{r}(H_n)\right)$. Next, we add a leaf $\mathrm{r}\left(H'_{n+1}\right)$ to v_{2k+3}, which will be our new root for the next tree H'_{n+1}; and another leaf y to v_{2k+2}. This completes the construction of $\tilde{H}_n$.

By the induction assumption, certain leaves of G_n have been coloured with one of the two colours in $R_n \cup B_n$, and also some leaves of H_n have been coloured with one of the two colours in $R_n \cup B_n$. In the above construction, we colour leaves of $\hat{H}_n$, $\hat{G}_n(r)$ and $\hat{G}_n(\hat{v})$ accordingly:

$$\begin{aligned}\tilde{R}_n &= \Big(R_n \cup \Big\{\hat{z} \in \hat{H}_n \cup \hat{G}_n(r) \cup \hat{G}_n(\hat{v})\colon z \in R_n\Big\}\Big) \setminus \Big\{\mathrm{r}(G_n), \mathrm{r}\Big(\hat{G}_n(r)\Big)\Big\},\\ \tilde{B}_n &= \Big(B_n \cup \Big\{\hat{z} \in \hat{H}_n \cup \hat{G}_n(r) \cup \hat{G}_n(\hat{v})\colon z \in B_n\Big\}\Big) \setminus \Big\{\mathrm{r}(H_n), \mathrm{r}\Big(\hat{H}_n\Big)\Big\}.\end{aligned} \tag{8.5.3}$$

Now put $M_n := \tilde{G}_n \cup \tilde{H}_n$ and consider the following promise structure $\mathcal{P} = \Big(M_n, \vec{P}, \mathcal{L}\Big)$ on M_n, consisting of four promise edges $\vec{P} = \{\vec{p}_1, \vec{p}_2, \vec{p}_3, \vec{p}_4\}$ and corresponding leaf sets $\mathcal{L} = \{L_1, L_2, L_3, L_4\}$, as follows:

$$\begin{aligned}&\bullet\ \vec{p}_1 \text{ pointing in } G_n \text{ towards } \mathrm{r}(G_n), \text{ with } L_1 = \tilde{R}_n,\\ &\bullet\ \vec{p}_2 \text{ pointing in } H_n \text{ towards } \mathrm{r}(H_n), \text{ with } L_2 = \tilde{B}_n,\\ &\bullet\ \vec{p}_3 \text{ pointing in } \tilde{G}_n \text{ towards } \mathrm{r}\big(G'_{n+1}\big), \text{ with } L_3 = \big\{\mathrm{r}\big(G'_{n+1}\big), y\big\},\\ &\bullet\ \vec{p}_4 \text{ pointing in } \tilde{H}_n \text{ towards } \mathrm{r}\big(H'_{n+1}\big), \text{ with } L_4 = \big\{\mathrm{r}\big(H'_{n+1}\big), g\big\}.\end{aligned} \tag{8.5.4}$$

Note that our construction so far has been tailored to provide us with a $\vec{P}$-respecting isomorphism

$$h\colon \tilde{G}_n - v \to \tilde{H}_n - \hat{v}. \tag{8.5.5}$$

Consider the closure $\mathrm{cl}(M_n)$ with respect to the above defined promise structure $\mathcal{P}$. Since $\mathrm{cl}(M_n)$ is a leaf-extension of M_n, it has two connected components, just as M_n. We now define

$$\begin{aligned}G'_{n+1} = &\text{ the component containing } G_n \text{ in } \mathrm{cl}(M_n),\\ H'_{n+1} = &\text{ the component containing } H_n \text{ in } \mathrm{cl}(M_n).\end{aligned} \tag{8.5.6}$$

It follows that $\mathrm{cl}(M_n) = G'_{n+1} \cup H'_{n+1}$. Further, since $\vec{p}_3$ and $\vec{p}_4$ are placeholder promises, $\mathrm{cl}(M_n)$ carries a corresponding promise structure, cf. Def. 8.3.5. We define

$$R_{n+1} = \mathrm{cl}(L_3) \text{ and } B_{n+1} = \mathrm{cl}(L_4). \tag{8.5.7}$$

Lastly, set

$$\begin{aligned}X_{n+1} &= X_n \cup \{v\},\\ Y_{n+1} &= Y_n \cup \{\hat{v}\},\\ \varphi_{n+1} &= \varphi_n \cup \{(v, \hat{v})\},\\ k_{n+1} &= 2(\tilde{k}_n + 1).\end{aligned} \tag{8.5.8}$$

We now build F_{n+1} in a similar fashion to the above procedure. That is, we take two copies of F_n and join them pairwise through their roots as indicated in Figure 8.7 to form a graph $\tilde{F}_n$. We consider the graph $N_n = \tilde{F}_n \cup \hat{\tilde{F}}_n$, and take F_{n+1} to be one of the components of $\mathrm{cl}(N_n)$ (unlike for $\mathrm{cl}(M_n)$, both components of $\mathrm{cl}(N_n)$ are isomorphic).

More precisely we take two copies of F_n, which we will denote by F_n^G and F_n^H. We extend F_n^G through the image of the $\mathrm{r}(G_n)$ under the bijection ψ_{G_n} by a path

$$\psi_{G_n}(\mathrm{r}(G_n)) = u_0, u_1, u_2, u_3 = \psi_{H_n}(\mathrm{r}(H_n))$$

of length three, where $\psi_{G_n}(\mathrm{r}(G_n))$ is taken in F_n^G and $\psi_{H_n}(\mathrm{r}(H_n))$ is taken in F_n^H. Further, we add a leaf x at u_1, and another leaf y at u_2. We will consider the graph $N_n = \tilde{F}_n \cup \hat{\tilde{F}}_n$ as in Figure 8.8 formed by taking two disjoint copies of $\tilde{F}_n$.

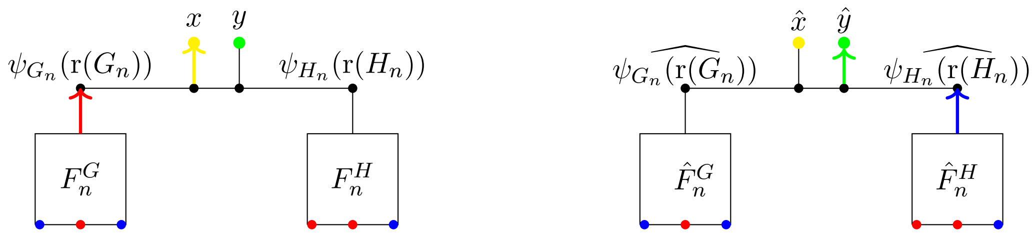

The graph $\tilde{F}_n$. The graph $\hat{\tilde{F}}_n$.

Figure 8.8: The graph $N_n = \tilde{F}_n \cup \hat{\tilde{F}}_n$.

By the induction assumption, certain leaves of F_n have been coloured with one of the two colours in $R'_n \cup B'_n$. In the above construction, we colour leaves of $F_n^G, F_n^H, \hat{F}_n^G$ and $\hat{F}_n^H$ accordingly:

$$\begin{aligned}
\tilde{R}'_n &= \left\{w \in F_n^G \cup F_n^H \cup \hat{F}_n^G \cup \hat{F}_n^H \;:\; w \in R'_n\right\} \setminus \left\{\psi_{G_n}(\mathrm{r}(G_n)), \psi_{G_n}\widehat{(\mathrm{r}(G_n))}\right\} \\
\tilde{B}'_n &= \left\{w \in F_n^G \cup F_n^H \cup \hat{F}_n^G \cup \hat{F}_n^H \;:\; w \in B'_n\right\} \setminus \left\{\psi_{H_n}(\mathrm{r}(H_n)), \psi_{H_n}\widehat{(\mathrm{r}(H_n))}\right\}.
\end{aligned} \tag{8.5.9}$$

Now consider the following promise structure $\mathcal{P}' = \left(N_n, \vec{P}', \mathcal{L}'\right)$ on N_n, consisting of four promise edges $\vec{P}' = \{\vec{r}_1, \vec{r}_2, \vec{r}_3, \vec{r}_4\}$ and corresponding leaf sets $\mathcal{L}' = \{L'_1, L'_2, L'_3, L'_4\}$, as follows:

$$\begin{aligned}
&\bullet\ \vec{r}_1 \text{ pointing in } F_n^G \text{ towards } \psi_{G_n}(\mathrm{r}(G_n)),\ \text{with } L'_1 = \tilde{R}'_n, \\
&\bullet\ \vec{r}_2 \text{ pointing in } \hat{F}_n^H \text{ towards } \psi_{H_n}(\mathrm{r}(H_n)),\ \text{with } L'_2 = \tilde{B}'_n, \\
&\bullet\ \vec{r}_3 \text{ pointing in } \tilde{F}_n \text{ towards } x,\ \text{with } L'_3 = \{x, \hat{x}\}, \\
&\bullet\ \vec{r}_4 \text{ pointing in } \hat{\tilde{F}}_n \text{ towards } \hat{y},\ \text{with } L'_4 = \{y, \hat{y}\}.
\end{aligned} \tag{8.5.10}$$

Consider the closure $\mathrm{cl}(N_n)$ with respect to the promise structure $\mathcal{P}'$ defined above. Since $\mathrm{cl}(N_n)$ is a leaf-extension of N_n, it has two connected components, and we define F_{n+1} to be the component containing F_n^G in $\mathrm{cl}(N_n)$. Since $\vec{r}_3$ and $\vec{r}_4$ are placeholder promises, $\mathrm{cl}(N_n)$ carries a corresponding promise structure, cf. Def. 8.3.5. We define

$$R'_{n+1} = \mathrm{cl}(L'_3) \cap F_{n+1} \quad \text{and} \quad B'_{n+1} = \mathrm{cl}(L'_4) \cap F_{n+1}. \tag{8.5.11}$$

Extending maps

In order to glue $F_{n+1} \times \mathbb{N}$ onto G'_{n+1} and H'_{n+1} we will need to show that that analogues of (†14) and (†15) hold for G'_{n+1}, H'_{n+1} and F_{n+1}. Our next lemma is essentially [29, Claim 4.13], and is an analogue of (†14). We briefly remind the reader of the details, as we need to know the nature of our extensions in our later claims.

Lemma 8.5.7. *There is a family of isomorphisms $\mathcal{H}'_{n+1} = \left\{h'_{n+1,x} \colon x \in X_{n+1}\right\}$ witnessing that $G'_{n+1} - x$ and $H'_{n+1} - \varphi_{n+1}(x)$ are isomorphic for all $x \in X_{n+1}$, such that $h'_{n+1,x}$ extends $h_{n,x}$ for all $x \in X_n$.*

Proof. The graphs G'_{n+1} and H'_{n+1} defined in (8.5.6) are obtained from $\tilde{G}_n$ and $\tilde{H}_n$ by attaching at every leaf in $\tilde{R}_n$ a copy of the rooted graph $\mathrm{cl}(M_n)(\vec{p}_1)$, and by attaching at every leaf in $\tilde{B}_n$ a copy of the rooted graph $\mathrm{cl}(M_n)(\vec{p}_2)$ by (cl.2).

From (8.5.5) we know that there is a $\vec{P}$-respecting isomorphism

$$h\colon \tilde{G}_n - v \to \tilde{H}_n - \varphi_{n+1}(v).$$

In other words, h maps promise leaves in $L_i \cap V(\tilde{G}_n)$ bijectively to the promise leaves in $L_i \cap V(\tilde{H}_n)$ for all $i = 1, 2, 3, 4$.

There is for each $\ell \in \tilde{R}_n \cup \tilde{B}_n \cup \{\mathrm{r}(G_n), \mathrm{r}(H_n)\}$ a $\mathrm{cl}(\vec{P})$-respecting isomorphism of rooted graphs

$$f_\ell\colon \mathrm{cl}(M_n)(\vec{q}_\ell) \cong \mathrm{cl}(M_n)(\vec{p}_i) \tag{8.5.12}$$

given by (cl.3) for $\ell \in (\tilde{R}_n \cup \tilde{B}_n)$, where i equals blue or red depending on whether $\ell \in \tilde{R}_n$ or $\tilde{B}_n$, and for the roots of G_n and H_n we have $\vec{q_r} = \vec{p_i}$ and the isomorphism is the identity. Hence, for each ℓ,

$$f^{-1}_{h(\ell)} \circ f_\ell\colon \mathrm{cl}(M_n)(\vec{q}_\ell) \cong \mathrm{cl}(M_n)(\vec{q}_{h(\ell)})$$

is a $\mathrm{cl}(\vec{P})$-respecting isomorphism of rooted graphs. By combining the isomorphism h between $\tilde{G}_n - v$ and $\tilde{H}_n - \varphi_{n+1}(v)$ with these isomorphisms between each $\mathrm{cl}(M_n)(\vec{q}_\ell)$ and $\mathrm{cl}(M_n)(\vec{q}_{h(\ell)})$ we get a $\mathrm{cl}(\vec{P})$-respecting isomorphism

$$h'_{n+1,v}\colon G'_{n+1} - v \to H'_{n+1} - \varphi_{n+1}(v).$$

To extend the old isomorphisms $h_{n,x}$ (for $x \in X_n$), note that G'_{n+1} and H'_{n+1} are obtained from G_n and H_n by attaching at every leaf in R_n a copy of the rooted graph $\mathrm{cl}(M_n)(\vec{p}_1)$, and similarly by attaching at every leaf in B_n a copy of the rooted graph $\mathrm{cl}(M_n)(\vec{p}_2)$. By induction assumption (†14), for each $x \in X_n$ the isomorphism

$$h_{n,x}\colon G_n - x \to H_n - \varphi_n(x)$$

maps the red leaves of G_n bijectively to the red leaves of H_n, and the blue leaves of G_n bijectively to the blue leaves of H_n. Thus, by (8.5.12),

$$f^{-1}_{h_{n,x}(\ell)} \circ f_\ell\colon \mathrm{cl}(M_n)(\vec{q}_\ell) \cong \mathrm{cl}(M_n)(\vec{q}_{h_{n,x}(\ell)})$$

are $\mathrm{cl}(\vec{P})$-respecting isomorphisms of rooted graphs for all $\ell \in (R_n \cup B_n) \cap V(G_n)$. By combining the isomorphism $h_{n,x}$ between $G_n - x$ and $H_n - \varphi_n(x)$ with these isomorphisms between each $\mathrm{cl}(M_n)(\vec{q}_\ell) = G'_{n+1}(\vec{q}_\ell)$ and $\mathrm{cl}(M_n)(\vec{q}_{h_{n,x}(l)}) = H'_{n+1}(\vec{q}_{h_{n,x}(l)})$, we obtain a $\mathrm{cl}(\vec{P})$-respecting extension

$$h'_{n+1,x}\colon G'_{n+1} - x \to H'_{n+1} - \varphi_n(x). \qquad \square$$

Our next claim should be seen as an approximation to property (†15). Recall that $\mathrm{cl}(N_n)$ has two components $F_{n+1} \cong \hat{F}_{n+1}$.

Lemma 8.5.8. *There are colour-preserving bijections*

$$\psi_{G'_{n+1}}\colon V(G'_{n+1}) \cap (R_{n+1} \cup B_{n+1}) \to R'_{n+1} \cup B'_{n+1},$$

$$\psi_{H'_{n+1}}\colon V(H'_{n+1}) \cap (R_{n+1} \cup B_{n+1}) \to \hat{R}'_{n+1} \cup \hat{B}'_{n+1},$$

and a family of isomorphisms

$$\hat{\Pi}_{n+1} = \left\{\hat{\pi}_{n+1,x}\colon F_{n+1} \to \hat{F}_{n+1} \colon x \in X_{n+1}\right\}$$

such that for each $x \in X_{n+1}$ the following diagram commutes.

$$\begin{tikzcd}
\mathcal{L}(G'_{n+1}) \arrow[r, "h'_{n+1,x} \restriction \mathcal{L}(G'_{n+1})"] \arrow[d, "\psi_{G'_{n+1}}"'] & \mathcal{L}(H'_{n+1}) \arrow[d, "\psi_{H'_{n+1}}"] \\
\mathcal{L}'(F_{n+1}) \arrow[r, "\hat{\pi}_{n+1,x} \restriction \mathcal{L}'(F_{n+1})"] & \mathcal{L}'(\hat{F}_{n+1})
\end{tikzcd}$$

Proof. Defining $\psi_{G'_{n+1}}$ and $\psi_{H'_{n+1}}$. By construction, we can combine the maps ψ_{G_n} and ψ_{H_n} to obtain a natural colour-preserving bijection

$$\psi\colon \mathcal{L}(M_n) \to \mathcal{L}'(N_n),$$

which satisfies the assumptions of Lemma 8.3.7. Thus, by Corollary 8.3.8, there are bijections

$$\alpha^i\colon \mathcal{L}(\mathrm{cl}(M_n)(\vec{p}_i)) \to \mathcal{L}'(\mathrm{cl}(N_n)(\vec{r}_i))$$

which are colour-preserving with respect to the promise structures $\mathrm{cl}(\mathcal{P})$ and $\mathrm{cl}(\mathcal{P}')$ on $\mathrm{cl}(M_n)$ and $\mathrm{cl}(N_n)$, respectively.

We now claim that ψ extends to a colour-preserving bijection (w.r.t. $\mathrm{cl}(\mathcal{P})$)

$$\mathrm{cl}(\psi)\colon \mathcal{L}(\mathrm{cl}(M_n)) \to \mathcal{L}'(\mathrm{cl}(N_n)).$$

Indeed, by (cl.3), for every $\ell \in \tilde{R}'_n \cup \tilde{B}'_n$, there is a $\vec{P}'$-respecting rooted isomorphism

$$g_\ell\colon \mathrm{cl}(N_n)(\vec{q}_\ell) \to \mathrm{cl}(N_n)(\vec{r}_i), \tag{8.5.13}$$

where i equals blue or red depending on whether $\ell \in \tilde{R}'_n$ or $\tilde{B}'_n$. As in the case of (8.5.12) we define the maps g_r with $\vec{q}_r = \vec{r}_i$ for the roots of F_n^G and $\hat{F}_n^H$ respectively to be the identity. Together with the rooted isomorphisms f_ℓ from (8.5.12), it follows that for each $\ell \in \tilde{R}_n \cup \tilde{B}_n \cup \{\mathrm{r}(G_n), \mathrm{r}(H_n)\}$, the map

$$\psi_\ell = g^{-1}_{\psi(\ell)} \circ \alpha^i \circ f_\ell\colon \mathcal{L}(\mathrm{cl}(M_n)(\vec{q}_\ell)) \to \mathcal{L}\big(\mathrm{cl}(N_n)(\vec{q}_{\psi(\ell)})\big)$$

is a colour-preserving bijection. Now combine ψ with the individual ψ_ℓ to obtain $\mathrm{cl}(\psi)$. We then put

$$\psi_{G'_{n+1}} = \mathrm{cl}(\psi) \restriction G'_{n+1} \quad \text{and} \quad \psi_{H'_{n+1}} = \mathrm{cl}(\psi) \restriction H'_{n+1}.$$

Defining isomorphisms $\hat{\Pi}_{n+1}$. To extend the old isomorphisms $\pi_{n,x}$, given by the induction assumption, note that by (cl.2), F_{n+1} is obtained from F_n by attaching at every leaf in R'_n a copy of the rooted graph $F_{n+1}(\vec{r}_1)$, and similarly by attaching at every leaf in B'_n a copy of the rooted graph $F_{n+1}(\vec{r}_2)$. For each $x \in X_n$ let us write $\hat{\pi}_{n,x}$ for the map sending each $z \in F_n^G$ to the copy of $\pi_{n,x}(z)$ in $\hat{F}_n^H$. By the induction assumption (†15), for each $x \in X_n$ the isomorphism

$$\hat{\pi}_{n,x}\colon F_n^G \to \hat{F}_n^H$$

preserves the colour of red and blue leaves. Thus, using the maps g_ℓ from (8.5.13), the mappings

$$g^{-1}_{\hat{\pi}_{n,x}(\ell)} \circ g_\ell\colon \mathrm{cl}(N_n)(\vec{q}_\ell) \cong \mathrm{cl}(N_n)(\vec{q}_{\hat{\pi}_{n,x}(\ell)})$$

are $\mathrm{cl}(\vec{P}')$-respecting isomorphisms of rooted graphs for all $\ell \in R'_n \cup B'_n$. By combining the isomorphism $\pi_{n,x}$ with these isomorphisms between each $F_{n+1}(\vec{q}_\ell)$ and $\hat{F}_{n+1}(\vec{q}_{\hat{\pi}_{n,x}(\ell)})$, we obtain a $\mathrm{cl}(\vec{P}')$-respecting extension

$$\hat{\pi}_{n+1,x}\colon F_{n+1} \to \hat{F}_{n+1}.$$

For the new isomorphism $\hat{\pi}_{n+1,v} \colon F_{n+1} \to \hat{F}_{n+1}$, we simply take the 'identity' map which extends the map sending each $z \in \tilde{F}_n$ to $\hat{z} \in \hat{F}_n$.

The diagram commutes. To see that the new diagram above commutes, for each $x \in X_n$ it suffices to check that for all $\ell \in (R_n \cup B_n) \cap V(G_n)$ we have

$$\hat{\pi}_{n+1,x} \circ \psi_{G'_{n+1}} \restriction \mathcal{L}\big(G'_{n+1}(\vec{q}_\ell)\big) = \psi_{H'_{n+1}} \circ h'_{n+1,x} \restriction \mathcal{L}\big(G'_{n+1}(\vec{q}_\ell)\big),$$

which by construction of $\mathrm{cl}(\psi)$ above is equivalent to showing that

$$\hat{\pi}_{n+1,x} \circ \psi_\ell = \psi_{h_{n,x}(\ell)} \circ h'_{n+1,x}.$$

By definition of ψ_ℓ this holds if and only if

$$\hat{\pi}_{n+1,x} \circ g^{-1}_{\psi(\ell)} \circ \alpha^i \circ f_\ell = g^{-1}_{\psi(h_{n,x}(\ell))} \circ \alpha^i \circ f_{h_{n,x}(\ell)} \circ h'_{n+1,x}.$$

Now by construction of $\hat{\pi}_{n+1,x}$ and $h'_{n+1,x}$, we have

$$\hat{\pi}_{n+1,x} \circ g^{-1}_{\psi(\ell)} = g^{-1}_{\hat{\pi}_{n,x}(\psi(\ell))} \quad \text{and} \quad f_{h_{n,x}(\ell)} \circ h'_{n+1,x} = f_\ell.$$

Hence, the above is true if and only if

$$g^{-1}_{\hat{\pi}_{n,x}(\psi(\ell))} \circ \alpha^i \circ f_\ell = g^{-1}_{\psi(h_{n,x}(\ell))} \circ \alpha^i \circ f_\ell.$$

Finally, this last line holds since $\psi(\ell) = \psi_{G_n}(\ell)$ and $\psi(h_{n,x}(\ell)) = \psi_{H_n}(h_{n,x}(\ell))$ by definition of ψ, and because

$$\hat{\pi}_{n,x} \circ \psi_{G_n}(\ell) = \psi_{H_n} \circ h_{n,x}(\ell)$$

by the induction assumption.

For $\hat{\pi}_{n+1,v}$ we see that, as above, it will be sufficient to show that for all $\ell \in (\tilde{R}_n \cup \tilde{B}_n) \cap V(\tilde{G}_n)$ we have

$$\hat{\pi}_{n+1,v} \circ \psi_\ell = \psi_{h'_{n+1,v}(\ell)} \circ h'_{n+1,v},$$

which reduces as before to showing that,

$$g^{-1}_{\hat{\pi}_{n+1,v}(\psi(\ell))} \circ \alpha^i \circ f_\ell = g^{-1}_{\psi(h'_{n+1,v}(\ell))} \circ \alpha^i \circ f_\ell.$$

Recall that, $\hat{\pi}_{n+1,v}$ sends each v to $\hat{v}$ and also, since $h'_{n+1,v} \restriction \tilde{G}_n = h$, the image of every leaf $\ell \in (\tilde{R}_n \cup \tilde{B}_n) \cap V(\tilde{G}_n)$ is simply $\hat{l} \in \hat{G}_n(v) \cup \hat{G}_n(r)$. Hence we wish to show that

$$g^{-1}_{(\hat{\psi(\ell)})} \circ \alpha^i \circ f_\ell = g^{-1}_{\psi(\hat{l})} \circ \alpha^i \circ f_\ell,$$

that is,

$$(\hat{\psi(\ell)}) = \psi(\hat{l}),$$

which follows from the construction of ψ. □

Gluing the graphs together

Let us take the cartesian product of F_{n+1} with a ray, which we simply denote by $F_{n+1} \times \mathbb{N}$. If we identify F_{n+1} with the subgraph $F_{n+1} \times \{0\}$, then we can interpret both $\psi_{G'_{n+1}}$ and $\psi_{H'_{n+1}}$ as maps from $\mathcal{L}(G'_{n+1})$ and $\mathcal{L}(H'_{n+1})$ to a set of vertices in $F_{n+1} \times \mathbb{N}$, under the natural isomorphism between $\hat{F}_{n+1}$ and F_{n+1}.

Instead of using the function $\psi_{G'_{n+1}}$ directly for our gluing operation, we identify, for every leaf l in $\mathcal{L}(G'_{n+1})$ the *unique neighbour* of l with $\psi_{G'_{n+1}}(l)$. Formally, define a bijection $\chi_{G_{n+1}}$ between the neighbours of $\mathcal{L}(G'_{n+1})$ and $\mathcal{L}'(F_{n+1})$ via

$$\chi_{G_{n+1}} = \Big\{(z_1, z_2) \colon \exists l \in \mathcal{L}(G'_{n+1}) \text{ s.t. } z_1 \in N(\ell) \text{ and } \psi_{G'_{n+1}}(l) = z_2\Big\}, \tag{8.5.14}$$

and similarly

$$\chi_{H_{n+1}} = \Big\{(z_1, z_2) \colon \exists l \in \mathcal{L}(H'_{n+1}) \text{ s.t. } z_1 \in N(\ell) \text{ and } \psi_{H'_{n+1}}(l) = z_2\Big\}. \tag{8.5.15}$$

Since two promise leaves in G'_{n+1} or H'_{n+1} are never adjacent to the same vertex, $\chi_{G_{n+1}}$ and $\chi_{H_{n+1}}$ are indeed bijections. Moreover, since all promise leaves were proper, the vertices in the domain of $\chi_{G_{n+1}}$ and $\chi_{H_{n+1}}$ have degree at least 3. Using our notion of gluing-sum (see Def. 8.4.1), we now define

$$G_{n+1} := G'_{n+1} \oplus_{\chi_{G_{n+1}}} (F_{n+1} \times \mathbb{N}) \text{ and } H_{n+1} := H'_{n+1} \oplus_{\chi_{H_{n+1}}} (F_{n+1} \times \mathbb{N}). \tag{8.5.16}$$

We consider R_{n+1}, B_{n+1}, X_{n+1} and Y_{n+1} as subsets of G_{n+1} and H_{n+1} in the natural way. Then $\psi_{G_{n+1}}$ and $\psi_{H_{n+1}}$ can be taken to be the maps $\psi_{G'_{n+1}}$ and $\psi_{H'_{n+1}}$, again identifying $\hat{F}_{n+1}$ with F_{n+1} in the natural way. We also take the roots of G_{n+1} and H_{n+1} to be the roots of G'_{n+1} and H'_{n+1} respectively

This completes the construction of graphs G_{n+1}, H_{n+1}, and F_{n+1}, the coloured leaf sets $R_{n+1}, B_{n+1}, R'_{n+1}$, and B'_{n+1}, the bijections $\psi_{G_{n+1}}$ and $\psi_{H_{n+1}}$, as well as $\varphi_{n+1} \colon X_{n+1} \to Y_{n+1}$, and $k_{n+1} = 2(k_n + 1)$. In the next section, we show the existence of families of isomorphisms $\mathcal{H}_{n+1}$ and Π_{n+1}, and verify that (†1)–(†15) are indeed satisfied for the $(n+1)^{\text{th}}$ instance.

8.5.7 The inductive step: verification

Lemma 8.5.9. *We have $G_n \subset G_{n+1}$, $H_n \subset H_{n+1}$, $\Delta(G_{n+1}), \Delta(H_{n+1}) \leqslant 5$, $\Delta(F_{n+1}) \leqslant 3$, and the roots of G_{n+1} and H_{n+1} are in R_{n+1} and B_{n+1} respectively.*

Proof. We note that $G_n \subset G'_{n+1}$ by construction. Hence, it follows that

$$G_n \subset G'_{n+1} \subset G'_{n+1} \oplus_{\chi_{G_{n+1}}} (F_{n+1} \times \mathbb{N}) = G_{n+1},$$

and similarly for H_n. Since we glued together degree 3 and degree 2 vertices, and $\Delta(G_n), \Delta(H_n) \leqslant 5$ and $\Delta(F_n) \leqslant 3$, it is clear that the same bounds hold for $n+1$. Finally, since the root of $\tilde{G}_n$ was a placeholder promise, and R_{n+1} was the corresponding set of promise leaves in $\mathrm{cl}(\tilde{G}_n)$, it follows that the root of G'_{n+1} is in R_{n+1}, and hence so is the root of G_{n+1}. A similar argument shows that the root of H_{n+1} is in B_{n+1}. □

Lemma 8.5.10. *We have $\sigma_0(G_{n+1}) = \sigma_0(H_{n+1}) = k_{n+1}$.*

Proof. By construction we have that $\sigma_0(\tilde{G}_n) = \sigma_0(\tilde{H}_n) = k_{n+1}$. Since G'_{n+1} and H'_{n+1} are realised as components of the promise closure of M_n, and this was a proper extension, it is a simple check that $\sigma_0(G'_{n+1}) = \sigma_0(H'_{n+1}) = k_{n+1}$. Also note that $F_{n+1} \times \mathbb{N}$ has no mii-paths of length bigger than two, since the vertices of degree two in $F_{n+1} \times \mathbb{N}$ are precisely those of the form $(\ell, 0)$ with ℓ a leaf of F_{n+1}.

Since $G'_{n+1} \oplus_{\chi_{G_{n+1}}} (F_{n+1} \times \mathbb{N})$ is formed by gluing a set of degree-two vertices of $F_{n+1} \times \mathbb{N}$ to a set of degree-three vertices in G'_{n+1}, it follows that $\sigma_0(G_{n+1}) = k_{n+1}$ as claimed. A similar argument shows that $\sigma_0(H_{n+1}) = k_{n+1}$. □

Lemma 8.5.11. *The graphs G_{n+1} and H_{n+1} are spectrally distinguishable.*

Proof. Since in $\tilde{G}_n$ we have that all long mii-paths except for those of length k_{n+1} are contained inside G_n or $\hat{H}_n$, it follows from our induction assumption (†5) that $\sigma_1(\tilde{G}_n) = k_n$. However, in $\tilde{H}_n$, we attached $\hat{G}_n(\hat{v})$ to generate an mii-path of length $\tilde{k}_n + 1$ in $\tilde{H}_n$ (see Fig. 8.7), implying that

$$\sigma_1(\tilde{H}_n) = \tilde{k}_n + 1 > k_n = \sigma_1(\tilde{G}_n).$$

As before, since the promise closures G'_{n+1} and H'_{n+1} are proper extensions of $\tilde{G}_n$ and $\tilde{H}_n$, they are spectrally distinguishable. Lastly, since $F_{n+1} \times \mathbb{N}$ has no leaves and no mii-paths of length bigger than two, the same is true for G_{n+1} and H_{n+1}. □

Lemma 8.5.12. *The graphs G_{n+1} and H_{n+1} have exactly one end, and $\Omega(G_{n+1} \cup H_{n+1}) \subset \overline{R_{n+1} \cup B_{n+1}}$.*

Proof. By the induction assumption (†8), we know that $\Omega(G_n \cup H_n) \subset \overline{R_n \cup B_n}$.

Claim. *The set $R_{n+1} \cup B_{n+1}$ is dense for G'_{n+1}.*

Consider a finite $S \subset V(G'_{n+1})$. We have to show that any infinite component C of $G'_{n+1} - S$ has non-empty intersection with $R_{n+1} \cup B_{n+1}$.

Let us consider the global structure of G'_{n+1} as being roughly that of an infinite regular tree, as in Figure 8.2. Specifically, we imagine a copy of G_n at the top level, at the next level are the copies of G_n and H_n that come from a blue or red leaf in the top level, at the next level are the copies attached to blue or red leaves from the previous level, and so on.

With this in mind, it is evident that either C contains an infinite component from some copy of $H_n - S$ or $G_n - S$, or C contains an infinite ray from this tree structure. In the first case, we have $|C \cap (R_n \cup B_n)| = \infty$ by induction assumption. Since any vertex from $R_n \cup B_n$ has a leaf from $R_{n+1} \cup B_{n+1}$ within distance $k_{n+1} + 1$ (cf. Figure 8.7), it follows that C also meets $R_{n+1} \cup B_{n+1}$ infinitely often. In the second case, the same conclusion follows, since between each level of our tree structure, there is a pair of leaves in $R_{n+1} \cup B_{n+1}$. This establishes the claim.

Claim. *The set $R_{n+1} \cup B_{n+1}$ is dense for H'_{n+1}.*

The proof of the second claim is entirely symmetric to the first claim.

To complete the proof of the lemma, observe that $F_{n+1} \times \mathbb{N}$ is one-ended, and with $R_{n+1} \cup B_{n+1}$, also $\mathrm{dom}(\chi_{G_{n+1}}) \cup \mathrm{dom}(\chi_{H_{n+1}})$ is dense for $G'_{n+1} \cup H'_{n+1}$ by our claims. So by Corollary 8.4.4, the graphs G_{n+1} and H_{n+1} have exactly one end. Moreover, since $R_{n+1} \cup B_{n+1}$ meets both graphs infinitely, it follows immediately that it is dense for $G_{n+1} \cup H_{n+1}$. □

Lemma 8.5.13. *The graph G_{n+1} is a proper mii-extension of infinite growth of G_n at $R_n \cup B_n$ to length $k_n + 1$, and $\mathrm{Ball}_{G_{n+1}}(G_n, k_n + 1)$ does not meet $R_{n+1} \cup B_{n+1}$. Similarly, H_{n+1} is a proper mii-extension of infinite growth of H_n at $R_n \cup B_n$ to length $k_n + 1$, and $\mathrm{Ball}_{H_{n+1}}(H_n, k_n + 1)$ does not meet $R_{n+1} \cup B_{n+1}$. Hence,* (†9) *and* (†10) *are satisfied at stage $n + 1$.*

Proof. We will just prove the statement for G_{n+1}, as the corresponding proof for H_{n+1} is analogous.

Since G'_{n+1} is an $\Big((\tilde{R}_n \cup \tilde{B}_n) \cap V(\tilde{G}_n)\Big)$-extension of $\tilde{G}_n$, it follows that G'_{n+1} is an

$$\Big(\Big((\tilde{R}_n \cup \tilde{B}_n) \cap V(G_n)\Big) \cup \mathrm{r}(G_n)\Big) = \big((R_n \cup B_n) \cap V(G_n)\big)\text{-extension of } G_n. \tag{8.5.17}$$

However, from the construction of the closure of a graph it is clear that that G'_{n+1} is also an L'-extension of the supergraph K of G_n formed by gluing a copy of $\tilde{G}_n(\vec{p_1})$ to every leaf in

$R_n \cap V(G_n)$ and a copy of $\tilde{H}_n(\vec{p_2})$ to every leaf in $B_n \cap V(G_n)$, where L' is defined as the set of inherited promise leaves from the copies of $\tilde{G}_n(\vec{p_1})$ and $\tilde{H}_n(\vec{p_2})$.

However, we note that every promise leaf in $\tilde{G}_n(\vec{p_1})$ and $\tilde{H}_n(\vec{p_2})$ is at distance at least $\tilde{k}_n+1$ from the respective root, and so $\mathrm{Ball}_{G'_{n+1}}(G_n, \tilde{k}_n) = \mathrm{Ball}_K(G_n, \tilde{k}_n)$. However, $\mathrm{Ball}_K(G_n, \tilde{k}_n)$ can be seen immediately to be an mii-extension of G_n at $R_n \cup B_n$ to length $\tilde{k}_n$, and since $\tilde{k}_n \geqslant k_n+1$ it follows that $\mathrm{Ball}_{G'_{n+1}}(G_n, k_n+1)$ is an mii-extension of G_n at $R_n \cup B_n$ to length k_n+1 as claimed.

Finally, we note that $R_{n+1} \cup B_{n+1}$ is the set of promise leaves $\mathrm{cl}(\mathcal{L}_n)$. By the same reasoning as before, $\mathrm{Ball}_{G'_{n+1}}(G_n, k_n+1)$ contains no promise leaf in $\mathrm{cl}(\mathcal{L}_n)$, and so does not meet $R_{n+1} \cup B_{n+1}$ as claimed. Furthermore, it doesn't meet any neighbours of $R_{n+1} \cup B_{n+1}$.

Recall that G_{n+1} is formed by gluing a set of vertices in $(F_{n+1} \times \mathbb{N})$ to neighbours of vertices in $R_{n+1} \cup B_{n+1}$. However, by the above claim, $\mathrm{Ball}_{G'_{n+1}}(G_n, k_n+1)$ does not meet any of the neighbours of $R_{n+1} \cup B_{n+1}$ and so $\mathrm{Ball}_{G_{n+1}}(G_n, k_n+1) = \mathrm{Ball}_{G'_{n+1}}(G_n, k_n+1)$, and the claim follows.

Finally, to see that G_{n+1} is a leaf extension of G_n of infinite growth, it suffices to observe that $G_{n+1} - G_n$ consists of one component only, which is a superset of the infinite graph $F_n \times \mathbb{N}$. $\square$

Lemma 8.5.14. *There is a family of isomorphisms*

$$\mathcal{H}_{n+1} = \{h_{n+1,x} \colon G_{n+1} - x \to H_{n+1} - \varphi_{n+1}(x) \colon x \in X_{n+1}\},$$

such that

- $h_{n+1,x} \upharpoonright (G_n - x) = h_{n,x}$ *for all* $x \in X_n$,
- *the image of* $R_{n+1} \cap V(G_{n+1})$ *under* $h_{n+1,x}$ *is* $R_{n+1} \cap V(H_{n+1})$,
- *the image of* $B_{n+1} \cap V(G_{n+1})$ *under* $h_{n+1,x}$ *is* $B_{n+1} \cap V(H_{n+1})$ *for all* $x \in X_{n+1}$.

Proof. Recall that Lemma 8.5.7 shows that the there exists such a family of isomorphisms between G'_{n+1} and H'_{n+1}. Furthermore, we have that

$$G_{n+1} := G'_{n+1} \oplus_{\chi_{G_{n+1}}} (F_{n+1} \times \mathbb{N}) \text{ and } H_{n+1} := H'_{n+1} \oplus_{\chi_{H_{n+1}}} (F_{n+1} \times \mathbb{N}).$$

where it is easy to check that $\chi_{G_{n+1}}$ and $\chi_{H_{n+1}}$ satisfy the assumptions of Lemma 8.4.2, since the functions $\psi_{G'_{n+1}}$ and $\psi_{H'_{n+1}}$ do by Lemma 8.5.8.

More precisely, given $x \in X_{n+1}$ and $h'_{n+1,x}$, it follows from Lemma 8.5.8 that

$$\chi_{H_{n+1}} \circ h'_{n+1,x} \circ \chi_{G_{n+1}}$$

extends to an isomorphism $\pi_{n+1,x}$ of F_{n+1}. Hence, by Lemma 8.4.2, $h'_{n+1,x}$ extends to an isomorphism $h_{n+1,x}$ from $G_{n+1} - x$ to $H_{n+1} - y$. That this isomorphism satisfies the three properties claimed follows immediately from Lemma 8.5.7 and the fact that $h_{n+1,x} \upharpoonright (G_n - x) = h'_{n+1,x} \upharpoonright (G_n - x)$. $\square$

Lemma 8.5.15. *There exist bijections*

$$\psi_{G_{n+1}} \colon V(G_{n+1}) \cap (R_{n+1} \cup B_{n+1}) \to R'_{n+1} \cup B'_{n+1}$$

and

$$\psi_{H_{n+1}} \colon V(H_{n+1}) \cap (R_{n+1} \cup B_{n+1}) \to R'_{n+1} \cup B'_{n+1},$$

and a family of isomorphisms

$$\Pi_{n+1} = \{\pi_{n+1,x} \colon F_{n+1} \to F_{n+1} \colon x \in X_{n+1}\},$$

such that

- *$\pi_{n+1,x} \restriction R'_{n+1}$ is a permutation of R'_{n+1} for each x,*
- *$\pi_{n+1,x} \restriction B'_{n+1}$ is a permutation of B'_{n+1} for each x, and*
- *for each $x \in X_{n+1}$, the corresponding diagram commutes:*

$$\begin{tikzcd}
\mathcal{L}(G_{n+1}) \arrow[r, "h_{n+1,x} \restriction \mathcal{L}(G_{n+1})"] \arrow[d, "\psi_{G_{n+1}}"'] & \mathcal{L}(H_{n+1}) \arrow[d, "\psi_{H_{n+1}}"] \\
\mathcal{L}(F_{n+1}) \arrow[r, "\pi_{n+1,x} \restriction \mathcal{L}(F_{n+1})"] & \mathcal{L}(F_{n+1})
\end{tikzcd}$$

I.e. for every $\ell \in V(G_{n+1}) \cap (R_{n+1} \cup B_{n+1})$ we have $\pi_{n+1,x}(\psi_{G_{n+1}}(\ell)) = \psi_{H_{n+1}}(h_{n+1,x}(\ell))$.

Proof. Since $R_{n+1}, B_{n+1} \subset G'_{n+1} \cup H'_{n+1}$, and $h_{n+1,x}$ extends $h'_{n+1,x}$ for each $x \in X_{n+1}$, this follows immediately from Lemma 8.5.8 after identifying $\hat{F}_{n+1}$ with F_{n+1}. □

This completes our recursive construction, and hence the proof of Theorem 8.1.2 is complete.

8.6 A non-reconstructible graph with countably many ends

In this section we will prove Theorem 8.1.3. Since the proof will follow almost exactly the same argument as the proof of Theorem 8.1.2, we will just indicate briefly here the parts which would need to be changed, and how the proof is structured.

The proof follows the same back and forth construction as in Section 8.5.2, however instead of starting with finite graphs G_0 and H_0 we will start with two infinite graphs, each containing one free end. For example we could start with the graphs in Figure 8.9.

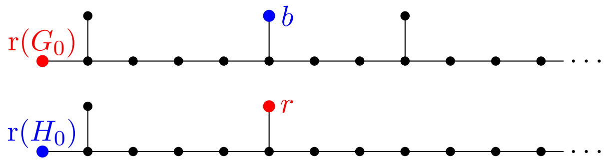

Figure 8.9: A possible choice for G_0 and H_0, where the dots indicate a ray.

The induction hypotheses remain the same, with the exception of (†7) and (†8) which are replaced by

(†7') G_n and H_n have exactly one limit end and infinitely many free ends when $n \geqslant 1$, and

(†8') $\overline{R_n \cup B_n} \cap \Omega(G_n \cup H_n) = \Omega'(G_n \cup H_n)$.

The arguments of Section 8.5.5 will then go through mutatis mutandis: for the proof of the analogue of Lemma 8.5.12, use Corollary 8.4.5 instead of Corollary 8.4.4.

To show that the construction then yields the desired non-reconstructible pair of graphs with countably many ends, we have to check that (†7') holds for the limit graphs G and H. It is clear that since $\overline{R_n \cup B_n} \cap \Omega(G_n \cup H_n) = \Omega'(G_n \cup H_n)$, every free end in a graph G_n or H_n remains free in the limit. Moreover, a similar argument to that in Section 8.5.3 shows that any pair of rays in G or H which were not in a free end in some G_n or H_n are equivalent in G or H, respectively.

However, since the end space of a locally finite connected graph is a compact metrizable space, and therefore has a countable dense subset, such a graph has at most countably many free ends, since they are isolated in $\Omega(G)$. Hence, both G and H have at most countably many free ends, and one limit end, and so both graphs have countably many ends.

Chapter 9

Topological ubiquity of trees

9.1 Introduction

Let $\lhd$ be a relation between graphs, for example the subgraph relation $\subseteq$, the topological minor relation $\leqslant$ or the minor relation $\preccurlyeq$. We say that a graph G is $\lhd$-*ubiquitous* if whenever Γ is a graph with $nG \lhd \Gamma$ for all $n \in \mathbb{N}$, then one also has $\aleph_0 G \lhd \Gamma$, where αG is the disjoint union of α many copies of G.

Two classic results of Halin [71, 72] say that both the ray and the double ray are $\subseteq$-ubiquitous, i.e. any graph which contains arbitrarily large collections of disjoint (double) rays must contain an infinite collection of disjoint (double) rays. However, even quite simple graphs can fail to be $\subseteq$ or $\leqslant$-ubiquitous, see e.g. [7, 130, 91], examples of which, due to Andreae [14], are depicted in Figures 9.1 and 9.2 below.

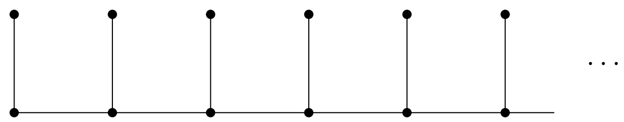

Figure 9.1: A graph which is not $\subseteq$-ubiquitous.

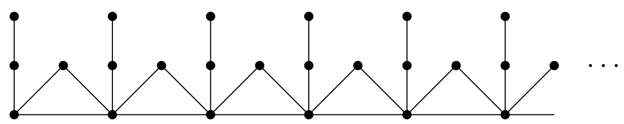

Figure 9.2: A graph which is not $\leqslant$-ubiquitous.

However, for the minor relation, no such simple examples of non-ubiquitous graphs are known. Indeed, one of the most important problems in the theory of infinite graphs is the so-called *Ubiquity Conjecture* due to Andreae [13].

Conjecture 9.1.1. *[The Ubiquity Conjecture] Every locally finite connected graph is $\preccurlyeq$-ubiquitous.*

In [13], Andreae constructed a graph that is not $\preccurlyeq$-ubiquitous. However, this construction relies on the existence of a counterexample to the well-quasi-ordering of infinite graphs under the minor relation, for which counterexamples are only known with very large cardinality [122]. In particular, it is still an open question whether or not there exists a countable connected graph which is not $\preccurlyeq$-ubiquitous.

In his most recent paper on ubiquity to date, Andreae [14] exhibited infinite families of locally finite graphs for which the ubiquity conjecture holds. The present paper is the first in a series of papers [25, 26, 27] making further progress towards the ubiquity conjecture, with the aim being to show that all graphs of bounded tree-width are ubiquitous.

As a first step towards this, we in particular need to deal with infinite trees, for which one even gets affirmative results regarding ubiquity under the topological minor relation. Halin showed in [73] that all trees of maximum degree 3 are $\leqslant$-ubiquitous. Andreae improved this result to show that all *locally finite* trees are $\leqslant$-ubiquitous [8], and asked if his result could be extended to arbitrary trees [8, p. 214]. Our main result of this paper answers this question in the affirmative.

Theorem 9.1.2. *Every tree is ubiquitous with respect to the topological minor relation.*

The proof will use some results about the well-quasi-ordering of trees under the topological minor relation of Nash-Williams [101] and Laver [93], as well as some notions about the topological structure of infinite graphs [49]. Interestingly, most of the work in proving Theorem 9.1.2 lies in dealing with the countable case, where several new ideas are needed. In fact, we will prove a slightly stronger statement in the countable case, which will allow us to derive the general result via transfinite induction on the cardinality of the tree, using some ideas from Shelah's singular compactness theorem [119].

To explain our strategy, let us fix some notation. When H is a subdivision of G we write $G \leqslant^* H$. Then, $G \leqslant \Gamma$ means that there is a subgraph $H \subseteq \Gamma$ which is a subdivision of G, that is, $G \leqslant^* H$. If H is a subdivision of G and v a vertex of G, then we denote by $H(v)$ the corresponding vertex in H. More generally, given a subgraph $G' \subseteq G$, we denote by $H(G')$ the corresponding subdivision of G' in H.

Now, suppose we have a rooted tree T and a graph Γ. Given a vertex $t \in T$, let T_t denote the subtree of T rooted in t. We say that a vertex $v \in \Gamma$ is *t-suitable* if there is some subdivision H of T_t in Γ with $H(t) = v$. For a subtree $S \subseteq T$ we say that a subdivision H of S in Γ is *T-suitable* if for each vertex $s \in V(S)$ the vertex $H(s)$ is s-suitable, i.e. for every $s \in V(S)$ there is a subdivision H' of T_s such that $H'(s) = H(s)$.

An *S-horde* is a sequence $(H_i \colon i \in \mathbb{N})$ of disjoint suitable subdivisions of S in Γ. If S' is a subtree of S, then we say that an S-horde $(H_i \colon i \in \mathbb{N})$ *extends* an S'-horde $(H'_i \colon i \in \mathbb{N})$ if for every $i \in \mathbb{N}$ we have $H_i(S') = H'_i$.

In order to show that an arbitrary tree T is $\leqslant$-ubiquitous, our rough strategy will be to build, by transfinite recursion, S-hordes for larger and larger subtrees S of T, each extending all the previous ones, until we have built a T-horde. However, to start the induction it will be necessary to show that we can build S-hordes for countable subtrees S of T. This will be done in the following key result of this paper:

Theorem 9.1.3. *Let T be a tree, S a countable subtree of T and Γ a graph such that $nT \leqslant \Gamma$ for every $n \in \mathbb{N}$. Then there is an S-horde in Γ.*

Note that Theorem 9.1.3 in particular implies $\leqslant$-ubiquity of countable trees.

We remark that whilst the relation $\preccurlyeq$ is a relaxation of the relation $\leqslant$, which is itself a relaxation of the relation $\subseteq$, it is not clear whether $\subseteq$-ubiquity implies $\leqslant$-ubiquity, or whether $\leqslant$-ubiquity implies $\preccurlyeq$-ubiquity. In the case of Theorem 9.1.2 however, it is true that arbitrary trees are also $\preccurlyeq$-ubiquitous, although the proof involves some extra technical difficulties that we will deal with in a later paper [27]. We note, however, that it is surprisingly easy to show that countable trees are $\preccurlyeq$-ubiquitous, since it can be derived relatively straightforwardly from Halin's grid theorem, see [25, Theorem 1.7].

This paper is structured as follows: In Section 9.2, we provide background on rooted trees, rooted topological embeddings of rooted trees (in the sense of Kruskal and Nash-Williams), and ends of graphs. In our graph theoretic notation we generally follow the textbook of Diestel

[43]. Next, Sections 9.3 to 9.5 introduce the key ingredients for our main ubiquity result. Section 9.3, extending ideas from Andreae's [8], lists three useful corollaries of Nash-Williams' and Laver's result that (labelled) trees are well-quasi-ordered under the topological minor relation, Section 9.4 investigates under which conditions a given family of disjoint rays can be rerouted onto another family of disjoint rays, and Section 9.5 shows that without loss of generality, we already have quite a lot of information about how exactly our copies of nG are placed in the host graph Γ.

Using these ingredients, we give a proof of the countable case, i.e. of Theorem 9.1.3, in Section 9.6. Finally, Section 9.7 contains the induction argument establishing our main result, Theorem 9.1.2.

9.2 Preliminaries

Definition 9.2.1. *A* rooted graph *is a pair (G, v) where G is a graph and $v \in V(G)$ is a vertex of G which we call the* root. *Often, when it is clear from the context which vertex is the root of the graph, we will refer to a rooted graph (G, v) as simply G.*

Given a rooted tree (T, v), we define a partial order $\leq$, which we call the tree-order*, on $V(T)$ by letting $x \leq y$ if the unique path between y and v in T passes through x. See [43, Section 1.5] for more background. For any edge $e \in E(T)$ we denote by e^- the endpoint closer to the root and by e^+ the endpoint further from the root. For any vertex t we denote by $N^+(t)$ the set of* children of t *in T, the neighbours s of t satisfying $t \leq s$. The subtree of T rooted at t is denoted by (T_t, t), that is, the induced subgraph of T on the set of vertices $\{s \in V(T) : t \leq s\}$.*

We say that a rooted tree (S, w) is a rooted subtree *of a rooted tree (T, v) if S is a subgraph of T such that the tree order on (S, w) agrees with the induced tree order from (T, v). In this case we write $(S, w) \subseteq_r (T, v)$.*

We say that a rooted tree (S, w) is a rooted topological minor *of a rooted tree (T, v) if there is a subgraph S' of T which is a subdivision of S such that for any $x \leq y \in V(S)$, $S'(x) \leq S'(y)$ in the tree-order on T. We call such an S' a* rooted subdivision of S. *In this case we write $(S, w) \leqslant_r (T, v)$, cf. [43, Section 12.2].*

Definition 9.2.2 (Ends of a graph, cf. [43, Chapter 8])**.** *An* end *in an infinite graph Γ is an equivalence class of rays, where two rays R and S are equivalent if and only if there are infinitely many vertex disjoint paths between R and S in Γ. We denote by $\Omega(\Gamma)$ the set of ends in Γ. Given any end $\epsilon \in \Omega(\Gamma)$ and a finite set $X \subseteq V(\Gamma)$ there is a unique component of $\Gamma - X$ which contains a tail of every ray in ϵ, which we denote by $C(X, \epsilon)$.*

A vertex $v \in V(\Gamma)$ dominates *an end ω if there is a ray $R \in \omega$ such that there are infinitely many vertex disjoint v–R-paths in Γ.*

Definition 9.2.3. *For a path or ray P and vertices $v, w \in V(P)$, let vPw denote the subpath of P with endvertices v and w. If P is a ray, let Pv denote the finite subpath of P between the initial vertex of P and v, and let vP denote the subray (or* tail*) of P with initial vertex v.*

Given two paths or rays P and Q which are disjoint but for one of their endvertices, we write PQ for the concatenation of P and Q*, that is the path, ray or double ray $P \cup Q$. Since concatenation of paths is associative, we will not use parentheses. Moreover, if we concatenate paths of the form vPw and wQx, then we omit writing w twice and denote the concatenation by $vPwQx$.*

9.3 Well-quasi-orders and κ-embeddability

Definition 9.3.1. *Let X be a set and let $\triangleleft$ be a binary relation on X. Given an infinite cardinal κ we say that an element $x \in X$ is* κ-embeddable (with respect to $\triangleleft$) in X *if there are at least κ many elements $x' \in X$ such that $x \triangleleft x'$.*

Definition 9.3.2 (well-quasi-order)**.** *A binary relation $\triangleleft$ on a set X is a* well-quasi-order *if it is reflexive and transitive, and for every sequence $x_1, x_2, \ldots \in X$ there is some $i < j$ such that $x_i \triangleleft x_j$.*

Lemma 9.3.3. *Let X be a set and let $\triangleleft$ be a well-quasi-order on X. For any infinite cardinal κ the number of elements of X which are not κ-embeddable with respect to $\triangleleft$ in X is less than κ.*

Proof. For $x \in X$ let $U_x = \{y \in X : x \triangleleft y\}$. Now suppose for a contradiction that the set $A \subseteq X$ of elements which are not κ-embeddable with respect to $\triangleleft$ in X has size at least κ. Then, we can recursively pick a sequence $(x_n \in A)_{n\in\mathbb{N}}$ such that $x_m \ntriangleleft x_n$ for $m < n$. Indeed, having chosen all x_m with $m < n$ it suffices to choose x_n to be any element of the set $A \setminus \bigcup_{m<n} U_{x_m}$, which is nonempty since A has size κ but each U_{x_m} has size $< \kappa$.

By construction we have $x_m \ntriangleleft x_n$ for $m < n$, contradicting the assumption that $\triangleleft$ is a well-quasi-order on X. □

We will use the following theorem of Nash-Williams on well-quasi-ordering of rooted trees, and its extension by Laver to labelled rooted trees.

Theorem 9.3.4 (Nash-Williams [101])**.** *The relation $\leqslant_r$ is a well-quasi order on the set of rooted trees.*

Theorem 9.3.5 (Laver [93])**.** *The relation $\leqslant_r$ is a well-quasi order on the set of rooted trees with finitely many labels, i.e. for every finite number $k \in \mathbb{N}$, whenever $(T_1, c_1), (T_2, c_2), \ldots$ is a sequence of rooted trees with k-colourings $c_i \colon T_i \to [k]$, there is some $i < j$ such that there exists a subdivision H of T_i with $H \subseteq_r T_j$ and $c_i(t) = c_j(H(t))$ for all $t \in T_i$.*[1]

Together with Lemma 9.3.3 these results give us the following three corollaries:

Definition 9.3.6. *Let (T, v) be an infinite rooted tree. For any vertex t of T and any infinite cardinal κ, we say that a child t' of t is κ-embeddable if there are at least κ children t'' of t such that $T_{t'}$ is a rooted topological minor of $T_{t''}$.*

Corollary 9.3.7. *Let (T, v) be an infinite rooted tree, $t \in V(T)$ and $\mathcal{T} = \{T_{t'} : t' \in N^+(t)\}$. Then for any infinite cardinal κ, the number of children of t which are not κ-embeddable is less than κ.*

Proof. By Theorem 9.3.4 the set $\mathcal{T} = \{T_{t'} : t' \in N^+(t)\}$ is well-quasi-ordered by $\leqslant_r$ and so the claim follows by Lemma 9.3.3 applied to $\mathcal{T}$, $\leqslant_r$, and κ. □

Corollary 9.3.8. *Let (T, v) be an infinite rooted tree, $t \in V(T)$ a vertex of infinite degree and $(t_i \in N^+(t) : i \in \mathbb{N})$ a sequence of countably many of its children. Then there exists $N_t \in \mathbb{N}$ such that for all $n \geqslant N_t$,*

$$\{t\} \cup \bigcup_{i>N_t} T_{t_i} \leqslant_r \{t\} \cup \bigcup_{i>n} T_{t_i}$$

(considered as trees rooted at t) fixing the root t.

[1] In fact, Laver showed that rooted trees labelled by a *better-quasi-order* are again better-quasi-ordered under $\leqslant_r$ respecting the labelling, but we shall not need this stronger result.

Proof. Consider a labelling $c\colon T_t \to [2]$ mapping t to 1, and all remaining vertices of T_t to 2. By Theorem 9.3.5, the set $\mathcal{T} = \{\{t\} \cup \bigcup_{i>n} T_{t_i} : n \in \mathbb{N}\}$ is well-quasi-ordered by $\leqslant_r$ respecting the labelling, and so the claim follows by applying Lemma 9.3.3 to $\mathcal{T}$ and $\leqslant_r$ with $\kappa = \aleph_0$. $\square$

Definition 9.3.9 (Self-similarity)**.** *A ray $R = r_1r_2r_3\ldots$ in a rooted tree (T, v) which is upwards with respect to the tree order* displays self-similarity of T *if there are infinitely many n such that there exists a subdivision H of T_{r_0} with $H \subseteq_r T_{r_n}$ and $H(R) \subseteq R$.*

Corollary 9.3.10. *Let (T, v) be an infinite rooted tree and let $R = r_1r_2r_3\ldots$ be a ray which is upwards with respect to the tree order. Then there is a $k \in \mathbb{N}$ such that r_kR displays self-similarity of T.*[2]

Proof. Consider a labelling $c\colon T \to [2]$ mapping the vertices on the ray R to 1, and labelling all remaining vertices of T with 2. By Theorem 9.3.5, the set $\mathcal{T} = \{(T_{r_i}, c_i) \colon i \in \mathbb{N}\}$, where c_i is the natural restriction of c to T_{r_i}, is well-quasi-ordered by $\leqslant_r$ respecting the labellings. Hence by Lemma 9.3.3, the number of indices i such that T_{r_i} is not $\aleph_0$-embeddable in $\mathcal{T}$ is finite. Let k be larger than any such i. Then, since T_{r_k} is $\aleph_0$-embeddable in $\mathcal{T}$, there are infinitely many $r_j \in r_kR$ such that $T_{r_k} \leqslant_r T_{r_j}$ respecting the labelling, i.e. mapping the ray to the ray, and hence r_kR displays the self similarity of T. $\square$

9.4 Linkages between rays

In this section we will establish a toolkit for constructing a disjoint system of paths from one family of disjoint rays to another.

Definition 9.4.1 (Tail of a ray)**.** *Given a ray R in a graph Γ and a finite set $X \subseteq V(\Gamma)$ the* tail of R after X*, denoted by $T(R, X)$, is the unique infinite component of R in $\Gamma - X$.*

Definition 9.4.2 (Linkage of families of rays)**.** *Let $\mathcal{R} = (R_i \colon i \in I)$ and $\mathcal{S} = (S_j \colon j \in J)$ be families of vertex disjoint rays, where the initial vertex of each R_i is denoted x_i. A family of paths $\mathcal{P} = (P_i \colon i \in I)$, is a* linkage *from $\mathcal{R}$ to $\mathcal{S}$ if there is an injective function $\sigma\colon I \to J$ such that*

- *each P_i joins a vertex $x'_i \in R_i$ to a vertex $y_{\sigma(i)} \in S_{\sigma(i)}$;*
- *the family $\mathcal{T} = (x_iR_ix'_iP_iy_{\sigma(i)}S_{\sigma(i)} \colon i \in I)$ is a collection of disjoint rays.*

We say that $\mathcal{T}$ is obtained by transitioning *from $\mathcal{R}$ to $\mathcal{S}$ along the linkage $\mathcal{P}$. Given a finite set of vertices $X \subseteq V(\Gamma)$, we say that $\mathcal{P}$ is* after X *if $x'_i \in T(R_i, X)$ and $x'_iP_iy_{\sigma(i)}S_{\sigma(i)}$ avoids X for all $i \in I$.*

Lemma 9.4.3 (Weak linking lemma)**.** *Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. Then for any families $\mathcal{R} = (R_i \colon i \in [n])$ and $\mathcal{S} = (S_j \colon j \in [n])$ of vertex disjoint rays in ϵ and any finite set X of vertices, there is a linkage from $\mathcal{R}$ to $\mathcal{S}$ after X.*

Proof. Let us write x_i for the initial vertex of each R_i and let x'_i be the initial vertex of the tail $T(R_i, X)$. Furthermore, let $X' = X \cup \bigcup_{i\in[n]} R_ix'_i$. For $i \in [n]$ we will construct inductively finite disjoint connected subgraphs $K_i \subseteq \Gamma$ for each $i \in [n]$ such that

- K_i meets $T(S_j, X')$ and $T(R_j, X')$ for every $j \in [n]$;
- K_i avoids X'.

[2] A slightly weaker statement, without the additional condition that $H(R) \subseteq R$ appeared in [8, Lemma 1].

Suppose that we have constructed $K_1, \ldots, K_{m-1}$ for some $m \leq n$. Let us write $X_m = X' \cup \bigcup_{i<m} V(K_i)$. Since $R_1, \ldots, R_n$ and $S_1, \ldots, S_n$ lie in the same end ϵ, there exist paths $Q_{i,j}$ between $T(R_i, X_m)$ and $T(S_j, X_m)$ avoiding X_m for all $i \neq j \in [n]$. Let $K_m = F \cup \bigcup_{i \neq j \in [n]} Q_{i,j}$, where F consists of an initial segment of each $T(R_i, X_m)$ sufficiently large to make K_m connected. Then it is clear that K_m is disjoint from all previous K_i and satisfies the claimed properties.

Let $K = \bigcup_{i=1}^n K_i$ and for each $j \in [n]$, let y_j be the initial vertex of $T(S_j, V(K))$. Note that by construction $T(S_j, V(K))$ avoids X for each j, since K_1 meets $T(S_j, X)$ and so $T(S_j, V(K)) \subseteq T(S_j, X)$.

We claim that there is no separator of size $< n$ between $\{x'_1, \ldots, x'_n\}$ and $\{y_1, \ldots, y_n\}$ in the subgraph $\Gamma' \subseteq \Gamma$ where $\Gamma' = K \cup \bigcup_{j=1}^n T(R_j, X') \cup T(S_j, X')$. Indeed, any set of $< n$ vertices must avoid at least one ray R_i, at least one graph K_m and one ray S_j. However, since K_m is connected and meets R_i and S_j, the separator does not separate x'_i from y_j.

Hence, by a version of Menger's theorem for infinite graphs [43, Proposition 8.4.1], there is a collection of n disjoint paths P_i from x'_i to $y_{\sigma(i)}$ in Γ'. Since Γ' is disjoint from X and meets each $R_i x'_i$ in x'_i only, it is clear that $\mathcal{P} = (P_i \colon i \in [n])$ is as desired. □

In some cases we will need to find linkages between families of rays which avoid more than just a finite subset X. For this we will use the following lemma, which is stated in slightly more generality than needed in this paper. Broadly the idea is that if we have a family of disjoint rays $(R_i \colon i \in [n])$ tending to an end ϵ and a number $a \in \mathbb{N}$, then there is some fixed number $N = N(a, n)$ such that if we have N disjoint graphs H_i, each with a specified ray S_i tending to ϵ, then we can 're-route' the rays $(R_i \colon i \in [n])$ to some of the rays $(S_j \colon j \in [N])$, in such a way that we totally avoid a of the graphs H_i.

Lemma 9.4.4 (Strong linking lemma)**.** *Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. Let X be a finite set of vertices, $a, n \in \mathbb{N}$, and $\mathcal{R} = (R_i \colon i \in [n])$ a family of vertex disjoint rays in ϵ. Let x_i be the initial vertex of R_i and let x'_i the initial vertex of the tail $T(R_i, X)$.*

Then there is a finite number $N = N(\mathcal{R}, X, a)$ with the following property: For every collection $(H_j \colon j \in [N])$ of vertex disjoint subgraphs of Γ, all disjoint from X and each including a specified ray S_j in ϵ, there is a set $A \subseteq [N]$ of size a and a linkage $\mathcal{P} = (P_i \colon i \in [n])$ from $\mathcal{R}$ to $(S_j \colon j \in [N])$ which is after X and such that the family

$$\mathcal{T} = \big(x_i R_i x'_i P_i y_{\sigma(i)} S_{\sigma(i)} \colon i \in [n]\big)$$

avoids $\bigcup_{k \in A} H_k$.

Proof. Let $X' = X \cup \bigcup_{i \in [n]} R_i x'_i$ and let $N_0 = |X'|$. We claim that the lemma holds with $N = N_0 + n^3 + a$.

Indeed suppose that $(H_j \colon j \in [N])$ is a collection of vertex disjoint subgraphs as in the statement of the lemma. Since the H_j are vertex disjoint, we may assume without loss of generality that the family $(H_j \colon j \in [n^3 + a])$ is disjoint from X'.

For each $i \in [n^2]$ we will build inductively finite, connected, vertex disjoint subgraphs $\hat{K}_i$ such that

- $\hat{K}_i$ contains $x'_{i \pmod{n}}$,
- $\hat{K}_i$ meets exactly n of the H_j, that is $|\{ j \in [n^3 + a] \,:\, \hat{K}_i \cap H_j \neq \emptyset\}| = n$, and
- $\hat{K}_i$ avoids X'.

Suppose we have done so for all $i < m$. Let $X_m = X' \cup \bigcup_{i<m} V(\hat{K}_i)$. We will build inductively for $t = 0, \ldots, n$ increasing connected subgraphs $\hat{K}^t_m$ that meet $R_{i \pmod{n}}$, meet exactly t of the H_j, and avoid X_m.

We start with $\hat{K}_m^0 = \emptyset$. For each $t = 0, \dots n-1$, if $T(R_{m \pmod n}, X_m)$ meets some H_j not met by $\hat{K}_m^t$ then there is some initial vertex $z_t \in T(R_{m \pmod n}, X_m)$ where it does so and we set $\hat{K}_m^{t+1} := \hat{K}_m^t \cup T(R_{m \pmod n}, X_m)z_t$. Otherwise we may assume $T(R_{m \pmod n}, X_m)$ does not meet any such H_j. In this case, let $j \in [n^3+a]$ be such that $\hat{K}_m^t \cap H_j = \emptyset$. Since $R_{m \pmod n}$ and S_j belong to the same end ϵ, there is some path P between $T(R_{m \pmod n}, X_m)$ and $T(S_j, X_m)$ which avoids X_m. Since this path meets some H_k with $k \in [n^3+a]$ which $\hat{K}_m^t$ does not, there is some initial segment P' which meets exactly one such H_k. To form $\hat{K}_m^{t+1}$ we add this path to $\hat{K}_m^t$ together with an appropriately large initial segment of $T(R_{m \pmod n}, X_m)$ such that $\hat{K}_m^{t+1}$ is connected and contains $x'_{m \pmod n}$. Finally we let $\hat{K}_m = \hat{K}_m^n$.

Let $K = \bigcup_{i \in [n^2]} \hat{K}_i$. Since each $\hat{K}_i$ meets exactly n of the H_j, the set

$$J = \{j \in [n^3+a] \, : \, H_j \cap K \neq \emptyset\}$$

satisfies $|J| \leqslant n^3$. For each $j \in J$ let y_j be the initial vertex of $T(S_j, V(K))$.

We claim that there is no separator of size $< n$ between $\{x'_1, \dots x'_n\}$ and $\{y_j \, : \, j \in J\}$ in the subgraph $\Gamma' \subseteq \Gamma$ where $\Gamma' = K \cup \bigcup_{j \in [n]} T(R_j, X') \cup \bigcup_{j \in J} H_j$. Suppose for a contradiction that there is such a separator S. Then S cannot meet every R_i, and hence avoids some R_q. Furthermore, there are n distinct $\hat{K}_i$ such that $i = q \pmod n$, all of which are disjoint. Hence there is some $\hat{K}_r$ with $r = q \pmod n$ disjoint from S. Finally, $|\{j \in J \, : \, \hat{K}_r \cap H_j \neq \emptyset\}| = n$ and so there is some H_s disjoint from S such that $\hat{K}_r \cap H_s \neq \emptyset$. Since $\hat{K}_r$ meets $T(R_q, X')$ and H_s, there is a path from x'_q to y_s in Γ', contradicting our assumption.

Hence, by a version of Menger's theorem for infinite graphs [43, Proposition 8.4.1], there is a family of disjoint paths $\mathcal{P} = (P_i \colon i \in [n])$ in Γ' from x'_i to $y_{\sigma(i)}$. Furthermore, since $|J| \leqslant n^3$ there is some subset $A \subseteq [n^3+a]$ of size a such that H_k is disjoint from K for each $k \in A$.

Therefore, since Γ' is disjoint from X' and meets each $R_i x'_i$ in x'_i only, the family $\mathcal{P}$ is a linkage from $\mathcal{R}$ to $(S_j)_{j \in [n^3+a]}$ which is after X such that

$$\mathcal{T} = \big(x_i R_i x'_i P_i y_{\sigma(i)} S_{\sigma(i)} \colon i \in [n]\big)$$

avoids $\bigcup_{k \in A} H_k$. □

We will also need the following result, which allows us to work with paths instead of rays if the end ϵ is dominated by infinitely many vertices.

Lemma 9.4.5. *Let Γ be a graph and ϵ an end of Γ which is dominated by infinitely many vertices. Let $x_1, x_2, \dots, x_k$ be distinct vertices. If there are disjoint rays from the x_i to ϵ then there are disjoint paths from the x_i to distinct vertices y_i which dominate ϵ.*

Proof. We argue by induction on k. The base case $k = 0$ is trivial, so let us assume $k > 0$.

Consider any family of disjoint rays R_i, each from x_i to ϵ. Let y_k be any vertex dominating ϵ. Let P be a y_k–$\bigcup_{i=1}^k R_i$-path. Without loss of generality the endvertex u of P in $\bigcup_{i=1}^k R_i$ lies on R_k. Then by the induction hypothesis applied to the graph $\Gamma - R_k u P$ we can find disjoint paths in that graph from the x_i with $i < k$ to vertices y_i which dominate ϵ. These paths together with $R_k u P$ then form the desired collection of paths. □

To go back from paths to rays we will use the following lemma.

Lemma 9.4.6. *Let Γ be a graph and ϵ an end of Γ which is dominated by infinitely many vertices. Let $y_1, y_2, \dots, y_k$ be vertices, not necessarily distinct, dominating Γ. Then there are rays R_i from the respective y_i to ϵ which are disjoint except at their initial vertices.*

Proof. We recursively build for each $n \in \mathbb{N}$ paths $P_1^n, \ldots, P_k^n$, each P_i^n from y_i to a vertex y_i^n dominating ϵ, disjoint except at their initial vertices, such that for $m < n$ each P_i^n properly extends P_i^m. We take P_i^0 to be a trivial path. For $n > 0$, build the P_i^n recursively in i: To construct P_i^n, we start by taking X_i^n to be the finite set of all the vertices of the P_j^n with $j < i$ or P_j^{n-1} with $j \geqslant i$. We then choose a vertex y_i^n outside of X_i^n which dominates ϵ and a path Q_i^n from y_i^{n-1} to y_i^n internally disjoint from X_i^n. Finally we let $P_i^n := P_i^{n-1} y_{n-1} Q_i^n$.

Finally, for each $i \leqslant k$, we let R_i be the ray $\bigcup_{n \in \mathbb{N}} P_i^n$. Then the R_i are disjoint except at their initial vertices, and they are in ϵ, since each of them contains infinitely many dominating vertices of ϵ. □

9.5 G-tribes and concentration of G-tribes towards an end

For showing that a given graph G is ubiquitous with respect to a fixed relation $\lhd$, we shall assume that $nG \lhd \Gamma$ for every $n \in \mathbb{N}$ and need to show that this implies that $\aleph_0 G \lhd \Gamma$. Since each subgraph witnessing that $nG \lhd \Gamma$ will be a collection of n disjoint subgraphs each being a witness for $G \lhd \Gamma$, it will be useful to introduce some notation for talking about these families of collections of n disjoint witnesses for each n.

To do this formally, we need to distinguish between a relation like the topological minor relation and the subdivision relation. Recall that we write $G \leqslant^* H$ if H is a subdivision of G and $G \leqslant \Gamma$ if G is a topological minor of Γ. We can interpret the topological minor relation as the composition of the subdivision relation and the subgraph relation.

Given two relations R and S, let their *composition* $S \circ R$ be the relation defined by $x(S \circ R)z$ if and only if there is a y such that xRy and ySz.

Hence we have that $G\,(\subseteq \circ \leqslant^*)\,\Gamma$ if and only if there exists H such that $G \leqslant^* H \subseteq \Gamma$, that is, if and only if $G \leqslant \Gamma$.

While in this paper we will only work with the topological minor relation, we will state the following definition and lemmas in greater generality, so that we may apply them in later papers in this series [25, 26, 27].

In general, we want to consider a pair $(\lhd, \blacktriangleleft)$ of binary relations of graphs with the following properties.

(R1) $\lhd \; = \; (\subseteq \circ \blacktriangleleft)$;

(R2) Given a set I and a family $(H_i : i \in I)$ of pairwise disjoint graphs with $G \blacktriangleleft H_i$ for all $i \in I$, then $|I| \cdot G \blacktriangleleft \bigcup \{H_i : i \in I\}$.

We call a pair $(\lhd, \blacktriangleleft)$ with these properties *compatible*.

Other examples of compatible pairs are $(\subseteq, \cong)$, where $\cong$ denotes the isomorphism relation, as well as $(\preccurlyeq, \preccurlyeq^*)$, where $G \preccurlyeq^* H$ if H is an inflated copy of G.

Definition 9.5.1 (G-tribes)**.** *Let G and Γ be graphs, and let $(\lhd, \blacktriangleleft)$ be a compatible pair of relations between graphs.*

- *A G-*tribe in Γ (with respect to $(\lhd, \blacktriangleleft)$) *is a collection $\mathcal{F}$ of finite sets F of disjoint subgraphs H of Γ such that $G \blacktriangleleft H$ for each* member of $\mathcal{F}$ *$H \in \bigcup \mathcal{F}$.*
- *A G-tribe $\mathcal{F}$ in Γ is called* thick*, if for each $n \in \mathbb{N}$ there is a* layer *$F \in \mathcal{F}$ with $|F| \geqslant n$; otherwise, it is called* thin.[3]

[3] A similar notion of *thick* and *thin families* was also introduced by Andreae in [8] (in German) and in [14]. The remaining notions, and in particular the concept of a *concentrated G-tribe*, which will be the backbone of essentially all our results in this series of papers, is new.

- *A G-tribe $\mathcal{F}'$ in Γ is a G-*subtribe *of a G-tribe $\mathcal{F}$ in Γ, denoted by $\mathcal{F}' \lhd \mathcal{F}$, if there is an injection $\Psi\colon \mathcal{F}' \to \mathcal{F}$ such that for each $F' \in \mathcal{F}'$ there is an injection $\varphi_{F'}\colon F' \to \Psi(F')$ such that $V(H') \subseteq V(\varphi_{F'}(H'))$ for each $H' \in F'$. The G-subtribe $\mathcal{F}'$ is called* flat, *denoted by $\mathcal{F}' \subseteq \mathcal{F}$, if there is such an injection Ψ satisfying $F' \subseteq \Psi(F')$.*

- *A thick G-tribe $\mathcal{F}$ in Γ is* concentrated at an end ϵ *of Γ, if for every finite vertex set X of Γ, the G-tribe $\mathcal{F}_X = \{F_X \colon F \in \mathcal{F}\}$ consisting of the layers $F_X = \{H \in F : H \not\subseteq C(X, \epsilon)\} \subseteq F$ is a thin subtribe of $\mathcal{F}$.*

Hence, for a given compatible pair $(\lhd, \blacktriangleleft)$, if we wish to show that G is $\lhd$-ubiquitous, we will need to show that the existence of a thick G-tribe in Γ with respect to $(\lhd, \blacktriangleleft)$ implies $\aleph_0 G \lhd \Gamma$. We first observe that removing a thin G-tribe from a thick G-tribe always leaves a thick G-tribe.

Lemma 9.5.2 (cf. [8, Lemma 3] or [14, Lemma 2])**.** *Let $\mathcal{F}$ be a thick G-tribe in Γ and let $\mathcal{F}'$ be a thin subtribe of $\mathcal{F}$, witnessed by $\Psi\colon \mathcal{F}' \to \mathcal{F}$ and $(\varphi_{F'}\colon F' \in \mathcal{F}')$. For $F \in \mathcal{F}$, if $F \in \Psi(\mathcal{F}')$, let $\Psi^{-1}(F) = \{F'_F\}$ and set $\hat{F} = \varphi_{F'_F}(F'_F)$. If $F \notin \Psi(\mathcal{F}')$, set $\hat{F} = \emptyset$. Then*

$$\mathcal{F}'' := \{F \setminus \hat{F}\colon F \in \mathcal{F}\}$$

is a thick flat G-subtribe of $\mathcal{F}$.

Proof. $\mathcal{F}''$ is obviously a flat subtribe of $\mathcal{F}$. As $\mathcal{F}'$ is thin, there is a $k \in \mathbb{N}$ such that $|F'| \leq k$ for every $F' \in \mathcal{F}'$. Thus $|\hat{F}| \leq k$ for all $F \in \mathcal{F}$. Let $n \in \mathbb{N}$. As $\mathcal{F}$ is thick, there is a layer $F \in \mathcal{F}$ satisfying $|F| \geq n + k$. Thus $|F \setminus \hat{F}| \geq n + k - k = n$. □

Given a thick G-tribe, the members of this tribe may have different properties, for example, some of them contain a ray belonging to a specific end ϵ of Γ whereas some of them do not. The next lemma allows us to restrict onto a thick subtribe, in which all members have the same properties, as long as we consider only finitely many properties. E.g. we find a subtribe in which either all members contain an ϵ-ray, or none of them contain such a ray.

Lemma 9.5.3 (Pigeon hole principle for thick G-tribes)**.** *Suppose for some $k \in \mathbb{N}$, we have a k-colouring $c\colon \bigcup \mathcal{F} \to [k]$ of the members of some thick G-tribe $\mathcal{F}$ in Γ. Then there is a monochromatic, thick, flat G-subtribe $\mathcal{F}'$ of $\mathcal{F}$.*

Proof. Since $\mathcal{F}$ is a thick G-tribe, there is a sequence $(n_i\colon i \in \mathbb{N})$ of natural numbers and a sequence $(F_i \in \mathcal{F}\colon i \in \mathbb{N})$ such that

$$n_1 \leqslant |F_1| < n_2 \leqslant |F_2| < n_3 \leqslant |F_3| < \cdots .$$

Now for each i, by pigeon hole principle, there is one colour $c_i \in [k]$ such that the subset $F'_i \subseteq F_i$ of elements of colour c_i has size at least n_i/k. Moreover, since $[k]$ is finite, there is one colour $c^* \in [k]$ and an infinite subset $I \subseteq \mathbb{N}$ such that $c_i = c^*$ for all $i \in I$. But this means that $\mathcal{F}' := \{F'_i\colon i \in I\}$ is a monochromatic, thick, flat G-subtribe. □

In this series of papers we will be interested in graph relations such as $\subseteq$, $\leqslant$ and $\preccurlyeq$. Given a connected graph G and a compatible pair of relations $(\lhd, \blacktriangleleft)$ we say that a G-tribe $\mathcal{F}$ w.r.t $(\lhd, \blacktriangleleft)$ is *connected* if every member H of $\mathcal{F}$ is connected. Note that for relations $\blacktriangleleft$ like $\cong, \leq^*, \preccurlyeq^*$, if G is connected and $G \blacktriangleleft H$, then H is connected. In this case, any G-tribe will be connected.

Lemma 9.5.4. *Let G be a connected graph (of arbitrary cardinality), $(\lhd, \blacktriangleleft)$ a compatible pair of relations of graphs and Γ a graph containing a thick connected G-tribe $\mathcal{F}$ w.r.t. $(\lhd, \blacktriangleleft)$. Then either $\aleph_0 G \lhd \Gamma$, or there is a thick flat subtribe $\mathcal{F}'$ of $\mathcal{F}$ and an end ϵ of Γ such that $\mathcal{F}'$ is concentrated at ϵ.*

Proof. For every finite vertex set $X \subseteq V(\Gamma)$, only a thin subtribe of $\mathcal{F}$ can meet X, so by Lemma 9.5.2 a thick flat subtribe $\mathcal{F}''$ is contained in the graph $\Gamma - X$. Since each member of $\mathcal{F}''$ is connected, any member H of $\mathcal{F}''$ is contained in a unique component of $\Gamma - X$. If for any X, infinitely many components of $\Gamma - X$ contain a $\blacktriangleleft$-copy of G, the union of all these copies is a $\blacktriangleleft$-copy of $\aleph_0 G$ in Γ by (R2), hence $\aleph_0 G \lhd \Gamma$. Thus, we may assume that for each X, only finitely many components contain elements from $\mathcal{F}''$, and hence, by colouring each H with a colour corresponding to the component of $\Gamma - X$ containing it, we may assume by the pigeon hole principle for G-tribes, Lemma 9.5.3, that at least one component of $\Gamma - X$ contains a thick flat subtribe of $\mathcal{F}$.

Let $C_0 = \Gamma$ and $\mathcal{F}_0 = \mathcal{F}$ and consider the following recursive process: If possible, we choose a finite vertex set X_n in C_n such that there are two components $C_{n+1} \neq D_{n+1}$ of $C_n - X_n$ where C_{n+1} contains a thick flat subtribe $\mathcal{F}_{n+1} \subseteq \mathcal{F}_n$ and D_{n+1} contains at least one $\blacktriangleleft$-copy H_{n+1} of G. Since by construction all H_n are pairwise disjoint, we either find infinitely many such H_n and thus, again by (R2), an $\aleph_0 G \lhd \Gamma$, or our process terminates at step N say. That is, we have a thick flat subtribe $\mathcal{F}_N$ contained in a subgraph C_N such that there is no finite vertex set X_N satisfying the above conditions.

Let $\mathcal{F}' := \mathcal{F}_N$. We claim that for every finite vertex set X of Γ, there is a unique component C_X of $\Gamma - X$ that contains a thick flat G-subtribe of $\mathcal{F}'$. Indeed, note that if for some finite $X \subseteq \Gamma$ there are two components C and C' of $\Gamma - X$ both containing thick flat G-subtribes of $\mathcal{F}'$, then since every G-copy in $\mathcal{F}'$ is contained in C_N, it must be the case that $C \cap C_N \neq \emptyset \neq C' \cap C_N$. But then $X_N = X \cap C_N \neq \emptyset$ is a witness that our process could not have terminated at step N.

Next, observe that whenever $X' \supseteq X$, then $C_{X'} \subseteq C_X$. By a theorem of Diestel and Kühn, [49], it follows that there is a unique end ϵ of Γ such that $C(X, \epsilon) = C_X$ for all finite $X \subseteq \Gamma$. It now follows easily from the uniqueness of $C_X = C(X, \epsilon)$ that $\mathcal{F}'$ is concentrated at this ϵ. □

We note that concentration towards an end ϵ is a robust property in the following sense:

Lemma 9.5.5. *Let G be a connected graph (of arbitrary cardinality), $(\lhd, \blacktriangleleft)$ a compatible pair of relations of graphs and Γ a graph containing a thick connected G-tribe $\mathcal{F}$ w.r.t. $(\lhd, \blacktriangleleft)$ concentrated at an end ϵ of Γ. Then the following assertions hold:*

1. *For every finite set X, the component $C(X, \epsilon)$ contains a thick flat G-subtribe of $\mathcal{F}$.*

2. *Every thick subtribe $\mathcal{F}'$ of $\mathcal{F}$ is concentrated at ϵ, too.*

Proof. Let X be a finite vertex set. By definition, if the G-tribe $\mathcal{F}$ is concentrated at ϵ, then $\mathcal{F}$ is thick, and the subtribe $\mathcal{F}_X$ consisting of the sets $F_X = \{H \in F : H \not\subseteq C(X, \epsilon)\} \subseteq F$ for $F \in \mathcal{F}$ is a thin subtribe of $\mathcal{F}$, i.e. there exists $k \in \mathbb{N}$ such that $|F_X| \leqslant k$ for all $F_X \in \mathcal{F}_X$.

For (1), observe that the G-tribe $\mathcal{F}' = \{F \setminus F_X \colon F \in \mathcal{F}\}$ is a thick flat subtribe of $\mathcal{F}$ by Lemma 9.5.2, and all its members are contained in $C(X, \epsilon)$ by construction.

For (2), observe that if $\mathcal{F}'$ is a subtribe of $\mathcal{F}$, then for every $F' \in \mathcal{F}'$ there is an injection $\varphi_{F'} \colon F' \to F$ for some $F \in \mathcal{F}$. Therefore, $|\varphi_{F'}^{-1}(F_X)| \leqslant k$ for $F_X \subseteq F$ as defined above, and so only a thin subtribe of $\mathcal{F}'$ is not contained in $C(X, \epsilon)$. □

9.6 Countable subtrees

In this section we prove Theorem 9.1.3. Let S be a countable subtree of T. Our aim is to construct an *S-horde* $(Q_i \colon i \in \mathbb{N})$ of disjoint suitable subdivisions of S in Γ inductively. By Lemma 9.5.4, we may assume without loss of generality that there are an end ϵ of Γ and a thick T-tribe $\mathcal{F}$ concentrated at ϵ.

In order to ensure that we can continue the construction at each stage, we will require the existence of additional structure for each n. But the details of what additional structure we use will vary depending on how many vertices dominate ϵ. So, after a common step of preprocessing, in Section 9.6.1, the proof of Theorem 9.1.3 splits into two cases according to whether the number of ϵ-dominating vertices in Γ is finite (Section 9.6.2) or infinite (Section 9.6.3).

9.6.1 Preprocessing

We begin by picking a root v for S, and also consider T as a rooted tree with root v. Let $V_\infty(S)$ be the set of vertices of infinite degree in S.

Definition 9.6.1. *Given S and T as above, define a spanning locally finite forest $S^* \subseteq S$ by*

$$S^* := S \setminus \bigcup_{t \in V_\infty(S)} \{tt_i \colon t_i \in N^+(t), i > N_t\},$$

where N_t is as in Corollary 9.3.8. We will also consider every component of S^* as a rooted tree given by the induced tree order from T.

Definition 9.6.2. *An edge e of S^* is an* extension edge *if there is a ray in S^* starting at e^+ which displays self-similarity of T. For each extension edge e we fix one such a ray R_e. Write $Ext(S^*) \subseteq E(S^*)$ for the set of extension edges.*

Consider the forest $S^* - Ext(S^*)$ obtained from S^* by removing all extension edges. Since every ray in S^* must contain an extension edge by Corollary 9.3.10, each component of $S^* - Ext(S^*)$ is a locally finite rayless tree and so is finite (this argument is inspired by [8, Lemma 2]). We enumerate the components of $S^* - Ext(S^*)$ as $S_0^*, S_1^*, \ldots$ in such a way that for every $n \geqslant 0$, the set

$$S_n := S\left[\bigcup_{i \leqslant n} V(S_i^*)\right]$$

is a finite subtree of S containing the root r. Let us write $\partial(S_n) = E_{S^*}(S_n, S^* \setminus S_n)$, and note that $\partial(S_n) \subseteq Ext(S^*)$. We make the following definitions:

- For a given T-tribe $\mathcal{F}$ and ray R of T, we say that R *converges to ϵ according to $\mathcal{F}$* if for all members H of $\mathcal{F}$ the ray $H(R)$ is in ϵ. We say that R is *cut from ϵ according to $\mathcal{F}$* if for all members H of $\mathcal{F}$ the ray $H(R)$ is not in ϵ. Finally we say that $\mathcal{F}$ *determines whether R converges to ϵ* if either R converges to ϵ according to $\mathcal{F}$ or R is cut from ϵ according to $\mathcal{F}$.

- Similarly, for a given T-tribe $\mathcal{F}$ and vertex t of T, we say that t *dominates ϵ according to $\mathcal{F}$* if for all members H of $\mathcal{F}$ the vertex $H(t)$ dominates ϵ. We say that t is *cut from ϵ according to $\mathcal{F}$* if for all members H of $\mathcal{F}$ the vertex $H(t)$ does not dominate ϵ. Finally we say that $\mathcal{F}$ *determines whether t dominates ϵ* if either t dominates ϵ according to $\mathcal{F}$ or t is cut from ϵ according to $\mathcal{F}$.

- Given $n \in \mathbb{N}$, we say a thick T-tribe $\mathcal{F}$ *agrees about $\partial(S_n)$* if for each extension edge $e \in \partial(S_n)$, it determines whether R_e converges to ϵ. We say that it *agrees about $V(S_n)$* if for each vertex t of S_n, it determines whether t dominates ϵ.

- Since $\partial(S_n)$ and $V(S_n)$ are finite for all n, it follows from Lemma 9.5.3 that given some $n \in \mathbb{N}$, any thick T-tribe has a flat thick T-subtribe $\mathcal{F}$ such that $\mathcal{F}$ agrees about $\partial(S_n)$

and $V(S_n)$. Under these circumstances we set

$$\begin{aligned}
\partial_\epsilon(S_n) &:= \{e \in \partial(S_n) \colon R_e \text{ converges to } \epsilon \text{ according to } \mathcal{F}\}\,, \\
\partial_{\neg\epsilon}(S_n) &:= \{e \in \partial(S_n) \colon R_e \text{ is cut from } \epsilon \text{ according to } \mathcal{F}\}\,, \\
V_\epsilon(S_n) &:= \{t \in V(S_n) \colon t \text{ dominates } \epsilon \text{ according to } \mathcal{F}\}\,, \text{ and} \\
V_{\neg\epsilon}(S_n) &:= \{t \in V(S_n) \colon t \text{ is cut from } \epsilon \text{ according to } \mathcal{F}\}\,.
\end{aligned}$$

- Also, under these circumstances, let us write $S_n^{\neg\epsilon}$ for the component of the forest $S - \partial_\epsilon(S_n) - \{e \in E_S(S_n, S \setminus S_n) \colon e^- \in V_\epsilon(S_n)\}$ containing the root of S. Note that $S_n \subseteq S_n^{\neg\epsilon}$.

The following lemma contains a large part of the work needed for our inductive construction.

Lemma 9.6.3 (T-tribe refinement lemma)**.** *Suppose we have a thick T-tribe $\mathcal{F}_n$ concentrated at ϵ which agrees about $\partial(S_n)$ and $V(S_n)$ for some $n \in \mathbb{N}$. Let f denote the unique edge from S_n to $S_{n+1} \setminus S_n$. Then there is a thick T-tribe $\mathcal{F}_{n+1}$ concentrated at ϵ with the following properties:*

(i) *$\mathcal{F}_{n+1}$ agrees about $\partial(S_{n+1})$ and $V(S_{n+1})$.*

(ii) *$\mathcal{F}_{n+1} \cup \mathcal{F}_n$ agree about $\partial(S_n) \setminus \{f\}$ and $V(S_n)$.*

(iii) *$S_{n+1}^{\neg\epsilon} \supseteq S_n^{\neg\epsilon}$.*

(iv) *For all $H \in \mathcal{F}_{n+1}$ there is a finite $X \subseteq \Gamma$ such that $H(S_{n+1}^{\neg\epsilon}) \cap (X \cup C_\Gamma(X, \epsilon)) = H(V_\epsilon(S_{n+1}))$.*

Moreover, if $f \in \partial_\epsilon(S_n)$, and $R_f = v_0v_1v_2 \ldots \subseteq S^*$ (with $v_0 = f^+$) denotes the ray displaying self-similarity of T at f, then we may additionally assume:

(v) *For every $H \in \mathcal{F}_{n+1}$ and every $k \in \mathbb{N}$, there is $H' \in \mathcal{F}_{n+1}$ with*

- *$H' \subseteq_r H$*
- *$H'(S_n) = H(S_n)$,*
- *$H'(T_{v_0}) \subseteq_r H(T_{v_k})$, and*
- *$H'(R_f) \subseteq H(R_f)$.*

Proof. Concerning (v), if $f \in \partial_\epsilon(S_n)$ recall that according to Definition 9.6.2, the ray R_f satisfies that for all $k \in \mathbb{N}$ we have $T_{v_0} \leqslant_r T_{v_k}$ such that R_f gets embedded into itself. In particular, there is a subtree $\hat{T}_1$ of T_{v_1} which is a rooted subdivision of T_{v_0} with $\hat{T}_1(R_f) \subseteq R_f$, considering $\hat{T}_1$ as a rooted tree given by the tree order in T_{v_1}. If we define recursively for each $k \in \mathbb{N}$ $\hat{T}_k = \hat{T}_{k-1}(\hat{T}_1)$ then it is clear that $(\hat{T}_k \colon k \in \mathbb{N})$ is a family of rooted subdivisions of T_{v_0} such that for each $k \in \mathbb{N}$

- $\hat{T}_k \subseteq T_{v_k}$;
- $\hat{T}_k \supseteq \hat{T}_{k+1}$;
- $\hat{T}_k(R_f) \subseteq R_f$

Hence, for every subdivision H of T with $H \in \bigcup \mathcal{F}_n$ and every $k \in \mathbb{N}$, the subgraph $H(\hat{T}_k)$ is also a rooted subdivision of T_{v_0}. Let us construct a subdivision $H^{(k)}$ of T by letting $H^{(k)}$ be the minimal subtree of H containing $H(T \setminus T_{v_0}) \cup H(\hat{T}_k)$, where $H^{(k)}(T \setminus T_{v_0}) = H(T \setminus T_{v_0})$ and $H^{(k)}(T_{v_0}) = H(\hat{T}_k)$. Note that

$$H^{(k)}(T_{v_0}) = H(\hat{T}_k) \subseteq_r H^{(k-1)}(T_{v_0}) = H(\hat{T}_{k-1}) \subseteq_r \ldots \subseteq_r H(T_{v_k}).$$

In particular, for every subdivision $H \in \bigcup \mathcal{F}_n$ of T and every $k \in \mathbb{N}$, there is a subdivision $H^{(k)} \subseteq H$ of T such that $H^{(k)}(S_n^{\neg\epsilon}) = H(S_n^{\neg\epsilon})$, $H^{(k)}(T_{v_0}) \subseteq_r H(T_{v_k})$, and $H^{(k)}(R_f) \subseteq H(R_f)$. By the pigeon hole principle, there is an infinite index set $K_H = \{k_1^H, k_2^H, \ldots\} \subseteq \mathbb{N}$ such that $\{\{H^{(k)}\} \colon k \in K_H\}$ agrees about $\partial(S_{n+1})$. Consider the thick subtribe $\mathcal{F}'_n = \{F'_i \colon F \in \mathcal{F}_n, i \in \mathbb{N}\}$ of $\mathcal{F}_n$ with

$$(\dagger) \qquad F'_i := \{H^{(k_i^H)} \colon H \in F\}.$$

Observe that $\mathcal{F}'_n \cup \mathcal{F}_n$ still agrees about $\partial(S_n)$ and $V(S_n)$. (If $f \in \partial_{\neg\epsilon}(S_n)$, then skip this part and simply let $\mathcal{F}'_n := \mathcal{F}_n$.)

Concerning (iii), observe that for every $H \in \bigcup \mathcal{F}'_n$, since the rays $H(R_e)$ for $e \in \partial_{\neg\epsilon}(S_n)$ do not tend to ϵ, there is a finite vertex set X_H such that $H(R_e) \cap C(X_H, \epsilon) = \emptyset$ for all $e \in \partial_{\neg\epsilon}(S_n)$. Furthermore, since X_H is finite, for each such extension edge e there exists $x_e \in R_e$ such that

$$H(T_{x_e}) \cap C(X_H, \epsilon) = \emptyset.$$

By definition of extension edges, cf. Definition 9.6.2, for each $e \in \partial_{\neg\epsilon}(S_n)$ there is a rooted embedding of T_{e^+} into $H(T_{x_e})$. Hence, there is a subdivision $\tilde{H}$ of T with $\tilde{H} \leqslant H$ and $\tilde{H}(S_n) = H(S_n)$ such that $\tilde{H}(T_{e^+}) \subseteq H(T_{x_e})$ for each $e \in \partial_{\neg\epsilon}(S_n)$.

Note that if $e \in \partial_{\neg\epsilon}(S_n)$ and g is an extension edge with $e \leqslant g \in \partial(S_{n+1}) \setminus \partial(S_n)$, then $\tilde{H}(R_g) \subseteq \tilde{H}(S_{e^+}) \subseteq H(S_{x_e})$, and so

$$(\ddagger) \qquad \tilde{H}(R_g) \text{ doesn't tend to } \epsilon.$$

Define $\tilde{\mathcal{F}}_n$ to be the thick T-subtribe of $\mathcal{F}'_n$ consisting of the $\tilde{H}$ for every H in $\bigcup \mathcal{F}'_n$.

Now use Lemma 9.5.3 to chose a maximal thick flat subtribe $\mathcal{F}^*_n$ of $\tilde{\mathcal{F}}_n$ which agrees about $\partial(S_{n+1})$ and $V(S_{n+1})$, so it satisfies (i) and (ii). By ($\ddagger$), the tribe $\mathcal{F}^*_n$ satisfies (iii), and by maximality and ($\dagger$), it satisfies (v).

In our last step, we now arrange for (iv) while preserving all other properties. For each $H \in \bigcup \mathcal{F}^*_n$. Since $H(S_{n+1})$ is finite, we may find a finite separator Y_H such that

$$H(S_{n+1}) \cap (Y_H \cup C(Y_H, \epsilon)) = H(V_\epsilon(S_{n+1})).$$

Since Y_H is finite, for every vertex $t \in V_{\neg\epsilon}(S_{n+1})$, say with $N^+(t) = (t_i)_{i \in \mathbb{N}}$, there exists $n_t \in \mathbb{N}$ such that $C(Y_H, \epsilon) \cap H(T_{t_j}) = \emptyset$ for all $j \geqslant n_t$. Using Corollary 9.3.8, for every such t there is a rooted embedding

$$\{t\} \cup \bigcup_{j > N_t} T_{t_j} \leqslant_r \{t\} \cup \bigcup_{j > n_t} T_{t_j}.$$

fixing the root t. Hence there is a subdivision H' of T with $H' \leqslant H$ such that $H'(T \setminus S) = H(T \setminus S)$ and for every $t \in V_{\neg\epsilon}(S_{n+1})$

$$H'\left[\{t\} \cup \bigcup_{j > N_t} T_{t_j}\right] \cap C(Y_H, \epsilon) = \emptyset.$$

Moreover, note that by construction of $\tilde{F}_n$, every such H' automatically satisfies that

$$H(S_{e^+}) \cap C(X_H \cup Y_H, \epsilon) = \emptyset$$

for all $e \in \partial_{\neg\epsilon}(S_{n+1})$. Let $\mathcal{F}_{n+1}$ consist of the set of H' as defined above for all $H \in \mathcal{F}^*_n$. Then $X_H \cup Y_H$ is a finite separator witnessing that $\mathcal{F}_{n+1}$ satisfies (iv). $\square$

9.6.2 Only finitely many vertices dominate ϵ

We first note as in Lemma 9.5.4, that for every finite vertex set $X \subseteq V(\Gamma)$ only a thin subtribe of $\mathcal{F}$ can meet X, so a thick subtribe is contained in the graph $\Gamma - X$. By removing the set of vertices dominating ϵ, we may therefore assume without loss of generality that no vertex of Γ dominates ϵ.

Definition 9.6.4 (Bounder, extender)**.** *Suppose that some thick T-tribe $\mathcal{F}$ which is concentrated at ϵ agrees about S_n for some given $n \in \mathbb{N}$, and $Q_1^n, Q_2^n, \dots, Q_n^n$ are disjoint subdivisions of $S_n^{\neg\epsilon}$ (note, $S_n^{\neg\epsilon}$ depends on $\mathcal{F}$).*

- *A* bounder *for the $(Q_i^n : i \in [n])$ is a finite set X of vertices in Γ separating all the Q_i from ϵ, i.e. such that*
$$C(X, \epsilon) \cap \bigcup_{i=1}^{n} Q_i^n = \emptyset.$$
- *An* extender *for the $(Q_i^n : i \in [n])$ is a family $\mathcal{E}_n = (E_{e,i}^n : e \in \partial_\epsilon(S_n), i \in [n])$ of rays in Γ tending to ϵ which are disjoint from each other and also from each Q_i^n except at their initial vertices, and where the start vertex of $E_{e,i}^n$ is $Q_i^n(e^-)$.*

To prove Theorem 9.1.3, we now assume inductively that for some $n \in \mathbb{N}$, with $r := \lfloor n/2 \rfloor$ and $s := \lceil n/2 \rceil$ we have:

1. A thick T-tribe $\mathcal{F}_r$ in Γ concentrated at ϵ which agrees about $\partial(S_r)$, with a boundary $\partial_\epsilon(S_r)$ such that $S_{r-1}^{\neg\epsilon} \subseteq S_r^{\neg\epsilon}$.[4]
2. a family $(Q_i^n : i \in [s])$ of s pairwise disjoint T-suitable subdivisions of $S_r^{\neg\epsilon}$ in Γ with $Q_i^n(S_{r-1}^{\neg\epsilon}) = Q_i^{n-1}$ for all $i \leqslant s-1$,
3. a bounder X_n for the $(Q_i^n : i \in [s])$, and
4. an extender $\mathcal{E}_n = (E_{e,i}^n : e \in \partial_\epsilon(S_r^{\neg\epsilon}), i \in [s])$ for the $(Q_i^n : i \in [s])$.

The base case $n = 0$ it easy, as we simply may choose $\mathcal{F}_0 \leqslant_r \mathcal{F}$ to be any thick T-subtribe in Γ which agrees about $\partial(S_0)$, and let all other objects be empty.

So, let us assume that our construction has proceeded to step $n \geqslant 0$. Our next task splits into two parts: First, if $n = 2k-1$ is odd, we extend the already existing k subdivisions $(Q_i^n : i \in [k])$ of $S_{k-1}^{\neg\epsilon}$ to subdivisions $(Q_i^{n+1} : i \in [k])$ of $S_k^{\neg\epsilon}$. And secondly, if $n = 2k$ is even, we construct a further disjoint copy Q_{k+1}^{n+1} of $S_k^{\neg\epsilon}$.

Construction part 1: $n = 2k-1$ is odd. By assumption, $\mathcal{F}_{k-1}$ agrees about $\partial(S_{k-1})$. Let f denote the unique edge from S_{k-1} to $S_k \setminus S_{k-1}$. We first apply Lemma 9.6.3 to $\mathcal{F}_{k-1}$ in order to find a thick T-tribe $\mathcal{F}_k$ concentrated at ϵ satisfying properties (i)–(v). In particular, $\mathcal{F}_k$ agrees about $\partial(S_k)$ and $S_{k-1}^{\neg\epsilon} \subseteq S_k^{\neg\epsilon}$

We first note that if $f \notin \partial_\epsilon(S_{k-1})$, then $S_{k-1}^{\neg\epsilon} = S_k^{\neg\epsilon}$, and we can simply take $Q_i^{n+1} := Q_i^n$ for all $i \in [k]$, $\mathcal{E}_{n+1} := \mathcal{E}_n$ and $X_{n+1} := X_n$.

Otherwise, we have $f \in \partial_\epsilon(S_{k-1})$. By Lemma 9.5.5(2) $\mathcal{F}_k$ is concentrated at ϵ, and so we may pick a collection $\{H_1, \dots, H_N\}$ of disjoint subdivisions of T from some $F \in \mathcal{F}_k$, all of which are contained in $C(X_n, \epsilon)$, where $N = |\mathcal{E}_n|$. By Lemma 9.4.3 there is some linkage $\mathcal{P} \subseteq C(X_n, \epsilon)$ from

$$\mathcal{E}_n \text{ to } (H_j(R_f) : j \in [N]),$$

[4] Note that since ϵ is undominated, every thick T-tribe agrees about the fact that $V_\epsilon(S_i) = \emptyset$ for all $i \in \mathbb{N}$.

which is after X_n. Let us suppose that the linkage $\mathcal{P}$ joins a vertex $x_{e,i} \in E^n_{e,i}$ to $y_{\sigma(e,i)} \in H_{\sigma(e,i)}(R_f)$ via a path $P_{e,i} \in \mathcal{P}$. Let $z_{\sigma(e,i)}$ be a vertex in R_f such that $y_{\sigma(e,i)} \leqslant H_{\sigma(e,i)}(z_{\sigma(e,i)})$ in the tree order on $H_{\sigma(e,i)}(T)$.

By property (v) of $\mathcal{F}_k$ in Lemma 9.6.3, we may assume without loss of generality that for each H_j there is a another member $H'_j \subseteq H_j$ of $\mathcal{F}_k$ such that $H'_j(T_{f^+}) \subseteq_r H_j(T_{z_j})$. Let $\hat{P}_j \subseteq H'_j$ denote the path from $H_j(y_j)$ to $H'_j(f^+)$.

Now for each $i \in [k]$, define

$$Q^{n+1}_i = Q^n_i \cup E^n_{f,i} x_{f,i} P_{f,i} y_{\sigma(f,i)} \hat{P}_{\sigma(f,i)} \cup H'_{\sigma(f,i)}(S^{\neg\epsilon}_k \setminus S^{\neg\epsilon}_{k-1}).$$

By construction, each Q^{n+1}_i is a T-suitable subdivision of $S^{\neg\epsilon}_k$.

By Lemma 9.6.3(iv) we may find a finite set $X_{n+1} \subseteq \Gamma$ with $X_n \subseteq X_{n+1}$ such that

$$C(X_{n+1}, \epsilon) \cap \Big(\bigcup_{i \in [k]} Q^{n+1}_i \Big) = \emptyset.$$

This set X_{n+1} will be our bounder.

Define an extender $\mathcal{E}_{n+1} = (E^{n+1}_{e,i} \colon e \in \partial_\epsilon(S_k), i \in [k])$ for the Q^{n+1}_i as follows:

- For $e \in \partial_\epsilon(S_{k-1}) \setminus \{f\}$, let $E^{n+1}_{e,i} := E^n_{e,i} x_{e,i} P_{e,i} y_{\sigma(e,i)} H_{\sigma(e,i)}(R_f)$.
- For $e \in \partial_\epsilon(S_k) \setminus \partial(S_{k-1})$, let $E^{n+1}_{e,i} := H'_{\sigma(e,i)}(R_e)$.

Since each $H_{\sigma(e,i)}, H'_{\sigma(e,i)} \in \bigcup \mathcal{F}_k$, and $\mathcal{F}_k$ determines that R_f converges to ϵ, these are indeed ϵ rays. Furthermore, since $H'_{\sigma(e,i)} \subseteq H_{\sigma(e,i)}$ and $\{H_1, \ldots, H_N\}$ are disjoint, it follows that the rays are disjoint.

Construction part 2: $n = 2k$ is even. If $\partial_\epsilon(S_k) = \emptyset$, then $S^{\neg\epsilon}_k = S$, and so picking any element Q^{n+1}_{k+1} from $\mathcal{F}_k$ with $Q^{n+1}_{k+1} \subseteq C(X_n, \epsilon)$ gives us a further copy of S disjoint from all the previous ones. Using Lemma 9.6.3(iv), there is a suitable bounder $X_{n+1} \supseteq X_n$ for Q^{n+1}_{k+1}, and we are done. Otherwise, pick $e_0 \in \partial_\epsilon(S_k)$ arbitrary.

Since $\mathcal{F}_k$ is concentrated at ϵ, we may pick a collection $\{H_1, \ldots, H_N\}$ of disjoint subdivisions of T from $\mathcal{F}_k$ all contained in $C(X_n, \epsilon)$, where N is large enough so that we may apply Lemma 9.4.4 to find a linkage $\mathcal{P} \subseteq C(X_n, \epsilon)$ from

$$\mathcal{E}_n \ \text{ to } \ (H_i(R_{e_0}) \colon i \in [N]),$$

after X_n, avoiding say H_1. Let us suppose the linkage $\mathcal{P}$ joins a vertex $x_{e,i} \in E^n_{e,i}$ to $y_{\sigma(e,i)} \in H_{\sigma(e,i)}(R_{e_0})$ via a path $P_{e,i} \in \mathcal{P}$. Define

$$Q^{n+1}_{k+1} = H_1(S^{\neg\epsilon}_k).$$

Note that Q^{n+1}_{k+1} is a T-suitable subdivision of $S^{\neg\epsilon}_k$.

By Lemma 9.6.3(iv) there is a finite set $X_{n+1} \subseteq \Gamma$ with $X_n \subseteq X_{n+1}$ such that $C(X_{n+1}, \epsilon) \cap Q^{n+1}_{k+1} = \emptyset$. This set X_{n+1} will be our new bounder.

Define the extender $\mathcal{E}_{n+1} = (E^{n+1}_{e,i} \colon e \in \partial_\epsilon(S_{k+1}), i \in [k+1])$ of ϵ-rays as follows:

- For $i \in [k]$, let $E^{n+1}_{e,i} := E^n_{e,i} x_{e,i} P_{e,i} y_{\sigma(e,i)} H_{\sigma(e,i)}(R_{e_0})$.
- For $i = k+1$, let $E^{n+1}_{e,k+1} := H_1(R_e)$ for all $e \in \partial_\epsilon(S_{k+1})$.

Once the construction is complete, let us define $H_i := \bigcup_{n \geqslant 2i-1} Q^n_i$.

Since $\bigcup_{n \in \mathbb{N}} S^{\neg\epsilon}_n = S$, and due to the extension property (2), the collection $(H_i)_{i \in \mathbb{N}}$ is an S-horde. □

We remark that our construction so far suffices to give a complete proof that countable trees are $\leqslant$-ubiquitous. Indeed, it is well-known that an end of Γ is dominated by infinitely many distinct vertices if and only if Γ contains a subdivision of $K_{\aleph_0}$ [43, Exercise 19, Chapter 8], in which case proving ubiquity becomes trivial:

Lemma 9.6.5. *For any countable graph G, we have $\aleph_0 \cdot G \subseteq K_{\aleph_0}$.*

Proof. By partitioning the vertex set of $K_{\aleph_0}$ into countably many infinite parts, we see that $\aleph_0 \cdot K_{\aleph_0} \subseteq K_{\aleph_0}$. Also, clearly $G \subseteq K_{\aleph_0}$. Hence, we have $\aleph_0 \cdot G \subseteq \aleph_0 \cdot K_{\aleph_0} \subseteq K_{\aleph_0}$. □

9.6.3 Infinitely many vertices dominate ϵ

The argument in this case is very similar to that in the previous subsection. We define bounders and extenders just as before. We once more assume inductively that for some $n \in \mathbb{N}$, with $r := \lfloor n/2 \rfloor$, we have objects given by (1)–(4) as in the last section, and which in addition satisfy

(5) $\mathcal{F}_r$ agrees about $V(S_r)$.

(6) For any $t \in V_\epsilon(S_r)$ the vertex $Q_i^n(t)$ dominates ϵ.

The base case is again trivial, so suppose that our construction has proceeded to step $n \geqslant 0$. The construction is split into two parts just as before, where the case $n = 2k$, in which we need to refine our T-tribe and find a new copy Q_{k+1}^{n+1} of $S_k^{\neg\epsilon}$, proceeds just as in the last section.

If $n = 2k - 1$ is odd, and if $f \in \partial_{\neg\epsilon}(S_{k-1})$ or $\partial_\epsilon(S_{k-1})$, then we proceed as in the last subsection. But these are no longer the only possibilities. It follows from the definition of $S_k^{\neg\epsilon}$ that there is one more option, namely that $f^- \in V_\epsilon(S_k)$. In this case we modify the steps of the construction as follows:

We first apply Lemma 9.6.3 to $\mathcal{F}_{k-1}$ in order to find a thick T-tribe $\mathcal{F}_{k-1}$ which agrees about $\partial(S_k)$ and $V(S_k)$.

Then, by applying Lemma 9.4.5 to tails of the rays $E_{e,i}^n$ in $C_\Gamma(X_n, \epsilon)$, we obtain a family $\mathcal{P}_{n+1}$ of paths $P_{e,i}^{n+1}$ which are disjoint from each other and from the Q_i^n except at their initial vertices, where the initial vertex of $P_{e,i}^{n+1}$ is $Q_i^n(e^-)$ and the final vertex $y_{e,i}^{n+1}$ of $P_{e,i}^{n+1}$ dominates ϵ.

Since $\mathcal{F}_k$ is concentrated at ϵ, we may pick a collection $\{H_1, \dots, H_k\}$ of disjoint subdivisions of T from $\mathcal{F}_k$ all contained in $C(X_n \cup \bigcup \mathcal{P}_{n+1}, \epsilon)$.

Now for each $i \in [k]$, define

$$\hat{Q}_i^{n+1} = Q_i^n \cup H_i(f^-) \cup H_i(S_k^{\neg\epsilon} \setminus S_{k-1}^{\neg\epsilon}).$$

These are almost T-suitable subdivisions of $S_k^{\neg\epsilon}$, except we need to add a path between $Q_i^n(f^-)$ and $H_i(f^-)$.

By applying Lemma 9.4.5 to tails of the rays $H_i(R_e)$ inside $C(X_n \cup \bigcup \mathcal{P}_{n+1}, \epsilon)$ with $e \in \partial_\epsilon(S_{k+1}) \setminus \partial(S_k)$ we can construct a family $\mathcal{P}'_{n+1} := \{P_{e,i}^{n+1} \colon e \in \partial_\epsilon(S_{k+1}) \setminus \partial_\epsilon(S_k), i \leqslant k\}$ of paths which are disjoint from each other and from the $\hat{Q}_i^{n+1}$ except at their initial vertices, where the initial vertex of $P_{e,i}^{n+1}$ is $H_i(e^-)$ and the final vertex $y_{e,i}^{n+1}$ of $P_{e,i}^{n+1}$ dominates ϵ. Therefore the family

$$\mathcal{P}_{n+1} \cup \mathcal{P}'_{n+1} = (P_{e,i}^{n+1} \colon e \in \partial_\epsilon(S_{k+1}), i \in [k])$$

is a family of disjoint paths, which are also disjoint from the $\hat{Q}_i^{n+1}$ except at their initial vertices, where the initial vertex of $P_{e,i}^{n+1}$ is $H_i(e^-)$ or $Q_i^n(e^-)$ and the final vertex $y_{e,i}^{n+1}$ of $P_{e,i}^{n+1}$ dominates ϵ.

Since $Q_i^n(f^-)$ and $H_i(f^-)$ both dominate ϵ for all i, we may recursively build a sequence $\hat{\mathcal{P}}_{n+1} = \{\hat{P}_i \colon 1 \leqslant i \leqslant k\}$ of disjoint paths $\hat{P}_i$ from $Q_i^n(f^-)$ to $H_i(f^-)$ with all internal vertices in $C(X_{n+1} \cup \left(\bigcup \mathcal{P}'_{n+1} \cup \bigcup \mathcal{P}_{n+1}\right), \epsilon)$. Letting $Q_i^{n+1} = \hat{Q}_i^{n+1} \cup \hat{P}_i$, we see that each Q_i^{n+1} is a T-suitable subdivision of $S_k^{\neg\epsilon}$ in Γ.

Our new bounder will be $X_{n+1} := X_n \cup \bigcup \hat{\mathcal{P}}_{n+1} \cup \bigcup \mathcal{P}'_{n+1} \cup \bigcup \mathcal{P}_{n+1}$.

Finally, let us apply Lemma 9.4.6 to $Y := \{y^{n+1}_{e,i} \colon e \in \partial_\epsilon(S_{n+1}), i \leqslant k\}$ in $\Gamma[Y \cup C(X_{n+1}, \epsilon)]$. This gives us a family of disjoint rays

$$\hat{\mathcal{E}}_{n+1} = (\hat{E}^{n+1}_{e,i} \colon e \in \partial_\epsilon(S_{k+1}), i \in [k])$$

such that $\hat{E}^{n+1}_{e,i}$ has initial vertex $y^{n+1}_{e,i}$. Let us define our new extender $\mathcal{E}_{n+1}$ given by

- $E^{n+1}_{e,i} = Q^n_i(e^-)P^{n+1}_{e,i}y^{n+1}_{e,i}\hat{E}^{n+1}_{e,i}$ if $e \in \partial_\epsilon(S_k), i \in [k]$;
- $E^{n+1}_{e,i} = H_i(e^-)P^{n+1}_{e,i}y^{n+1}_{e,i}\hat{E}^{n+1}_{e,i}$ if $e \in \partial_\epsilon(S_{k+1}) \setminus \partial(S_k), i \in [k]$.

This concludes the proof of Theorem 9.1.3. □

9.7 The induction argument

We consider T as a rooted tree with root r. In Section 9.6 we constructed an S-horde for any countable subtree S of T. In this section we will extend an S-horde for some specific countable subtree S to a T-horde, completing the proof of Theorem 9.1.2.

Recall that for a vertex t of T and an infinite cardinal κ we say that a child t' of t is κ-embeddable if there are at least κ children t'' of t such that $T_{t'}$ is a (rooted) topological minor of $T_{t''}$ (Definition 9.3.6). By Corollary 9.3.7, the number of children of t which are not κ-embeddable is less than κ.

Definition 9.7.1 (κ-closure)**.** *Let T be an infinite tree with root r.*

- *If S is a subtree of T and S' is a subtree of S, then we say that S' is* κ-closed in S *if for any vertex t of S' all children of t in S are either in S' or else are κ-embeddable.*
- *The* κ-closure of S' in S *is the smallest κ-closed subtree of S including S'.*

Lemma 9.7.2. *Let S' be a subtree of S. If κ is a uncountable regular cardinal and S' has size less than κ, then the κ-closure of S' in S also has size less than κ.*

Proof. Let $S'(0) := S'$ and define inductively $S'(n+1)$ to consist of $S'(n)$ together with all non-κ-embeddable children contained in S for all vertices of $S'(n)$. It is clear that $\bigcup_{n\in\mathbb{N}} S'(n)$ is the κ-closure of S'. If κ_n denotes the size of $S'(n)$, then $\kappa_n < \kappa$ by induction with Corollary 9.3.7. Therefore, the size of the κ-closure is bounded by $\sum_{n\in\mathbb{N}} \kappa_n < \kappa$, since κ has uncountable cofinality. □

We will construct the desired T-horde via transfinite induction on the cardinals $\mu \leqslant |T|$. Our first lemma illustrates the induction step for regular cardinals.

Lemma 9.7.3. *Let κ be an uncountable regular cardinal. Let S be a rooted subtree of T of size at most κ and let S' be a κ-closed rooted subtree of S of size less than κ. Then any S'-horde $(H_i \colon i \in \mathbb{N})$ can be extended to an S-horde.*

Proof. Let $(s_\alpha \colon \alpha < \kappa)$ be an enumeration of the vertices of S such that the parent of any vertex appears before that vertex in the enumeration, and for any α let S_α be the subtree of T with vertex set $V(S') \cup \{s_\beta \colon \beta < \alpha\}$. Let $\bar{S}_\alpha$ denote the κ-closure of S_α in S, and observe that $|\bar{S}_\alpha| < \kappa$ by Lemma 9.7.2.

We will recursively construct for each α an $\bar{S}_\alpha$-horde $(H^\alpha_i \colon i \in \mathbb{N})$ in Γ, where each of these hordes extends all the previous ones. For $\alpha = 0$ we let $H^0_i = H_i$ for each $i \in \mathbb{N}$. For any limit ordinal λ we have $\bar{S}_\lambda = \bigcup_{\beta<\lambda} \bar{S}_\beta$, and so we can take $H^\lambda_i = \bigcup_{\beta<\lambda} H^\beta_i$ for each $i \in \mathbb{N}$.

For any successor ordinal $\alpha = \beta + 1$, if $s_\beta \in \bar{S}_\beta$, then $\bar{S}_\alpha = \bar{S}_\beta$, and so we can take $H_i^\alpha = H_i^\beta$ for each $i \in \mathbb{N}$. Otherwise, $\bar{S}_\alpha$ is the κ-closure of $\bar{S}_\beta + s_\beta$, and so $\bar{S}_\alpha - \bar{S}_\beta$ is a subtree of T_{s_β}. Furthermore, since s_β is not contained in $\bar{S}_\beta$, it must be κ-embeddable.

Let s be the parent of s_β. By suitability of the H_i^β, we can find for each $i \in \mathbb{N}$ some subdivision $\hat{H}_i$ of T_s with $\hat{H}_i(s) = H_i^\beta(s)$. We now build the H_i^α recursively in i as follows:

Let t_i be a child of s such that T_{t_i} has a rooted subdivision K of T_{s_β}, and such that $\hat{H}_i(T_{t_i} + s) - \hat{H}_i(s)$ is disjoint from all H_j^α with $j < i$ and from all H_j^β. Since there are κ disjoint possibilities for K, and all H_j^α with $j < i$ and all H_j^β cover less than κ vertices in Γ, such a choice of K is always possible. Then let H_i^α be the union of H_i^β with $\hat{H}_i(K(\bar{S}_\alpha - \bar{S}_\beta) + st_i)$.

This completes the construction of the $(H_i^\alpha \colon i \in \mathbb{N})$. Obviously, each H_i^α for $i \in \mathbb{N}$ is a subdivision of $\bar{S}_\alpha$ with $H_i^\alpha(\bar{S}_\gamma) = H_i^\gamma$ for all $\gamma < \alpha$, and all of them are pairwise disjoint for $i \neq j \in \mathbb{N}$. Moreover, H_i^α is T-suitable since for all vertices $H_i^\alpha(t)$ whose t-suitability is not witnessed in previous construction steps, their suitability is witnessed now by the corresponding subtree of $\hat{H}_i$. Hence $(\bigcup_{\alpha<\kappa} H_i^\alpha \colon i \in \mathbb{N})$ is the desired S-horde extending $(H_i \colon i \in \mathbb{N})$. □

Our final lemma will deal with the induction step for singular cardinals. The crucial ingredient will be to represent a tree S of singular cardinality μ as a continuous increasing union of $<\mu$-sized subtrees $(S_\varrho \colon \varrho < \mathrm{cf}(\mu))$ where each S_ϱ is $|S_\varrho|^+$-closed in S. This type of argument is based on Shelah's singular compactness theorem, see e.g. [119], but can be read without knowledge of the paper.

Definition 9.7.4 (S-representation)**.** *For a tree S with $|S| = \mu$, we call a sequence $\mathcal{S} = (S_\varrho \colon \varrho < \mathrm{cf}(\mu))$ of subtrees of S with $|S_\varrho| = \mu_\varrho$ an* S-representation *if*

- *$(\mu_\varrho \colon \varrho < \mathrm{cf}(\mu))$ is a strictly increasing continuous sequence of cardinals less than μ which is cofinal for μ,*
- *$S_\varrho \subseteq S_{\varrho'}$ for all $\varrho < \varrho'$, i.e. $\mathcal{S}$ is* increasing*,*
- *for every limit $\lambda < \mathrm{cf}(\mu)$ we have $\bigcup_{\varrho<\lambda} S_\varrho = S_\lambda$, i.e. $\mathcal{S}$ is* continuous*,*
- *$\bigcup_{\varrho<\mathrm{cf}(\mu)} S_\varrho = S$, i.e. $\mathcal{S}$ is* exhausting*,*
- *S_ϱ is μ_ϱ^+-closed in S for all $\varrho < \mathrm{cf}(\mu)$, where μ_ϱ^+ is the successor cardinal of μ_ϱ.*

Moreover, for a tree $S' \subseteq S$ we say that $\mathcal{S}$ is an S-representation extending S' *if additionally*

- *$S' \subseteq S_\varrho$ for all $\varrho < \mathrm{cf}(\mu)$.*

Lemma 9.7.5. *For every tree S of singular cardinality and every subtree S' of S with $|S'| < |S|$ there is an S-representation extending S'.*

Proof. Let $|S| = \mu$ be singular, and let $|S'| = \kappa$. Let $(s_\alpha \colon \alpha < \mu)$ be an enumeration of the vertices of S. Let γ be the cofinality of μ and let $(\mu_\varrho \colon \varrho < \gamma)$ be a strictly increasing continuous cofinal sequence of cardinals less than μ with $\mu_0 > \gamma$ and $\mu_0 > \kappa$. By recursion on i we choose for each $i \in \mathbb{N}$ a sequence $(S_\varrho^i \colon \varrho < \gamma)$ of subtrees of S of cardinality μ_ϱ, where the vertices of each S_ϱ^i are enumerated as $(s_{\varrho,\alpha}^i \colon \alpha < \mu_\varrho)$, such that:

1. S_ϱ^i is μ_ϱ^+-closed.
2. S' is a subtree of S_ϱ^i.
3. $S_{\varrho'}^i$ is a subtree of S_ϱ^i for $\varrho' < \varrho$.
4. $s_\alpha \in S_\varrho^i$ for $\alpha < \mu_\varrho$.

5. $s^j_{\varrho',\alpha} \in S^i_\varrho$ for any $j < i$, $\varrho \leqslant \varrho' < \gamma$ and $\alpha < \mu_\varrho$

This is achieved by recursion on ϱ as follows: For any given $\varrho < \gamma$, let X^i_ϱ be the set of all vertices which are forced to lie in S^i_ϱ by conditions 2–5, that is, all vertices of S' or of $S^i_{\varrho'}$ with $\varrho' < \varrho$, all s_β with $\beta < \mu_\varrho$ and all $s^j_{\varrho',\alpha}$ with $j < i$, $\varrho \leqslant \varrho' < \gamma$ and $\alpha < \mu_\varrho$. Then X^i_ϱ has cardinality μ_ϱ and so it is included in a subtree of S of cardinality μ_ϱ. We take S^i_ϱ to be the μ^+_ϱ-closure of this subtree in S. Note that, since μ^+_ϱ is regular, it follows from Lemma 9.7.2 that S^i_ϱ has cardinality μ_ϱ.

For each $\varrho < \gamma$, let $S_\varrho := \bigcup_{i\in\mathbb{N}} S^i_\varrho$. Then each S_ϱ is a union of μ^+_ϱ-closed trees and so is μ^+_ϱ-closed itself. Furthermore, each S_ϱ clearly has cardinality μ_ϱ.

It follows from 4 that $S = \bigcup_{\varrho<\gamma} S_\varrho$. Thus, it remains to argue that our sequence is indeed continuous, i.e. that for any limit ordinal $\lambda < \gamma$ we have $S_\lambda = \bigcup_{\varrho<\lambda} S_\varrho$. The inclusion $\bigcup_{\varrho<\lambda} S_\varrho \subseteq S_\lambda$ is clear from 3. For the other inclusion, let s be any element of S_λ. Then there is some $i \in \mathbb{N}$ with $s \in S^i_\lambda$ and so there is some $\alpha < \mu_\alpha$ with $s = s^i_{\lambda,\alpha}$. Then by continuity there is some $\sigma < \lambda$ with $\alpha < \mu_\sigma$ and so $s \in S^{i+1}_\sigma \subseteq S_\sigma \subseteq \bigcup_{\varrho<\lambda} S_\varrho$. □

Lemma 9.7.6. *Let μ be a cardinal. Then for any rooted subtree S of T of size μ and any uncountable regular cardinal $\kappa \leqslant \mu$, any S'-horde $(H_i : i \in \mathbb{N})$ of a κ-closed rooted subtree S' of S of size less than κ can be extended to an S-horde.*

Proof. The proof is by transfinite induction on μ. If μ is regular, we let S'' be the μ-closure of S' in S. Thus S'' has size less than μ. So by the induction hypothesis $(H_i : i \in \mathbb{N})$ can be extended to an S''-horde, which by Lemma 9.7.3 can be further extended to an S-horde.

So let us assume that μ is singular, and write $\gamma = \mathrm{cf}(\mu)$. By Lemma 9.7.5, fix an S-representation $\mathcal{S} = (S_\varrho : \varrho < \mathrm{cf}(\mu))$ extending S' with $|S'| < |S_0|$.

We now recursively construct for each $\varrho < \gamma$ an S_ϱ-horde $(H^\varrho_i : i \in \mathbb{N})$, where each of these hordes extends all the previous ones and $(H_i : i \in \mathbb{N})$. Using that each S_ϱ is μ^+_ϱ-closed in S, we can find $(H^0_i : i \in \mathbb{N})$ by the induction hypothesis, and if ϱ is a successor ordinal we can find $(H^\varrho_i : i \in \mathbb{N})$ by again using the induction hypothesis. For any limit ordinal λ we set $H^\lambda_i = \bigcup_{\varrho<\lambda} H^\varrho_i$ for each $i \in \mathbb{N}$, which yields an S_λ-horde by the continuity of $\mathcal{S}$.

This completes the construction of the H^ϱ_i. Then $(\bigcup_{\varrho<\gamma} H^\varrho_i : i \in \mathbb{N})$ is an S-horde extending $(H_i : i \in \mathbb{N})$. □

Finally, with the right induction start we obtain the following theorem and hence a proof of Theorem 9.1.2.

Theorem 9.7.7. *Let T be a tree and Γ a graph such that $nT \leqslant \Gamma$ for every $n \in \mathbb{N}$. Then there is a T-horde, and hence $\aleph_0 T \leqslant \Gamma$.*

Proof. By Theorem 9.1.3, we may assume that T is uncountable. Let S' be the $\aleph_1$-closure of the root $\{r\}$ in T. Then S' is countable by Lemma 9.7.2 and so there is an S'-horde in Γ by Theorem 9.1.3. This can be extended to a T-horde in Γ by Lemma 9.7.6 with $\mu = |T|$. □

Chapter 10

Ubiquity of graphs with non-linear end structure

10.1 Introduction

This paper is the second in a series of papers making progress towards a conjecture of Andreae on the *ubiquity* of graphs. Given a graph G and some relation $\triangleleft$ between graphs we say that G is $\triangleleft$-*ubiquitous* if whenever Γ is a graph such that $nG \triangleleft \Gamma$ for all $n \in \mathbb{N}$, then $\aleph_0 G \triangleleft \Gamma$, where αG denotes the disjoint union of α many copies of G. For example, a classic result of Halin [71] says that the ray is $\subseteq$-ubiquitous, where $\subseteq$ is the subgraph relation.

Examples of graphs which are not ubiquitous with respect to the subgraph or topological minor relation are known (see [14] for some particularly simple examples). In [13] Andreae initiated the study of ubiquity of graphs with respect to the minor relation $\preccurlyeq$. He constructed a graph which is not $\preccurlyeq$-ubiquitous, however the construction relied on the existence of a counterexample to the well-quasi-ordering of infinite graphs under the minor relation, for which only examples of very large cardinality are known [122]. In particular, the question of whether there exists a countable graph which is not $\preccurlyeq$-ubiquitous remains open. Most importantly, however, Andreae [13] conjectured that at least all locally finite graphs, those with all degrees finite, should be $\preccurlyeq$-ubiquitous.

Conjecture 9.1.1. *[The Ubiquity Conjecture] Every locally finite connected graph is $\preccurlyeq$-ubiquitous.*

In [14] Andreae proved that his conjecture holds for a large class of locally finite graphs. The exact definition of this class is technical, but in particular his result implies the following.

Theorem 10.1.1 (Andreae, [14, Corollary 2])**.** *Let G be a connected, locally finite graph of finite tree-width such that every block of G is finite. Then G is $\preccurlyeq$-ubiquitous.*

Note that every end in such a graph G must have degree[1] one.

Andreae's proof employs deep results about well-quasi-orderings of labelled (infinite) trees [92]. Interestingly, the way these tools are used does not require the extra condition in Theorem 10.1.1 that every block of G is finite and so it is natural to ask if his proof can be adapted to remove this condition. And indeed, it is the purpose of the present and subsequent paper in our series, [26], to show that this is possible, i.e. that all connected, locally finite graphs of finite tree-width are $\preccurlyeq$-ubiquitous.

The present paper lays the groundwork for this extension of Andreae's result. The fundamental obstacle one encounters when trying to extend Andreae's methods is the following: Let

[1] A precise definitions of rays, the ends of a graph, their degree, and what it means for a ray to converge to an end can be found in Section 10.2.

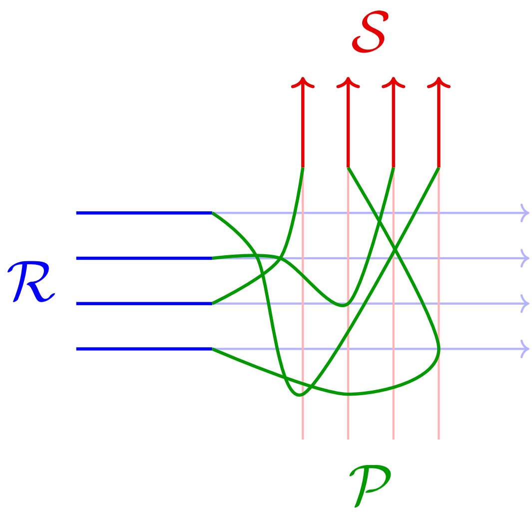

Figure 10.1: A linkage between $\mathcal{R}$ and $\mathcal{S}$.

$[n] = \{1, 2, \ldots, n\}$. In the proof we often have two families of disjoint rays $\mathcal{R} = (R_i : i \in [n])$ and $\mathcal{S} = (S_j : j \in [m])$ in Γ, which we may assume all converge[1] to a common end of Γ, and we wish to find a *linkage* between $\mathcal{R}$ and $\mathcal{S}$, that is, an injective function $\sigma : [n] \to [m]$ and a set $\mathcal{P}$ of disjoint finite paths P_i from $x_i \in R_i$ to $y_{\sigma(i)} \in S_{\sigma(i)}$ such that the walks

$$\mathcal{T} = (R_i x_i P_i y_{\sigma(i)} S_{\sigma(i)} : i \in [n])$$

formed by following each R_i along to x_i, then following the path P_i to $y_{\sigma(i)}$, then following the tail of $S_{\sigma(i)}$, form a family of disjoint rays (see Figure 10.1). Broadly, we can think of this as 're-routing' the rays $\mathcal{R}$ to some subset of the rays in $\mathcal{S}$. Since all the rays in $\mathcal{R}$ and $\mathcal{S}$ converge to the same end of Γ, it is relatively simple to show that, as long as $n \leqslant m$, there is enough connectivity between the rays in Γ so that such a linkage always exists.

However, in practice it is not enough for us to be guaranteed the existence of some injection σ giving rise to a linkage, but instead we want to choose σ in advance, and be able to find a corresponding linkage afterwards.

In general, however, it is quite possible that for certain choices of σ no suitable linkage exists. Consider for example the case where Γ is the *half grid* (briefly denoted by $\mathbb{Z}\square\mathbb{N}$), which is the graph whose vertex set is $\mathbb{Z} \times \mathbb{N}$ and where two vertices are adjacent if they differ in precisely one co-ordinate and the difference in that co-ordinate is one. If we consider two sufficiently large families of disjoint rays $\mathcal{R}$ and $\mathcal{S}$ in Γ, then it is not hard to see that both $\mathcal{R}$ and $\mathcal{S}$ inherit a linear ordering from the planar structure of Γ, which must be preserved by any linkage between them.

Analysing this situation gives rise to the following definition: We say that an end ϵ of a graph G is *linear* if for every finite set $\mathcal{R}$ of at least three disjoint rays in G which converge to ϵ we can order the elements of $\mathcal{R}$ as $\mathcal{R} = \{R_1, R_2, \ldots, R_n\}$ such that for each $1 \leqslant k < i < \ell \leqslant n$, the rays R_k and R_ℓ belong to different ends of $G - V(R_i)$.

Thus the half grid has a unique end and it is linear. On the other end of the spectrum, let us say that a graph G has *nowhere-linear end structure* if no end of G is linear. Since ends of degree at most two are automatically linear, every end of a graph with nowhere-linear end structure must have degree at least three.

Our main theorem in this paper is the following.

Theorem 10.1.2. *Every locally finite connected graph with nowhere-linear end structure is $\preccurlyeq$-ubiquitous.*

Roughly, if we assume that every end of G has nonlinear structure, then the fact that $nG \preccurlyeq \Gamma$ for all $n \in \mathbb{N}$ allows us to deduce that Γ must also have some end with a sufficiently complicated structure that we can always find suitable linkages for all σ as above. In fact, this property is so strong that we do not need to follow Andreae's strategy for such graphs. We can use the linkages to directly build a $K_{\aleph_0}$-minor of Γ, and it follows that $\aleph_0 G \preccurlyeq \Gamma$.

In later papers in the series, we shall need to make more careful use of the ideas developed here. We shall analyse the possible kinds of linkages which can arise between two families of rays converging to a given end. If some end of Γ admits many different kinds of linkages, then we can again find a $K_{\aleph_0}$-minor. If not, then we can use the results of the present paper to show that certain ends of G are linear. This extra structure allows us to carry out an argument like that of Andreae, but using only the limited collection of these maps σ which we know to be present. This technique will be key to extending Theorem 10.1.1 in [26].

Independently of these potential later developments, our methods already allow us to establish new ubiquity results for many natural graphs and graph classes.

As a first concrete example, let G be the full grid, a graph not previously known to be ubiquitous. The *full grid* (briefly denoted by $\mathbb{Z}\square\mathbb{Z}$) is analogously defined as the half grid but with $\mathbb{Z} \times \mathbb{Z}$ as vertex set. The grid G is one-ended, and for any ray R in G, the graph $G - V(R)$ still has at most one end. Hence the unique end of G is non-linear, and so Theorem 10.1.2 has the following corollary:

Corollary 10.1.3. *The full grid is $\preccurlyeq$-ubiquitous.*

Using an argument similar in spirit to that of Halin [73], we also establish the following theorem in this paper:

Theorem 10.1.4. *Any connected minor of the half grid $\mathbb{N}\square\mathbb{Z}$ is $\preccurlyeq$-ubiquitous.*

Since every countable tree is a minor of the half grid, Theorem 10.1.4 implies that all countable trees are $\preccurlyeq$-ubiquitous, see Corollary 10.7.4. We remark that while all trees are ubiquitous with respect to the topological minor relation, [24], the problem whether all uncountable trees are $\preccurlyeq$-ubiquitous has remained open, and we hope to resolve this in a paper in preparation [27].

In a different direction, if G is any locally finite connected graph, then it is possible to show that $G\square\mathbb{Z}$ or $G\square\mathbb{N}$ either have nowhere-linear end structure, or are a subgraph of the half grid respectively. Hence, Theorems 10.1.2 and 10.1.4 together have the following corollary.

Theorem 10.1.5. *For every locally finite connected graph G, both $G\square\mathbb{Z}$ and $G\square\mathbb{N}$ are $\preccurlyeq$-ubiquitous.*

Finally, we will also show the following result about non-locally finite graphs. For $k \in \mathbb{N}$, we let the *k-fold dominated ray* be the graph DR_k formed by taking a ray together with k additional vertices, each of which we make adjacent to every vertex in the ray. For $k \leqslant 2$, DR_k is a minor of the half grid, and so ubiquitous by Theorem 10.1.4. In our last theorem, we show that DR_k is ubiquitous for all $k \in \mathbb{N}$.

Theorem 10.1.6. *The k-fold dominated ray DR_k is $\preccurlyeq$-ubiquitous for every $k \in \mathbb{N}$.*

The paper is structured as follows: In Section 10.2 we introduce some basic terminology for talking about minors. In Section 10.3 we introduce the concept of a *ray graph* and *linkages* between families of rays, which will help us to describe the structure of an end. In Sections 10.4 and 10.5 we introduce a pebble-pushing game which encodes possible linkages between families of rays and use this to give a sufficient condition for an end to contain a countable clique minor.

In Section 10.6 we re-introduce some concepts from [24] and show that we may assume that the G-minors in Γ are *concentrated* towards some end ϵ of Γ. In Section 10.7 we use the results of the previous section to prove Theorem 10.1.4 and finally in Section 10.8 we prove Theorem 10.1.2 and its corollaries.

10.2 Preliminaries

In our graph theoretic notation we generally follow the textbook of Diestel [43]. Given two graphs G and H the *cartesian product* $G \square H$ is a graph with vertex set $V(G) \times V(H)$ with an edge between (a, b) and (c, d) if and only if $a = c$ and $(b, d) \in E(H)$ or $(a, c) \in E(G)$ and $b = d$.

Definition 10.2.1. *A one-way infinite path is called a* ray *and a two-way infinite path is called a* double ray*.*

For a path or ray P and vertices $v, w \in V(P)$, let vPw denote the subpath of P with endvertices v and w. If P is a ray, let Pv denote the finite subpath of P between the initial vertex of P and v, and let vP denote the subray (or tail*) of P with initial vertex v.*

Given two paths or rays P and Q which are disjoint but for one of their endvertices, we write PQ for the concatenation of P and Q*, that is the path, ray or double ray $P \cup Q$. Moreover, if we concatenate paths of the form vPw and wQx, then we omit writing w twice and denote the concatenation by $vPwQx$.*

Definition 10.2.2 (Ends of a graph, cf. [43, Chapter 8])**.** *An* end *of an infinite graph Γ is an equivalence class of rays, where two rays R and S are equivalent if and only if there are infinitely many vertex disjoint paths between R and S in Γ. We denote by $\Omega(\Gamma)$ the set of ends of Γ.*

We say that a ray $R \subseteq \Gamma$ converges *(or* tends*) to an end ϵ of Γ if R is contained in ϵ. In this case we call R an* ϵ-ray*.*

Given an end $\epsilon \in \Omega(\Gamma)$ and a finite set $X \subseteq V(\Gamma)$ there is a unique component of $\Gamma - X$ which contains a tail of every ray in ϵ, which we denote by $C(X, \epsilon)$.

For an end $\epsilon \in \Gamma$ we define the degree *of ϵ in Γ as the supremum of all sizes of sets containing vertex disjoint ϵ-rays. If an end has finite degree, we call it* thin*. Otherwise, we call it* thick*.*

A vertex $v \in V(\Gamma)$ dominates *an end $\epsilon \in \Omega(\Gamma)$ if there is a ray $R \in \omega$ such that there are infinitely many v – R-paths in Γ that are vertex disjoint except from v.*

We will use the following two basic facts about infinite graphs.

Proposition 10.2.3. *[43, Proposition 8.2.1] An infinite connected graph contains either a ray or a vertex of infinite degree.*

Proposition 10.2.4. *[43, Exercise 8.19] A graph G contains a subdivided $K_{\aleph_0}$ as a subgraph if and only if G has an end which is dominated by infinitely many vertices.*

Definition 10.2.5 (Inflated graph, branch set)**.** *Given a graph G we say that a pair (H, φ) is an* inflated copy of G*, or an IG, if H is a graph and $\varphi \colon V(H) \to V(G)$ is a map such that:*

- *For every $v \in V(G)$ the* branch set *$\varphi^{-1}(v)$ induces a non-empty, connected subgraph of H;*
- *There is an edge in H between $\varphi^{-1}(v)$ and $\varphi^{-1}(w)$ if and only if $(v, w) \in E(G)$ and this edge, if it exists, is unique.*

When there is no danger of confusion we will simply say that H is an IG instead of saying that (H, φ) is an IG, and denote by $H(v) = \varphi^{-1}(v)$ the branch set of v.

Definition 10.2.6 (Minor)**.** *A graph G is a minor of another graph Γ, written $G \preccurlyeq \Gamma$, if there is some subgraph $H \subseteq \Gamma$ such that H is an inflated copy of G.*

Definition 10.2.7 (Extension of inflated copies)**.** *Suppose $G \subseteq G'$ as subgraphs, and that H is an IG and H' is an IG'. We say that H'* extends *H (or that H' is an extension of H) if $H \subseteq H'$ as subgraphs and $H(v) \subseteq H'(v)$ for all $v \in V(G) \cap V(G')$.*

Note that since $H \subseteq H'$, for every edge $(v, w) \in E(G)$, the unique edge between the branch sets $H'(v)$ and $H'(w)$ is also the unique edge between $H(v)$ and $H(w)$.

Definition 10.2.8 (Tidiness)**.** *An IG (H, φ) is called* tidy *if*

- *$H[\varphi^{-1}(v)]$ is a tree for all $v \in V(G)$;*
- *$H(v)$ is finite if $d_G(v)$ is finite.*

Note that every IG H contains a subgraph H' such that $(H', \varphi \restriction V(H'))$ is a tidy IG, although this choice may not be unique. In this paper we will always assume without loss of generality that each IG is tidy.

Definition 10.2.9 (Restriction)**.** *Let G be a graph, $M \subseteq G$ a subgraph of G, and let (H, φ) be an IG. The* restriction of H to M*, denoted by $H(M)$, is the IG given by $(H(M), \varphi')$ where $\varphi'^{-1}(v) = \varphi^{-1}(v)$ for all $v \in V(M)$ and $H(M)$ consists of union of the subgraphs of H induced on each branch set $\varphi^{-1}(v)$ for each $v \in V(M)$ together with the edge between $\varphi^{-1}(u)$ and $\varphi^{-1}(v)$ for each $(u, v) \in E(M)$.*

Note that if H is tidy, then $H(M)$ will be tidy. Given a ray $R \subseteq G$ and a tidy IG H in a graph Γ, the restriction $H(R)$ is a one-ended tree, and so every ray in $H(R)$ will share a tail. Later in the paper we will want to make this correspondence between rays in G and Γ more explicit, with use of the following definition:

Definition 10.2.10 (Pullback)**.** *Let G be a graph, $R \subseteq G$ a ray, and let H be a tidy IG. The* pullback of R to H *is the subgraph $H^{\downarrow}(R) \subseteq H$ where $H^{\downarrow}(R)$ is subgraph minimal such that $(H^{\downarrow}(R), \varphi \restriction V(H^{\downarrow}(R)))$ is an IM.*

Note that, since H is tidy, $H^{\downarrow}(R)$ is well defined. As well shall see, $H^{\downarrow}(R)$ will be a ray.

Lemma 10.2.11. *Let G be a graph and let H be a tidy IG. If $R \subseteq G$ is a ray, then the pullback $H^{\downarrow}(R)$ is also a ray.*

Proof. Let $R = x_1 x_2 \ldots$. For each integer $i \geqslant 1$ there is a unique edge $(v_i, w_i) \in E(H)$ between the branch sets $H(x_i)$ and $H(x_{i+1})$. By the tidiness assumption, $H(x_{i+1})$ induces a tree in H, and so there is a unique path $P_i \subset H(x_{i+1})$ from w_i to v_{i+1} in H.

By minimality of $H^{\downarrow}(R)$, it follows that $H^{\downarrow}(R)(x_1) = \{v_1\}$ and $H^{\downarrow}(R)(x_{i+1}) = V(P_i)$ for each $i \geqslant 1$. Hence $H^{\downarrow}(R)$ is a ray. □

10.3 The Ray Graph

Definition 10.3.1 (Ray graph)**.** *Given a finite family of disjoint rays $\mathcal{R} = (R_i : i \in I)$ in a graph Γ the* ray graph *$RG_\Gamma(\mathcal{R}) = RG_\Gamma(R_i : i \in I)$ is the graph with vertex set I and with an edge between i and j if there is an infinite collection of vertex disjoint paths from R_i to R_j in Γ which meet no other R_k. When the host graph Γ is clear from the context we will simply write $RG(\mathcal{R})$ for $RG_\Gamma(\mathcal{R})$.*

The following lemmas are simple exercises. For a family $\mathcal{R}$ of disjoint rays in G tending to the same end and $H \subseteq \Gamma$ being an IG the aim is to establish the following: if $\mathcal{S}$ is a family of disjoint rays in Γ which contains the pullback $H^{\downarrow}(R)$ of each $R \in \mathcal{R}$, then the subgraph of the ray graph $RG_\Gamma(\mathcal{S})$ induced on the vertices given by $\{H^{\downarrow}(R) : R \in \mathcal{R}\}$ is connected.

Lemma 10.3.2. *Let G be a graph and let $\mathcal{R} = (R_i \colon i \in I)$ be a finite family of disjoint rays in G. Then $RG_G(\mathcal{R})$ is connected if and only if all rays in $\mathcal{R}$ tend to a common end $\omega \in \Omega(G)$.*

Lemma 10.3.3. *Let G be a graph, $\mathcal{R} = (R_i \colon i \in I)$ be a finite family of disjoint rays in G and let H be an IG. If $\mathcal{R}' = (H^{\downarrow}(R_i) \colon i \in I)$ is the set of pullbacks of the rays in $\mathcal{R}$ in H, then $RG_G(\mathcal{R}) = RG_H(\mathcal{R}')$.*

Lemma 10.3.4. *Let G be a graph, $H \subseteq G$, $\mathcal{R} = (R_i \colon i \in I)$ be a finite disjoint family of rays in H and let $\mathcal{S} = (S_j \colon j \in J)$ be a finite disjoint family of rays in $G - V(H)$, where I and J are disjoint. Then $RG_H(\mathcal{R})$ is a subgraph of $RG_G(\mathcal{R} \cup \mathcal{S})\big[I\big]$. In particular, if all rays in $\mathcal{R}$ tend to a common end in H, then $RG_G(\mathcal{R} \cup \mathcal{S})\big[I\big]$ is connected.*

Recall that an end ω of a graph G is called *linear* if for every finite set $\mathcal{R}$ of at least three disjoint ω-rays in G we can order the elements of $\mathcal{R}$ as $\mathcal{R} = \{R_1, R_2, \ldots, R_n\}$ such that for each $1 \leqslant k < i < \ell \leqslant n$, the rays R_k and R_ℓ belong to different ends of $G - V(R_i)$.

Lemma 10.3.5. *An end ω of a graph G is linear if and only if the ray graph of every finite family of disjoint ω-rays is a path.*

Proof. For the forward direction suppose ω is linear and $\{R_1, R_2, \ldots, R_n\}$ converge to ω, with the order given by the definition of linear. It follows that there is no $1 \leqslant k < i < \ell \leqslant n$ such that (k, ℓ) is an edge in $RG(R_j \colon j \in [n])$. However, by Lemma 10.3.2 $RG(R_j \colon j \in [n])$ is connected, and hence it must be the path $12 \ldots n$.

Conversely, suppose that the ray graph of every finite family of ω-rays is a path. Then, every such family $\mathcal{R}$ can be ordered as $\{R_1, R_2, \ldots, R_n\}$ such that $RG(\mathcal{R})$ is the path $12 \ldots n$. It follows that, for each i, $(k, \ell) \notin E(RG(\mathcal{R}))$ whenever $1 \leqslant k < i < \ell \leqslant n-1$, and so by definition of $RG(\mathcal{R})$ there is no infinite collection of vertex disjoint paths from R_k to R_ℓ in $G - V(R_i)$. Therefore R_k and R_ℓ belong to different ends of $G - V(R_i)$. $\square$

Definition 10.3.6 (Tail of a ray after a set)**.** *Given a ray R in a graph G and a finite set $X \subseteq V(G)$ the* tail of R after X*, denoted by $T(R, X)$, is the unique infinite component of R in $G - X$.*

Definition 10.3.7 (Linkage of families of rays)**.** *Let $\mathcal{R} = (R_i \colon i \in I)$ and $\mathcal{S} = (S_j \colon j \in J)$ be families of disjoint rays of Γ, where the initial vertex of each R_i is denoted x_i. A family $\mathcal{P} = (P_i \colon i \in I)$ of paths in Γ is a* linkage *from $\mathcal{R}$ to $\mathcal{S}$ if there is an injective function $\sigma \colon I \to J$ such that*

- *Each P_i goes from a vertex $x_i' \in R_i$ to a vertex $y_{\sigma(i)} \in S_{\sigma(i)}$;*
- *The family $\mathcal{T} = (x_i R_i x_i' P_i y_{\sigma(i)} S_{\sigma(i)} \colon i \in I)$ is a collection of disjoint rays.*

We say that $\mathcal{T}$ is obtained by transitioning *from $\mathcal{R}$ to $\mathcal{S}$ along the linkage. We say the linkage $\mathcal{P}$* induces *the mapping σ. Given a vertex set $X \subseteq V(G)$ we say that the linkage is* after X *if $X \cap V(R_i) \subseteq V(x_i R_i x_i')$ for all $i \in I$ and no other vertex in X is used by $\mathcal{T}$. We say that a function $\sigma \colon I \to J$ is a* transition function *from $\mathcal{R}$ to $\mathcal{S}$ if for any finite vertex set $X \subseteq V(G)$ there is a linkage from $\mathcal{R}$ to $\mathcal{S}$ after X that induces σ.*

We will need the following lemma from [24], which asserts the existence of linkages.

Lemma 10.3.8 (Weak linking lemma)**.** *Let Γ be a graph, $\omega \in \Omega(\Gamma)$ and let $n \in \mathbb{N}$. Then for any two families $\mathcal{R} = (R_i \colon i \in [n])$ and $\mathcal{S} = (S_j \colon j \in [n])$ of vertex disjoint ω-rays and any finite vertex set $X \subseteq V(G)$, there is a linkage from $\mathcal{R}$ to $\mathcal{S}$ after X.*

10.4 A pebble-pushing game

Suppose we have a family of disjoint rays $\mathcal{R} = (R_i : i \in I)$ in a graph G and a subset $J \subseteq I$. Often we will be interested in which functions we can obtain as transition functions between $(R_i : i \in J)$ and $(R_i : i \in I)$. We can think of this as trying to 're-route' the rays $(R_i : i \in J)$ to a different set of $|J|$ rays in $(R_i : i \in I)$.

To this end, it will be useful to understand the following pebble-pushing game on a graph.

Definition 10.4.1 (Pebble-pushing game). *Let $G = (V, E)$ be a finite graph. For any fixed positive integer k we call a tuple $(x_1, x_2, \dots, x_k) \in V^k$ a* game state *if $x_i \neq x_j$ for all $i, j \in [k]$ with $i \neq j$.*

The pebble-pushing game *(on G) is a game played by a single player. Given a game state $Y = (y_1, y_2, \dots, y_k)$, we imagine k labelled pebbles placed on the vertices $(y_1, y_2, \dots, y_k)$. We move between game states by moving a pebble from a vertex to an adjacent vertex which does not contain a pebble, or formally, a Y*-move *is a game state $Z = (z_1, z_2 \dots, z_k)$ such that there is an $\ell \in [k]$ such that $y_\ell z_\ell \in E$ and $y_i = z_i$ for all $i \in [k] \setminus \{\ell\}$.*

Let $X = (x_1, x_2 \dots, x_k)$ be a game state. The X-pebble-pushing game *(on G) is a pebble-pushing game where we start with k labelled pebbles placed on the vertices $(x_1, x_2 \dots, x_k)$.*

We say a game state Y is achievable *in the X-pebble-pushing game if there is a sequence $(X_i : i \in [n])$ of game states for some $n \in \mathbb{N}$ such that $X_1 = X$, $X_n = Y$ and X_{i+1} is an X_i-move for all $i \in [n-1]$, that is, if it is a sequence of moves that pushes the pebbles from X to Y.*

A graph G is k-pebble-win *if Y is an achievable game state in the X-pebble-pushing game on G for every two game states X and Y.*

The following lemma shows that achievable game states on the ray graph $RG(\mathcal{R})$ yield transition functions from a subset of $\mathcal{R}$ to itself. Therefore, it will be useful to understand which game states are achievable, and in particular the structure of graphs on which there are unachievable game states.

Lemma 10.4.2. *Let Γ be a graph, $\omega \in \Omega(\Gamma)$, $m \geqslant k$ be positive integers and let $(S_j : j \in [m])$ be a family of disjoint rays in ω. For every achievable game state $Z = (z_1, z_2, \dots, z_k)$ in the $(1, 2, \dots, k)$-pebble-pushing game on $RG(S_j : j \in [m])$, the map σ defined via $\sigma(i) := z_i$ for every $i \in [k]$ is a transition function from $(S_i : i \in [k])$ to $(S_j : j \in [m])$.*

Proof. We first note that if σ is a transition function from $(S_i : i \in [k])$ to $(S_j : j \in [m])$ and τ is a transition function from $(S_i : i \in \sigma([k]))$ to $(S_j : j \in [m])$, then clearly $\tau \circ \sigma$ is a transition function from $(S_i : i \in [k])$ to $(S_j : j \in [m])$.

Hence, it will be sufficient to show the statement holds when σ is obtained from $(1, 2, \dots, k)$ by a single move, that is, there is some $t \in [k]$ and a vertex $\sigma(t) \notin [k]$ such that $\sigma(t)$ is adjacent to t in $RG(S_j : j \in [m])$ and $\sigma(i) = i$ for $i \in [k] \setminus \{t\}$.

So, let $X \subseteq V(G)$ be a finite set. We will show that there is a linkage from $(S_i : i \in [k])$ to $(S_j : j \in [m])$ after X that induces σ. By assumption there is an edge $(t, \sigma(t)) \in E(RG(S_j : j \in [m]))$. Hence, there is a path P between $T(S_t, X)$ and $T(S_{\sigma(t)}, X)$ which avoids X and all other S_j.

Then the family $\mathcal{P} = (P_1, P_2, \dots, P_k)$ where $P_t = P$ and $P_i = \emptyset$ for each $i \neq t$ is a linkage from $(S_i : i \in [k])$ to $(S_j : j \in [m])$ after X that induces σ. $\square$

We note that this pebble-pushing game is sometimes known in the literature as "permutation pebble motion" [87] or "token reconfiguration" [32]. Previous results have mostly focused on computational questions about the game, rather than the structural questions we are interested in, but we note that in [87] the authors give an algorithm that decides whether or not a graph is k-pebble-win, from which it should be possible to deduce the main result in this section, Lemma 10.4.9. However, since a direct derivation was shorter and self contained, we will not use their results. We present the following simple lemmas without proof.

Lemma 10.4.3. *Let G be a finite graph and X a game state.*

- *If Y is an achievable game state in the X-pebble-pushing game on G, then X is an achievable game state in the Y-pebble-pushing game on G.*

- *If Y is an achievable game state in the X-pebble-pushing game on G and Z is an achievable game state in the Y-pebble-pushing game on G, then Z is an achievable game state in the X-pebble-pushing game on G.*

Definition 10.4.4. *Let G be a finite graph and let $X = (x_1, x_2, \dots, x_k)$ be a game state. Given a permutation σ of $[k]$ let us write $X^\sigma = (x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(k)})$. We define the* pebble-permutation group *of (G, X) to be the set of permutations σ of $[k]$ such that X^σ is an achievable game state in the X-pebble-pushing game on G.*

Note that by Lemma 10.4.3, the pebble-permutation group of (G, X) is a subgroup of the symmetric group S_k.

Lemma 10.4.5. *Let G be a graph and let X be a game state. If Y is an achievable game state in the X-pebble-pushing game and σ is in the pebble-permutation group of Y, then σ is in the pebble-permutation group of X.*

Lemma 10.4.6. *Let G be a finite connected graph and let X be a game state. Then G is k-pebble-win if and only if the pebble-permutation group of (G, X) is S_k.*

Proof. Clearly, if the pebble-permutation group is not S_k then G is not k-pebble-win. Conversely, since G is connected, for any game states X and Y there is some τ such that Y^τ is an achievable game state in the X-pebble-pushing game, since we can move the pebbles to any set of k vertices, up to some permutation of the labels. We know by assumption that $X^{\tau^{-1}}$ is an achievable game state in the X-pebble-pushing game. Therefore, by Lemma 10.4.3 Y is an achievable game state in the X-pebble-pushing game. □

Lemma 10.4.7. *Let G be a finite connected graph and let $X = (x_1, x_2, \dots, x_k)$ be a game state. If G is not k-pebble-win, then there is a two colouring $c\colon X \to \{r, b\}$ such that both colour classes are non trivial and for all $i, j \in [k]$ with $c(x_i) = r$ and $c(x_j) = b$ the transposition (ij) is not in the pebble-permutation group.*

Proof. Let us draw a graph H on $\{x_1, x_2, \dots, x_k\}$ by letting (x_i, x_j) be an edge if and only if (ij) is in the pebble-permutation group of (G, X). It is a simple exercise to show that the pebble-permutation group of (G, X) is S_k if and only if H has a single component.

Since G is not k-pebble-win, we therefore know by Lemma 10.4.6 that there are at least two components in H. Let us pick one component C_1 and set $c(x) = r$ for all $x \in V(C_1)$ and $c(x) = b$ for all $x \in X \setminus V(C_1)$. □

Definition 10.4.8. *Given a graph G, a path $x_1 x_2 \dots x_m$ in G is a* bare path *if $d_G(x_i) = 2$ for all $2 \leqslant i \leqslant m - 1$.*

Lemma 10.4.9. *Let G be a finite connected graph with vertex set V which is not k-pebble-win and with $|V| \geqslant k + 2$. Then there is a bare path $P = p_1 p_2 \dots p_n$ in G such that $|V \setminus V(P)| \leqslant k$. Furthermore, either every edge in P is a bridge in G, or G is a cycle.*

Proof. Let $X = (x_1, x_2, \ldots, x_k)$ be a game state. Since G is not k-pebble-win, by Lemma 10.4.7 there is a two colouring $c\colon \{x_i\colon i \in [k]\} \to \{r, b\}$ such that both colour classes are non trivial and for all $i, j \in [k]$ with $c(x_i) = r$ and $c(x_j) = b$ the transposition (ij) is not in the pebble permutation group. Let us consider this as a three colouring $c\colon V \to \{r, b, 0\}$ where $c(v) = 0$ if $v \notin \{x_1, x_2, \ldots, x_k\}$.

For every achievable game state $Z = (z_1, z_2, \ldots, z_k)$ in the X-pebble-pushing game we define a three colouring c_Z given by $c_Z(z_i) = c(x_i)$ for all $i \in [k]$ and by $c_Z(v) = 0$ for all $v \notin \{z_1, z_2, \ldots, z_k\}$. We note that, for any achievable game state Z there is no $z_i \in c_Z^{-1}(r)$ and $z_j \in c_Z^{-1}(b)$ such that (ij) is in the pebble permutation group of (G, Z). Indeed, if it were, then by Lemma 10.4.3 $X^{(ij)}$ is an achievable game state in the X-pebble-pushing game, contradicting the fact that $c(x_i) = r$ and $c(x_j) = b$.

Since G is connected, for every achievable game state Z there is a path $P = p_1p_2 \ldots p_m$ in G with $c_Z(p_1) = r$, $c_Z(p_m) = b$ and $c_Z(p_i) = 0$ otherwise. Let us consider an achievable game state Z for which G contains such a path P of maximal length.

We first claim that there is no $v \notin P$ with $c_Z(v) = 0$. Indeed, suppose there is such a vertex v. Since G is connected there is some v–P path Q in G and so, by pushing pebbles towards v on Q, we can achieve a game state Z' such that $c_{Z'} = c_Z$ on P and there is a vertex v' adjacent to P such that $c_{Z'}(v') = 0$. Clearly v' cannot be adjacent to p_1 or p_m, since then we can push the pebble on p_1 or p_m onto v' and achieve a game state Z'' for which G contains a longer path than P with the required colouring. However, if v' is adjacent to p_ℓ with $2 \leqslant \ell \leqslant m-1$, then we can push the pebble on p_1 onto p_ℓ and then onto v', then push the pebble from p_m onto p_1 and finally push the pebble on v' onto p_ℓ and then onto p_m.

However, if $Z' = (z'_1, z'_2, \ldots, z'_k)$ with $p_1 = z'_i$ and $p_m = z'_j$, then above shows that (ij) is in the pebble-permutation group of (G, Z'). However, $c_{Z'}(z'_i) = c_Z(p_1) = r$ and $c_{Z'}(z'_j) = c_Z(p_m) = b$, contradicting our assumptions on $c_{Z'}$.

Next, we claim that each p_i with $3 \leqslant i \leqslant m-2$ has degree 2. Indeed, suppose first that p_i with $3 \leqslant i \leqslant m-2$ is adjacent to some other p_j with $1 \leqslant j \leqslant m$ such that p_i and p_j are not adjacent in P. Then it is easy to find a sequence of moves which exchanges the pebbles on p_1 and p_m, contradicting our assumptions on c_Z.

Suppose then that p_i is adjacent to a vertex v not in P. Then, $c_Z(v) \neq 0$, say without loss of generality $c_Z(v) = r$. However then, we can push the pebble on p_m onto p_{i-1}, push the pebble on v onto p_i and then onto p_m and finally push the pebble on p_{i-1} onto p_i and then onto v. As before, this contradicts our assumptions on c_Z.

Hence $P' = p_2p_3 \ldots p_{m-1}$ is a bare path in G, and since every vertex in $V - V(P')$ is coloured using r or using b, there are at most k such vertices.

Finally, suppose that there is some edge in P' which is not a bridge of G, and so no edge of P' is a bridge of G. We wish to show that G is a cycle. We first make the following claim:

Claim 10.4.10. *There is no achievable game state $W = (w_1, w_2, \ldots, w_k)$ such that there is a cycle $C = c_1c_2 \ldots c_rc_1$ and a vertex $v \notin C$ such that:*

- *There exist distinct positive integers i, j, s and t such that $c_W(c_i) = r$, $c_W(c_j) = b$ and $c_W(c_s) = c_W(c_t) = 0$;*
- *v adjacent to some $c_v \in C$.*

Proof of Claim 10.4.10. Suppose for a contradiction there exists such an achievable game state W. Since C is a cycle we may assume without loss of generality that $c_i = c_1, c_s = c_2 = c_v$, $c_t = c_3$ and $c_j = c_4$. If $c_W(v) = b$, then we can push the pebble at v to c_2 and then to c_3, push the pebble at c_1 to c_2 and then to v, and then push the pebble at c_3 to c_1. This contradicts our assumptions on c_W. The case where $c_W(v) = r$ is similar. Finally if $c_W(v) = 0$, then we can push the pebble at c_1 to c_2 and then to v, then push the pebble at c_4 to c_1, then push the pebble at v to c_2 and then to c_4. Again this contradicts our assumptions on c_W. $\square$

Since no edge of P' is a bridge, it follows that G contains a cycle C containing P'. If G is not a cycle, then there is a vertex $v \in V \setminus C$ which is adjacent to C. However by pushing the pebble on p_1 onto p_2 and the pebble on p_m onto p_{m-1}, which is possible since $|V| \geqslant k+2$, we achieve a game state Z' such that C and v satisfy the assumptions of the above claim, a contradiction. $\square$

10.5 Pebbly ends

Definition 10.5.1 (Pebbly)**.** *Let Γ be a graph and ω an end of Γ. We say ω is* pebbly *if for every $k \in \mathbb{N}$ there is an $n \geqslant k$ and a family $\mathcal{R} = (R_i : i \in [n])$ of disjoint rays in ω such that $RG(\mathcal{R})$ is k-pebble-win. If for some k there is no such family $\mathcal{R}$, we say ω is* not k-pebble-win.

The following is an immediate corollary of Lemma 10.4.9.

Corollary 10.5.2. *Let ω be an end of a graph Γ which is not k-pebble-win and let $\mathcal{R} = (R_i : i \in [m])$ be a family of $m \geqslant k+2$ disjoint rays in ω. Then there is a bare path $P = p_1p_2 \dots p_n$ in $RG(R_i : i \in [m])$ such that $|[m] \setminus V(P)| \leqslant k$. Furthermore, either each edge in P is a bridge in $RG(R_i : i \in [m])$, or $RG(R_i : i \in [m])$ is a cycle.*

Hence, if an end in Γ is not pebbly, then we have some constraint on the behaviour of rays towards this ends. In a later paper [26] we will investigate more precisely what can be said about the structure of the graph towards this end. For now, the following lemma allows us to easily find any countable graph as a minor of a graph with a pebbly end.

Lemma 10.5.3. *Let Γ be a graph and let $\omega \in \Omega(\Gamma)$ be a pebbly end. Then $K_{\aleph_0} \preccurlyeq \Gamma$.*

Proof. By assumption, there exists a sequence $\mathcal{R}_1, \mathcal{R}_2, \dots$ of families of disjoint ω-rays such that, for each $k \in \mathbb{N}$, $RG(\mathcal{R}_k)$ is k-pebble-win. Let us suppose that

$$\mathcal{R}_i = (R^i_1, R^i_2, \dots, R^i_{m_i}) \text{ for each } i \in \mathbb{N}.$$

Let us enumerate the vertices and edges of $K_{\aleph_0}$ by choosing some bijection $\sigma : \mathbb{N} \cup \mathbb{N}^{(2)} \to \mathbb{N}$ such that $\sigma(i,j) > \sigma(i), \sigma(j)$ for every $\{i,j\} \in \mathbb{N}^{(2)}$ and also $\sigma(1) < \sigma(2) < \cdots$. For each $k \in \mathbb{N}$ let G_k be the graph on vertex set $V_k = \{i \in \mathbb{N} : \sigma(i) \leqslant k\}$ and edge set $E_k = \{\{i,j\} \in \mathbb{N}^{(2)} : \sigma(i,j) \leqslant k\}$.

We will inductively construct subgraphs H_k of Γ such that H_k is an IG_k extending H_{k-1}. Furthermore for each $k \in \mathbb{N}$ if $V(G_k) = [n]$ then there will be tails $T_1, T_2, \dots, T_n$ of n distinct rays in $\mathcal{R}_n$ such that for every $i \in [n]$ the tail T_i meets H_k in a vertex of the branch set of i, and is otherwise disjoint from H_k. We will assume without loss of generality that T_i is a tail of R^n_i.

Since $\sigma(1) = 1$ we can take H_1 to be the initial vertex of R^1_1. Suppose then that $V(G_{n-1}) = [r]$ and we have already constructed H_{n-1} together with appropriate tails T_i of R^r_i for each $i \in [r]$. Suppose firstly that $\sigma^{-1}(n) = r+1 \in \mathbb{N}$.

Let $X = V(H_{n-1})$. There is a linkage from $(T_i : i \in [r])$ to $(R^{r+1}_1, R^{r+1}_2, \dots, R^{r+1}_r)$ after X by Lemma 10.3.8, and, after relabelling, we may assume this linkage induces the identity on $[r]$. Let us suppose the linkage consists of paths P_i from $x_i \in T_i$ to $y_i \in R^{r+1}_i$.

Since $X \cup \bigcup_i P_i \cup \bigcup_i T_i x_i$ is a finite set, there is some vertex y_{r+1} on R^{r+1}_{r+1} such that the tail $y_{r+1}R^{r+1}_{r+1}$ is disjoint from $X \cup \bigcup_i P_i \cup \bigcup_i T_i x_i$.

To form H_n we add the paths $T_i x_i \cup P_i$ to the branch set of each $i \leqslant r$ and set y_{r+1} as the branch set for $r+1$. Then H_n is an IG_n extending H_{n-1} and the tails $y_j R^{r+1}_j$ are as claimed.

Suppose then that $\sigma^{-1}(n) = \{u, v\} \in \mathbb{N}^{(2)}$ with $u, v \leqslant r$. We have tails T_i of R_i^r for each $i \in [r]$ which are disjoint from H_{n-1} apart from their initial vertices. Let us take tails T_j of R_j^r for each $j > r$ which are also disjoint from H_{n-1}. Since $RG(\mathcal{R}_r)$ is r-pebble-win, it follows that $RG(T_i\colon i \in [m_r])$ is also r-pebble-win. Furthermore, since by Lemma 10.3.2 $RG(T_i\colon i \in [m_r])$ is connected, there is some neighbour $w \in [m_r]$ of u in $RG(T_i\colon i \in [m_r])$.

Let us first assume that $w \notin [r]$. Since $RG(T_i\colon i \in [m_r])$ is r-pebble-win, the game state $(1, 2, \dots, v-1, w, v+1, \dots, r)$ is an achievable game state in the $(1, 2, \dots, r)$- pebble-pushing game and hence by Lemma 10.4.2 the function φ_1 given by $\varphi_1(i) = i$ for all $i \in [r] \setminus \{v\}$ and $\varphi_1(v) = w$ is a transition function from $(T_i\colon i \in [r])$ to $(T_i\colon i \in [m_r])$.

Let us take a linkage from $(T_i\colon i \in [r])$ to $(T_i\colon i \in [m_r])$ inducing φ_1 which is after $V(H_{n-1})$. Let us suppose the linkage consists of paths P_i from $x_i \in T_i$ to $y_i \in T_i$ for $i \neq v$ and P_v from $x_v \in T_v$ to $y_v \in T_w$. Let

$$X = V(H_{n-1}) \cup \bigcup_{i \in [r]} P_i \cup \bigcup_{i \in [r]} T_i x_i$$

Since u is adjacent to w in $RG(T_i\colon i \in [m_r])$ there is a path $\hat{P}$ between $T(T_u, X)$ and $T(T_w, X)$ which is disjoint from X and from all other T_i, say $\hat{P}$ is from $\hat{x} \in T_u$ to $\hat{y} \in T_w$.

Finally, since $RG(T_i\colon i \in [m_r])$ is r-pebble-win, the game state $(1, 2, \dots, r)$ is an achievable game state in the $(1, 2, \dots, v-1, w, v+1, \dots, r)$-pebble-pushing game and hence by Lemma 10.4.2 the function φ_2 given by $\varphi_2(i) = i$ for all $i \in [r] \setminus \{v\}$ and $\varphi_2(w) = v$ is a transition function from $(T_i\colon i \in [r] \setminus \{v\} \cup \{w\})$ to $(T_i\colon i \in [m_r])$.

Let us take a further linkage from $(T_i\colon i \in [r] \setminus \{v\} \cup \{w\})$ to $(T_i\colon i \in [m_r])$ inducing φ_2 which is after $X \cup \hat{P} \cup T_u\hat{x} \cup y_v T_w \hat{y}$. Let us suppose the linkage consists of paths P'_i from $x'_i \in T_i$ to $y'_i \in T_i$ for $i \in [r] \setminus \{v\}$ and P'_v from $x'_v \in T_w$ to $y'_v \in T_v$.

In the case that $w \in [r]$, $w < v$, say, the game state

$$(1, 2, \dots, w - 1, v, w + 1, \dots, v - 1, w, v + 1, \dots r)$$

is an achievable game state in the $(1, 2, \dots, r)$-pebble pushing-game and we get, by a similar argument, all $P_i, x_i, y_i, P'_i, x'_i, y'_i$ and $\hat{P}$.

We build H_n from H_{n-1} by adjoining the following paths:

- for each $i \neq v$ we add the path $T_i x_i P_i y_i T_i x'_i P'_i y'_i$ to H_{n-1}, adding the vertices to the branch set of i;
- we add $\hat{P}$ to H_{n-1}, adding the vertices of $V(\hat{P}) \setminus \{\hat{y}\}$ to the branch set of u;
- we add the path $T_v x_v P_v y_v T_w x'_v P'_v y'_v$ to H_{n-1}, adding the vertices to the branch set of v.

We note that, since $\hat{y} \in y_v T_w x'_v$ the branch sets for u and v are now adjacent. Hence H_n is an IG_n extending H_{n-1}. Finally the rays $y'_i T_i$ for $i \in [r]$ are appropriate tails of the used rays of $\mathcal{R}_r$. □

As every countable graph is a subgraph of $K_{\aleph_0}$, a graph with a pebbly end contains every countable graph as a minor. Thus, as $\aleph_0 G$ is countable, if G is countable, we obtain the following corollary:

Corollary 10.5.4. *Let Γ be a graph with a pebbly end ω and let G be a countable graph. Then $\aleph_0 G \preccurlyeq \Gamma$.*

10.6 G-tribes and concentration of G-tribes towards an end

To show that a given graph G is $\preccurlyeq$-ubiquitous, we shall assume that $nG \preccurlyeq \Gamma$ holds for every $n \in \mathbb{N}$ an show that this implies $\aleph_0 G \preccurlyeq \Gamma$. To this end we use the following notation for such collections of nG in Γ, most of which we established in [24].

Definition 10.6.1 (G-tribes). *Let G and Γ be graphs.*

- *A* G-tribe in Γ (with respect to the minor relation) *is a family $\mathcal{F}$ of finite collections F of disjoint subgraphs H of Γ such that each* member *H of $\mathcal{F}$ is an IG.*
- *A G-tribe $\mathcal{F}$ in Γ is called* thick*, if for each $n \in \mathbb{N}$ there is a* layer *$F \in \mathcal{F}$ with $|F| \geqslant n$; otherwise, it is called* thin.
- *A G-tribe $\mathcal{F}'$ in Γ is a* G-subtribe [1] *of a G-tribe $\mathcal{F}$ in Γ, denoted by $\mathcal{F}' \preccurlyeq \mathcal{F}$, if there is an injection $\Psi \colon \mathcal{F}' \to \mathcal{F}$ such that for each $F' \in \mathcal{F}'$ there is an injection $\varphi_{F'} \colon F' \to \Psi(F')$ such that $V(H') \subseteq V(\varphi_{F'}(H'))$ for each $H' \in F'$. The G-subtribe $\mathcal{F}'$ is called* flat*, denoted by $\mathcal{F}' \subseteq \mathcal{F}$, if there is such an injection Ψ satisfying $F' \subseteq \Psi(F')$.*
- *A thick G-tribe $\mathcal{F}$ in Γ is* concentrated at an end ϵ *of Γ, if for every finite vertex set X of Γ, the G-tribe $\mathcal{F}_X = \{F_X \colon F \in \mathcal{F}\}$ consisting of the layers $F_X = \{H \in F : H \not\subseteq C(X,\epsilon)\} \subseteq F$ is a thin subtribe of $\mathcal{F}$. It is* strongly concentrated at ϵ *if additionally, for every finite vertex set X of Γ, every member H of $\mathcal{F}$ intersects $C(X,\epsilon)$.*

We note that, every thick G-tribe $\mathcal{F}$ contains a thick subtribe $\mathcal{F}'$ such that every $H \in \bigcup \mathcal{F}$ is a tidy IG. We will use the following lemmas from [24].

Lemma 10.6.2 (Removing a thin subtribe, [24, Lemma 5.2]). *Let $\mathcal{F}$ be a thick G-tribe in Γ and let $\mathcal{F}'$ be a thin subtribe of $\mathcal{F}$, witnessed by $\Psi \colon \mathcal{F}' \to \mathcal{F}$ and $(\varphi_{F'} \colon F' \in \mathcal{F}')$. For $F \in \mathcal{F}$, if $F \in \Psi(\mathcal{F}')$, let $\Psi^{-1}(F) = \{F'_F\}$ and set $\hat{F} = \varphi_{F'_F}(F'_F)$. If $F \notin \Psi(\mathcal{F}')$, set $\hat{F} = \emptyset$. Then*

$$\mathcal{F}'' := \{F \setminus \hat{F} \colon F \in \mathcal{F}\}$$

is a thick flat G-subtribe of $\mathcal{F}$.

Lemma 10.6.3 (Pigeon hole principle for thick G-tribes, [24, Lemma 5.3]). *Suppose for some $k \in \mathbb{N}$, we have a k-colouring $c \colon \bigcup \mathcal{F} \to [k]$ of the members of some thick G-tribe $\mathcal{F}$ in Γ. Then there is a monochromatic, thick, flat G-subtribe $\mathcal{F}'$ of $\mathcal{F}$.*

Note that, in the following lemma, it is necessary that G is connected, so that every member of the G-tribe is a connected graph.

Lemma 10.6.4 ([24, Lemma 5.4]). *Let G be a connected graph and Γ a graph containing a thick G-tribe $\mathcal{F}$. Then either $\aleph_0 G \preccurlyeq \Gamma$, or there is a thick flat subtribe $\mathcal{F}'$ of $\mathcal{F}$ and an end ϵ of Γ such that $\mathcal{F}'$ is concentrated at ϵ.*

Lemma 10.6.5 ([24, Lemma 5.5]). *Let G be a connected graph and Γ a graph containing a thick G-tribe $\mathcal{F}$ concentrated at an end ϵ of Γ. Then the following assertions hold:*

1. *For every finite set X, the component $C(X,\epsilon)$ contains a thick flat G-subtribe of $\mathcal{F}$.*
2. *Every thick subtribe $\mathcal{F}'$ of $\mathcal{F}$ is concentrated at ϵ, too.*

[1] When G is clear from the context we will often refer to a G-subtribe as simply a subtribe.

Lemma 10.6.6. *Let G be a connected graph and Γ a graph containing a thick G-tribe $\mathcal{F}$ concentrated at an end $\epsilon \in \Omega(\Gamma)$. Then either $\aleph_0 G \preccurlyeq \Gamma$, or there is a thick flat subtribe of $\mathcal{F}$ which is strongly concentrated at ϵ.*

Proof. Suppose that no thick flat subtribe of $\mathcal{F}$ is strongly concentrated at ϵ. We construct an $\aleph_0 G \preccurlyeq \Gamma$ by recursively choosing disjoint IGs $H_1, H_2, \dots$ in Γ as follows: Having chosen $H_1, H_2, \dots, H_n$ such that for some finite set X_n we have

$$H_i \cap C(X_n, \epsilon) = \emptyset$$

for all $i \in [n]$, then by Lemma 10.6.5(1), there is still a thick flat subtribe $\mathcal{F}'_n$ of $\mathcal{F}$ contained in $C(X_n, \epsilon)$. Since by assumption, $\mathcal{F}'_n$ is not strongly concentrated at ϵ, we may pick $H_{n+1} \in \mathcal{F}'_n$ and a finite set $X_{n+1} \supseteq X_n$ with $H_{n+1} \cap C(X_{n+1}, \epsilon) = \emptyset$. Then the union of all the H_i is an $\aleph_0 G \preccurlyeq \Gamma$. □

The following lemma will show that we can restrict ourself to thick G-tribes which are concentrated at thick ends.

Lemma 10.6.7. *Let G be a connected graph and Γ a graph containing a thick G-tribe $\mathcal{F}$ concentrated at an end $\epsilon \in \Omega(\Gamma)$ which is thin. Then $\aleph_0 G \preccurlyeq \Gamma$.*

Proof. Since ϵ is thin, we know by Proposition 10.2.4 that only finitely many vertices dominate ϵ. Deleting these yields a subgraph of Γ in which there is still a thick G-tribe concentrated at ϵ. Hence we may assume without loss of generality that ϵ is not dominated by any vertex in Γ.

Let $k \in \mathbb{N}$ be the degree of ϵ. By [69, Corollary 5.5] there is a sequence of vertex sets $(S_n : n \in \mathbb{N})$ such that:

- $|S_n| = k$,
- $C(S_{n+1}, \epsilon) \subseteq C(S_n, \epsilon)$, and
- $\bigcap_{n \in \mathbb{N}} C(S_n, \epsilon) = \emptyset$.

Suppose there is a thick subtribe $\mathcal{F}'$ of $\mathcal{F}$ which is strongly concentrated at ϵ. For any $F \in \mathcal{F}'$ there is an $N_F \in \mathbb{N}$ such that $H \setminus C(S_{N_F}, \epsilon) \neq \emptyset$ for all $H \in F$ by the properties of the sequence. Furthermore, since $\mathcal{F}'$ is strongly concentrated, $H \cap C(S_{N_F}, \epsilon) \neq \emptyset$ as well for each $H \in F$.

Let $F \in \mathcal{F}'$ be such that $|F| > k$. Since G is connected, so is H, and so from the above it follows that $H \cap S_{N_F} \neq \emptyset$ for each $H \in F$, contradicting the fact that $|S_{N_F}| = k < |F|$. Thus $\aleph_0 G \preccurlyeq \Gamma$ by Lemma 10.6.6. □

Note that, whilst concentration is hereditary for subtribes, strong concentration is not. However if we restrict to *flat* subtribes, then strong concentration is a hereditary property.

Let us show see how ends of the members of a strongly concentrated tribe relate to ends of the host graph Γ. Let G be a connected graph and $H \subseteq \Gamma$ an IG. By Lemmas 10.3.2 and 10.3.4, if $\omega \in \Omega(G)$ and R_1 and $R_2 \in \omega$ then the pullbacks $H^{\downarrow}(R_1)$ and $H^{\downarrow}(R_2)$ belong to the same end $\omega' \in \Omega(\Gamma)$. Hence, H determines for every end $\omega \in G$ a *pullback end* $H(\omega) \in \Omega(\Gamma)$. The next lemma is where we need to use the assumption that G is locally finite.

Lemma 10.6.8. *Let G be a locally finite connected graph and Γ a graph containing a thick G-tribe $\mathcal{F}$ strongly concentrated at an end $\epsilon \in \Omega(\Gamma)$ where every member is a tidy IG. Then either $\aleph_0 G \preccurlyeq \Gamma$, or there is a flat subtribe $\mathcal{F}'$ of $\mathcal{F}$ such that for every $H \in \bigcup \mathcal{F}'$ there is an end $\omega_H \in \Omega(G)$ such that $H(\omega_H) = \epsilon$.*

Proof. Since G is locally finite and every $H \in \bigcup \mathcal{F}$ is tidy, the branch sets $H(v)$ are finite for each $v \in V(G)$. If ϵ is dominated by infinitely many vertices, then we know by Proposition 10.2.4 that Γ contains a topological $K_{\aleph_0}$ minor, in which case $\aleph_0 G \preccurlyeq \Gamma$, since every locally finite connected graph is countable. If this is not the case, then there is some $k \in \mathbb{N}$ such that ϵ is dominated by k vertices and so for every $F \in \mathcal{F}$ at most k of the $H \in F$ contain vertices which dominate ϵ in Γ. Therefore, there is a thick flat subtribe $\mathcal{F}'$ of $\mathcal{F}$ such that no $H \in \bigcup \mathcal{F}'$ contains a vertex dominating ϵ in Γ. Note that $\mathcal{F}'$ is still strongly concentrated at ϵ, and every branch set of every $H \in \bigcup \mathcal{F}'$ is finite.

Since $\mathcal{F}'$ is strongly concentrated at ϵ, for every finite vertex set X of Γ every $H \in \bigcup \mathcal{F}'$ intersects $C(X, \epsilon)$. By a standard argument, since H as a connected infinite graph does not contain a vertex dominating ϵ in Γ, instead H contains a ray $R_H \in \epsilon$.

Since each branch set $H(v)$ is finite, R_H meets infinitely many branch sets. Let us consider the subgraph $K \subseteq G$ consisting of all the edges (v, w) such that R_H uses an edge between $H(v)$ and $H(w)$. Note that, since there is a edge in H between $H(v)$ and $H(w)$ if and only if $(v, w) \in E(G)$, K is well-defined and connected.

K is then an infinite connected subgraph of a locally finite graph, and as such contains a ray S_H in G. Since the edges between $H(v)$ and $H(w)$, if they exist, were unique, it follows that the pullback $H^{\downarrow}(S_H)$ of S_H has infinitely many edges in common with R_H, and so tends to ϵ in Γ. Therefore, if S_H tends to ω_H in $\Omega(G)$, then $H(\omega_H) = \epsilon$.

□

10.7 Ubiquity of minors of the half grid

Here, and in the following, we denote by $\mathbb{H}$ the infinite, one-ended, cubic hexagonal half grid (see Figure 10.2). The following theorem of Halin is one of the cornerstones of infinite graph theory.

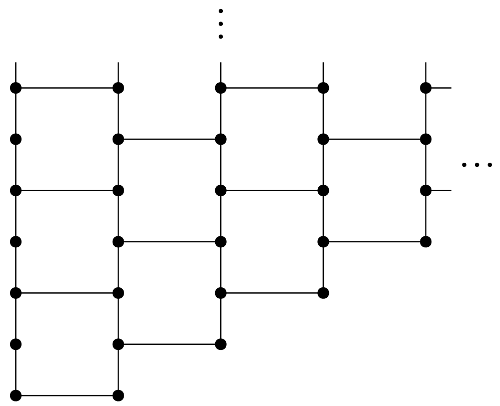

Figure 10.2: The hexagonal half grid $\mathbb{H}$.

Theorem 10.7.1 (Halin, see [43, Theorem 8.2.6]). *Whenever a graph Γ contains a thick end, then $\mathbb{H} \leqslant \Gamma$.* □

In [73], Halin used this result to show that every topological minor of $\mathbb{H}$ is ubiquitous with respect to the topological minor relation $\leqslant$. In particular, trees of maximum degree 3 are ubiquitous with respect to $\leqslant$.

However, the following argument, which is a slight adaptation of Halin's, shows that every connected minor of $\mathbb{H}$ is ubiquitous with respect to the minor relation. In particular, the dominated ray, the dominated double ray, and all countable trees are ubiquitous with respect to the minor relation.

The main difference to Halin's original proof is that, since he was only considering locally finite graphs, he was able to assume that the host graph Γ was also locally finite.

Lemma 10.7.2 ([73, (4) in Section 3])**.** $\aleph_0\mathbb{H}$ *is a topological minor of* $\mathbb{H}$.

Theorem 10.1.4. *Any connected minor of the half grid* $\mathbb{N}\square\mathbb{Z}$ *is* $\preccurlyeq$*-ubiquitous.*

Proof. Suppose $G \preccurlyeq \mathbb{N}\square\mathbb{Z}$ is a minor of the half grid, and Γ is a graph such that $nG \preccurlyeq \Gamma$ for each $n \in \mathbb{N}$. By Lemma 10.6.4 we may assume there is an end ϵ of Γ and a thick G-tribe $\mathcal{F}$ which is concentrated at ϵ. By Lemma 10.6.7 we may assume that ϵ is thick. Hence $\mathbb{H} \leqslant \Gamma$ by Theorem 10.7.1, and with Lemma 10.7.2 we obtain

$$\aleph_0 G \preccurlyeq \aleph_0(\mathbb{N}\square\mathbb{Z}) \preccurlyeq \aleph_0\mathbb{H} \leqslant \mathbb{H} \leqslant \Gamma. \qquad \square$$

Lemma 10.7.3. $\mathbb{H}$ *contains every countable tree as a minor.*

Proof. It is easy to see that the infinite binary tree T_2 embeds into $\mathbb{H}$ as a topological minor. It is also easy to see that countably regular tree T_∞ where every vertex has infinite degree embeds into T_2 as a minor. And obviously, every countable tree T is a subgraph of T_∞. Hence we have

$$T \subseteq T_\infty \preccurlyeq T_2 \leqslant \mathbb{H}$$

from which the result follows. $\square$

Corollary 10.7.4. *All countable trees are ubiquitous with respect to the minor relation.*

Proof. This is an immediate consequence of Lemma 10.7.3 and Theorem 10.1.4. $\square$

10.8 Proof of main results

Lemma 10.8.1. *Let* ϵ *be a non-pebbly end of* Γ *and let* $\mathcal{F}$ *be a* G*-tribe such that for every* $H \in \bigcup\mathcal{F}$ *there is an end* $\omega_H \in \Omega(G)$ *such that* $H(\omega_H) = \epsilon$*. Then there is a thick flat subtribe* $\mathcal{F}'$ *such that* ω_H *is linear for every* $H \in \bigcup\mathcal{F}'$ *.*

Proof. Let $\mathcal{F}'$ be the flat subtribe of $\mathcal{F}$ given by $\mathcal{F}' = \{F' \colon F \in \mathcal{F}\}$ with

$$F' = \{H \colon H \in F \text{ and } \omega_H \text{ is not linear}\}.$$

Suppose for a contradiction that $\mathcal{F}'$ is thick. Then, there is some $F \in \mathcal{F}$ which contains $k+2$ disjoint IGs, $H_1, H_2, \ldots, H_{k+2}$, where k is such that ϵ is not k-pebble-win. By assumption ω_{H_i} is not linear for each i, and so for each i there is a family of disjoint rays $\{R^i_1, R^i_2, \ldots, R^i_{m_i}\}$ in G tending to ω_{H_i} whose ray graph in G is not a path. Let

$$\mathcal{S} = (H_i^{\downarrow}(R^i_j) \colon i \in [k+2], j \in [m_i]).$$

By construction $\mathcal{S}$ is a disjoint family of rays which tend to ϵ in Γ and by Lemma 10.3.3 and Lemma 10.3.4 $RG_\Gamma(\mathcal{S})$ contains disjoint subgraphs $K_1, K_2, \ldots, K_{k+2}$ such that $K_i \cong RG_G(R^i_j \colon j \in [m_i])$. However, by Corollary 10.5.2, there is a set X of vertices of size at most k such that $RG_\Gamma(\mathcal{S}) - X$ is a bare path P. However, then some $K_i \subseteq P$ is a path, a contradiction.

Since $\mathcal{F}$ is the union of $\mathcal{F}'$ and $\mathcal{F}''$ where $\mathcal{F}'' = \{F'' \colon F \in \mathcal{F}\}$ with

$$F'' = \{H \colon H \in F \text{ and } \omega_H \text{ is linear}\},$$

it follows that $\mathcal{F}''$ is thick. $\square$

Theorem 10.1.2. *Every locally finite connected graph with nowhere-linear end structure is $\preccurlyeq$-ubiquitous.*

Proof. Let Γ be a graph such that $nG \preccurlyeq \Gamma$ holds for every $n \in \mathbb{N}$. Hence, Γ contains a thick G-tribe $\mathcal{F}$. By Lemmas 10.6.4 and 10.6.6 we may assume that $\mathcal{F}$ is strongly concentrated at an end ϵ of Γ and so by Lemma 10.6.8 we may assume that for every $H \in \bigcup \mathcal{F}$ there is an end $\omega_H \in \Omega(G)$ such that $H(\omega_H) = \epsilon$.

Since ω_H is not linear for each $H \in \bigcup \mathcal{F}$, it follows by Lemma 10.8.1 that ϵ is pebbly, and hence by Corollary 10.5.4 $\aleph_0 G \preccurlyeq \Gamma$. □

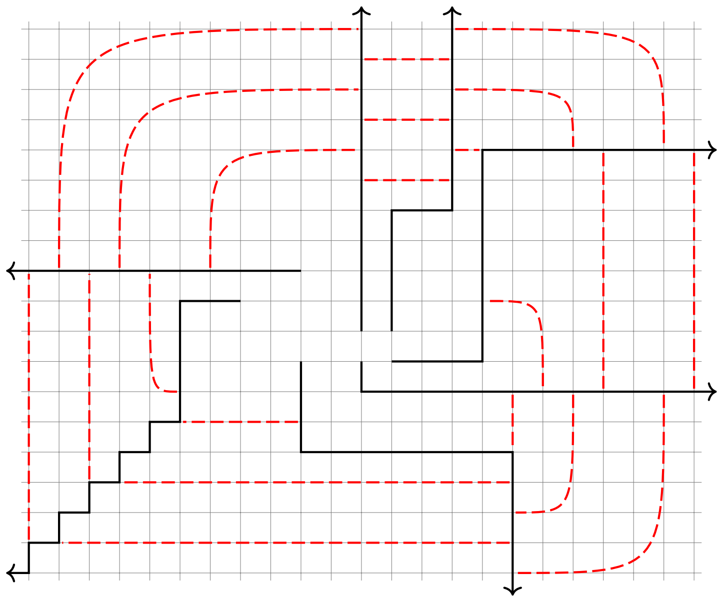

Figure 10.3: The ray graphs in the full grid are cycles.

Corollary 10.1.3. *The full grid is $\preccurlyeq$-ubiquitous.*

Proof. Let G be the full grid. Since $G - R$ has at most one end for any ray $R \in G$, by Lemma 10.3.2 the ray graph $RG(\mathcal{R})$ is 2-connected for any finite family of three or more rays. Hence, by Theorem 10.1.2 G is $\preccurlyeq$-ubiquitous □

Remark. *In fact, every ray graph in the full grid is a cycle (see Figure 10.3).*

Theorem 10.1.5. *For every locally finite connected graph G, both $G\Box\mathbb{Z}$ and $G\Box\mathbb{N}$ are $\preccurlyeq$-ubiquitous.*

Proof. If G is a path or a ray, then $G\Box\mathbb{Z}$ is a subgraph of the half grid $\mathbb{N}\Box\mathbb{Z}$ and thus $\preccurlyeq$-ubiquitous by Theorem 10.1.4. If G is a double ray then $G\Box\mathbb{Z}$ is the full grid and thus $\preccurlyeq$-ubiquitous by Corollary 10.1.3. Otherwise let G' be a finite connected subgraph of G which is not a path. For any end ω of $G\Box\mathbb{Z}$ there is a ray R of $\mathbb{Z}$ such that all rays of the form $\{v\}\Box R$ for $v \in V(G)$ go to ω. But then G' is a subgraph of $RG_{G\Box\mathbb{Z}}((\{v\}\Box R)_{v\in V(G')})$, so this ray-graph is not a path, hence by Lemma 10.3.5 $G\Box\mathbb{Z}$ has nowhere-linear end structure and is therefore $\preccurlyeq$-ubiquitous by Theorem 10.1.2. □

Finally let us prove Theorem 10.1.6. Recall that for $k \in \mathbb{N}$ let DR_k denote the graph formed by taking a ray R together with k vertices $v_1, v_2, \ldots, v_k$ adjacent to every vertex in R. We shall need the following strengthening of Proposition 10.2.3.

A *comb* is a union of a ray R with infinitely many disjoint finite paths, all having precisely their first vertex on R. The last vertices of these paths are the *teeth* of the comb.

Proposition 10.8.2. *[43, Proposition 8.2.2] Let U be an infinite set of vertices in a connected graph G. Then G either contains a comb with all teeth in U or a subdivision of an infinite star with all leaves in U.*

Theorem 10.1.6. *The k-fold dominated ray DR_k is $\preccurlyeq$-ubiquitous for every $k \in \mathbb{N}$.*

Proof. Note that if $k \leqslant 2$ then DR_k is a minor of the half grid, and hence ubiquity follows from Theorem 10.1.4.

Suppose then that $k \geqslant 3$ and Γ is a graph which contains a thick DR_k-tribe $\mathcal{F}$ each of whose members is tidy. By Lemma 10.6.6 we may assume that there is an end ϵ of Γ such that $\mathcal{F}$ is concentrated at ϵ. If there are infinitely many vertices dominating ϵ, then $\aleph_0 DR_k \preccurlyeq K_{\aleph_0} \leqslant \Gamma$ holds by Proposition 10.2.4. So we may assume that only finitely many vertices dominate ϵ. By taking a thick subtribe if necessary, we may assume that no member of $\mathcal{F}$ contains such a vertex.

As before, if we can show that ϵ is pebbly, then we will be done by Corollary 10.5.4. So suppose for a contradiction that ϵ is not r-pebble-win for some $r \in \mathbb{N}$.

Let R be the ray as stated in the definition of DR_k and let $v_1, v_2, \ldots, v_k \in V(DR_k)$ be the vertices adjacent to each vertex of R. For each $H \in \bigcup \mathcal{F}$ and each $i \in [k]$ we have the $H(v_i)$ is a connected subgraph of Γ. Let U be the set of all vertices in $H(v_i)$ which are the endpoint of some edge in H between $H(v_i)$ and $H(w)$ with $w \in R$. Since v_i dominated R, U is infinite, and so by Proposition 10.8.2 $H(v_i)$ either contains a comb with all teeth in U or a subdivision of an infinite star with all leaves in U. However in the latter case the centre of the star would dominate ϵ, and so each $H(v_i)$ contains such a comb, whose spine we denote by $R_{H,i}$. Let $R_H = H^{\downarrow}(R)$ be the pullback of the ray R in H. Now we set $\mathcal{R}_H = (R_{H,1}, R_{H,2}, \ldots, R_{H,k}, R_H)$.

Since $R_{H,i}$ is the spine of a comb, all of whose leaves are in U, it follows that in the graph $RG_H(\mathcal{R}_H)$ each $R_{H,i}$ is adjacent to R_H. Hence $RG_H(\mathcal{R}_H)$ contains a vertex of degree $k \geqslant 3$.

There is some layer $F \in \mathcal{F}$ of size $\ell \geqslant r+1$, say $F = (H_i \colon i \in [\ell])$. For every $i \in [r+1]$ we set $\mathcal{R}_{H_i} = (R_{H_i,1}, R_{H_i,2}, \ldots, R_{H_i,k}, R_{H_i})$. Let us now consider the family of disjoint rays

$$\mathcal{R} = \bigcup_{i=1}^{r+1} \mathcal{R}_{H_i}.$$

By construction $\mathcal{R}$ is a family of disjoint rays which tend to ϵ in Γ and by Lemma 10.3.3 and Lemma 10.3.4 $RG_\Gamma(\mathcal{R})$ contains $r+1$ vertices whose degree is at least $k \geqslant 3$. However, by Corollary 10.5.2, there is a vertex set X of size at most r such that $RG_\Gamma(\mathcal{R}) - X$ is a bare path P. But then some vertex whose degree is at least 3 is contained in the bare path, a contradiction. $\square$

Chapter 11

Ubiquity of locally finite graphs with extensive tree decompositions

11.1 Introduction

Given a graph G and some relation $\lhd$ between graphs we say that G is $\lhd$-*ubiquitous* if whenever Γ is a graph such that $nG \lhd \Gamma$ for all $n \in \mathbb{N}$, then $\aleph_0 G \lhd \Gamma$, where αG is the disjoint union of α many copies of G. A classic result of Halin [71, Satz 1] says that the ray is $\subseteq$-ubiquitous, where $\subseteq$ is the subgraph relation. That is, any graph which contains arbitrarily large collections of vertex-disjoint rays must contain an infinite collection of vertex-disjoint rays. Later, Halin showed that the double ray is also $\subseteq$-ubiquitous [72].

However, not all graphs are $\subseteq$-ubiquitous, and in fact even trees can fail to be $\subseteq$-ubiquitous (see for example [130]). The question of ubiquity for classes of graphs has also been considered for other graph relations. In particular, whilst there are still reasonably simple examples of graphs which are not $\leqslant$-ubiquitous (see [91, 7]), where $\leqslant$ is the topological minor relation, it was shown by Andreae that all rayless countable graphs [9] and all locally finite trees [8] are $\leqslant$-ubiquitous. The latter result was recently extended to the class of all trees by the authors [24].

In [13] Andreae initiated the study of ubiquity of graphs with respect to the minor relation, $\preccurlyeq$. He constructed a graph which was not $\preccurlyeq$-ubiquitous, however the construction relied on the existence of a counterexample to the well-quasi-ordering of infinite graphs under the minor relation, for which only examples of very large cardinality are known [122]. In particular, the question whether there exists a countable graph which is not $\preccurlyeq$-ubiquitous remains open.

Andreae conjectured that at least all *locally finite* graphs, those with all degrees finite, should be $\preccurlyeq$-ubiquitous.

Conjecture 9.1.1. *[The Ubiquity Conjecture] Every locally finite connected graph is $\preccurlyeq$-ubiquitous.*

In [14] Andreae proved that his conjecture holds for a large class of locally finite graphs. The exact definition of this class is technical, but in particular his result implies the following.

Theorem 11.1.1 (Andreae, [14, Corollary 1])**.** *Let G be a locally finite, connected graph with finitely many ends such that every block of G is finite. Then G is $\preccurlyeq$-ubiquitous.*

Theorem 11.1.2 (Andreae, [14, Corollary 2])**.** *Let G be a locally finite, connected graph of finite tree-width such that every block of G is finite. Then G is $\preccurlyeq$-ubiquitous.*

Note, in particular, that if G is such a graph, then the degree of every end in G must be one.[1] In this paper we will extend Andreae's approach to prove that an even larger class of locally finite graphs is $\preccurlyeq$-ubiquitous, removing the assumption of finite blocks. Again, the exact definition of this class will be technical, but in particular it will imply the following results, extending Theorems 11.1.1 and 11.1.2:

Theorem 11.1.3. *Let G be a locally finite, connected graph with finitely many ends such that every end of G has finite degree. Then G is $\preccurlyeq$-ubiquitous.*

Theorem 11.1.4. *Every locally finite, connected graph of finite tree-width is $\preccurlyeq$-ubiquitous.*

The proof uses in an essential way some known results about the well-quasi-ordering of graphs under the minor relation, including Thomas' result [123] that graphs of bounded tree width are well-quasi-ordered under the minor relation. Our methods, building on Andreae's, give a blueprint by which stronger results about the well-quasi-ordering of graphs can be used to prove the ubiquity of larger classes of graphs. A more precise discussion of this connection will be given in Section 11.10.

In Section 11.2 we will give a sketch of the key ideas in the proof, at the end of which we will give a short overview of the structure of the paper.

11.2 Proof sketch

To give a flavour of the main ideas involved in the proof, let's begin by considering the case of a locally finite connected graph G with a single end ω, where ω has finite degree d (this means that there is a family $(A_i : 1 \leqslant i \leqslant d)$ of d disjoint rays in ω, but no family of more than d such rays). Our construction will exploit the fact that graphs of this kind have a very particular structure. More precisely, there is a tree-decomposition $(S, \mathcal{V})$ of G, where $S = s_0s_1s_2\dots$ is a ray and such that, if we denote V_{s_n} by V_n and $\bigcup_{l \geqslant n} V_l$ by G_n for each n, the following holds:

1. each V_n is finite;
2. every vertex of G appears in only finitely many V_n;
3. all the A_i begin in V_0, and
4. for each $m \geqslant 1$ there are infinitely many $n > m$ such that G_m is a minor of G_n, in such a way that for any edge e of G_m and any $i \leqslant d$, e is an edge of A_i if and only if the edge representing it in this minor is.

Property 4 seems rather strong, and the reason it can always be achieved has to do with the well-quasi-ordering of finite graphs. For details of how this works, see Section 11.5. The skeptical reader who does not yet see how to achieve this may consider the argument in this section as showing ubiquity simply for graphs G with a decomposition of the above kind.

Now we suppose that we are given some graph Γ such that $nG \preccurlyeq \Gamma$ for each n, and we wish to show that $\aleph_0 G \preccurlyeq \Gamma$. Consider a G-minor H in Γ. Any ray R of G can then be expanded to a ray $H(R)$ in the copy H of G in Γ, and since G only has one end, all rays $H(R)$ go to the same end ϵ_H of Γ; we shall say that H *goes to* the end ϵ_H.

We now show that we can suppose without loss of generality that all G-minors go to the same end ϵ of Γ. For suppose that there are two G-minors H and H' with $\epsilon_H \neq \epsilon_{H'}$. Since G is locally finite, we may assume that all branch sets of H and H' are finite. Thus there is a finite set X such that each of H and H' only uses vertices from one component of $\Gamma - X$. In

[1] A precise definitions of the ends of a graph and their degree can be found in Section 11.3.

any $(|X| + 2n)G$-minor of Γ, only at most $|X|$ of the G-minors involved can meet X, and each of the remaining $2n$ must be included in some component of $G - X$. Without loss of generality at most n of them are in the component that meets H, and so $\Gamma - H$ has an nG-minor.

Thus there is a G-minor H_0 of Γ such that $\Gamma_1 := \Gamma - H_0$ still has an nG-minor for each n. If there are two G-minors going to different ends of Γ_1 then we may as above find a G-minor H_1 of Γ_1 such that $\Gamma_2 := \Gamma_1 - H_1$ has an nG-minor for any n. Proceeding in this way we either find infinitely many disjoint G-minors $H_0, H_1, H_2, \ldots$, giving an $\aleph_0 G$-minor, or else after finitely many steps we find a subgraph Γ_k of Γ which has an nG-minor for any n and in which all G-minors go to the same end ϵ.

So from now on we will assume that all G-minors of Γ go to the same end ϵ. From any G-minor H we obtain rays $H(A_i)$ corresponding to our marked rays A_i in G. We will call this collection of rays the *bundle* of rays given by H.

Our aim now is to build up an $\aleph_0 G$-minor of Γ recursively. At stage n we hope to construct n disjoint $G[\bigcup_{m \leqslant n} V_m]$-minors $H_1^n, H_2^n, \ldots H_n^n$, such that for each such H_m^n there is a family $(R_{m,i}^n : i \leqslant k)$ of disjoint rays to ϵ, where the path in H_m^n corresponding to the initial segment of the ray A_i in $\bigcup_{m \leqslant n} G_m$ is an initial segment of $R_{m,i}^n$, but these rays are otherwise disjoint from the various H_l^n and from each other. We aim to do this in such a way that each H_m^n extends all previous H_m^l for $l \leqslant n$, so that at the end of our construction we can obtain infinitely many disjoint G-minors as $(\bigcup_{n \geqslant m} H_m^n : m \in \mathbb{N})$. The rays chosen at later stages need not bear any relation to those chosen at earlier stages; we just need them to exist so that there is some hope of continuing the construction.

We will again refer to the families $(R_{m,i}^n : i \leqslant k)$ of rays starting at the various H_m^n as the *bundles* of rays from those H_m^n.

The rough idea for getting from the n^{th} to the $n + 1^{\text{st}}$ stage of this construction is now as follows: we choose a very large family $\mathcal{H}$ of disjoint G-minors in Γ. We throw away all those which meet any previous H_m^n and we consider the family of rays corresponding to the A_i in the remaining minors. Then it is possible to find a collection of paths transitioning from the $R_{m,i}^n$ from stage n onto these new rays. Precisely what we need is captured in the following definition, which also introduces some helpful terminology for dealing with such transitions:

Definition 11.2.1 (Linkage of families of rays)**.** *Let $\mathcal{R} = (R_i \colon i \in I)$ and $\mathcal{S} = (S_j \colon j \in J)$ be families of disjoint rays, where the initial vertex of each R_i is denoted x_i. A family of paths $\mathcal{P} = (P_i \colon i \in I)$, is a* linkage *from $\mathcal{R}$ to $\mathcal{S}$ if there is an injective function $\sigma \colon I \to J$ such that*

- *Each P_i goes from a vertex $x'_i \in R_i$ to a vertex $y_{\sigma(i)} \in S_{\sigma(i)}$;*

- *The family $\mathcal{T} = (x_i R_i x'_i P_i y_{\sigma(i)} S_{\sigma(i)} : i \in I)$ is a collection of disjoint rays.*[2] *We write $\mathcal{R} \circ_{\mathcal{P}} \mathcal{S}$ for the family $\mathcal{T}$ as well $R_i \circ_{\mathcal{P}} \mathcal{S}$ for the ray in $\mathcal{T}$ with initial vertex x_i.*

We say that $\mathcal{T}$ is obtained by transitioning *from $\mathcal{R}$ to $\mathcal{S}$ along the linkage. We say the linkage $\mathcal{P}$* induces *the mapping σ. We further say that $\mathcal{P}$* links *$\mathcal{R}$ to $\mathcal{S}$. Given a set X we say that the linkage is* after *X if $X \cap R_i \subseteq x_i R_i x'_i$ for all $i \in I$ and no other point in X is used by $\mathcal{T}$.*

Thus our aim is to find a linkage from the $R^n_{m,i}$ to the new rays after all the H^n_m. That this is possible is guaranteed by the following lemma from [24]:

Lemma 11.2.2 (Weak linking lemma [24, Lemma 4.3])**.** *Let Γ be a graph and $\omega \in \Omega(\Gamma)$. Then for any collections $\mathcal{R} = (R_1, \ldots, R_n)$ and $\mathcal{S} = (S_1, \ldots, S_n)$ of vertex disjoint rays in ω and any finite set X of vertices, there is a linkage from $\mathcal{R}$ to $\mathcal{S}$ after X.*

The aim is now to use property 4 of our tree decomposition of G to find copies of V_{n+1} sufficiently far along the new rays that we can stick them on to our H^n_m to obtain suitable H^{n+1}_m. There are two difficulties at this point in this argument. The first is that, as well as extending the existing H^n_m to H^{n+1}_m we also need to introduce H^{n+1}_{n+1}. To achieve this, we ensure that one of the G-minors in $\mathcal{H}$ is disjoint from all the paths in the linkage, so that we may take an initial segement of it as H^{n+1}_{n+1}. This is possible because of a slight strengthening of the linking lemma above; see [24, Lemma 4.4] or 9.4.4 for a precise statement.

A more serious difficulty is that in order to stick the new V_{n+1} onto H^n_m we need the following property:

> For each of the bundles corresponding to an H^n_m, the rays in the bundle are linked precisely to the rays in the bundle coming from some $H \in \mathcal{H}$. This happens in such a way that each $R^n_{m,i}$ is linked to $H(A_i)$. $(*)$

Thus we need a great deal of control over which rays get linked to which. We can keep track of which rays are linked to which as follows:

Definition 11.2.3 (Transition function)**.** *Let $\mathcal{R} = (R_i : i \in I)$ and $\mathcal{S} = (S_j : j \in J)$ be families of disjoint rays, where the initial vertex of each R_i is denoted x_i. We say that a function $\sigma : I \to J$ is a* transition function *from $\mathcal{R}$ to $\mathcal{S}$ if for any finite set X of vertices there is a linkage from $\mathcal{R}$ to $\mathcal{S}$ after X that induces σ.*

So our aim is to find a transition function assigning new rays to the R^n_m so as to achieve $(*)$. One reason for expecting this to be possible is that the new rays all go to the same end, and so they are joined up by many paths. We might hope to be able to use these paths to move between the rays, allowing us some control over which rays are linked to which. The structure of possible jumps is captured by a graph whose vertex set is the set of rays:

Definition 11.2.4 (Ray graph)**.** *Given a finite family of disjoint rays $\mathcal{R} = (R_i : i \in I)$ in a graph Γ the* ray-graph*, $RG_\Gamma(\mathcal{R}) = RG_\Gamma(R_i : i \in I)$ is the graph with vertex set I and with an edge between i and j if there is an infinite collection of vertex disjoint paths from R_i to R_j which meet no other R_k. When the host graph Γ is clear from the context we will simply write $RG(\mathcal{R})$ for $RG_\Gamma(\mathcal{R})$.*

[2] Where we use the notation as in [43], see also Definition 11.3.3.

Unfortunately, the collection of possible transition functions can be rather limited. Consider, for example, the case of families of disjoint rays in the grid. Any such family has a natural cyclic order, and any transition function must preserve this cyclic order. This paucity of transition functions is reflected in the sparsity of the ray graphs, which are all just cycles.

In Sections 11.6 and 11.7 we therefore carefully analyse the possibilities for how the ray graphs and transition functions associated to a given thick[3] end may look. We find that there are just 3 possibilities.

The easiest case is that in which the rays to the end are very joined up, in the sense that any injective function between two families of rays is a transition function. This case was already dealt with in [25]. The second possibility is that which we saw above for the grid: all ray graphs are cycles, and all transition functions between them preserve the cyclic order. The third possibility is that all ray graphs consist of a path together with a bounded number of further junk vertices, where these junk vertices are hanging at the ends of the paths (formally: all interior vertices on this *central paths* have degree 2 in the ray graph). In this case, the transition functions must preserve the linear order along the paths.

The second and third cases can be dealt with using similar ideas, so we will focus on the third one here.

The structure of the ray graphs and transition functions can be used to get around the problem discussed above, by slightly strengthening the properties required of the rays in the recursive construction. More precisely, we want that the ray graph of a slightly larger family $\mathcal{R}$ of disjoint rays, consisting of the $R^n_{m,i}$ and some extra 'junk' rays, should have all the $R^n_{m,i}$ on the central path, arranged in such a way that for each n and m the $R^n_{m,i}$ are consecutive in order from $R^n_{m,1}$ to $R^n_{m,k}$.

Of course, in order that this is possible we must first ensure that the A_i are arranged in order so that for every n we can find n disjoint G-minors H such that there is some ray graph in which, for each H, the rays $H(A_i)$ appear in order along the central path. Since there are only finitely many possible orders, there must be an order with this property.

Then our extra order assumptions ensure that, by transitioning between rays using edges of the ray graph, we can modify the linkage so that $(*)$ holds.

There is one last subtle difficulty which we have to address, once more relating to the fact that we want to introduce a new H^{n+1}_{n+1} together with its private bundle of rays corresponding to its copies of A_i's, disjoint from all the other H^{n+1}_m and their bundles. Recall that the strong linking lemma allows us to find a linkage which avoids one of the G-minors in $\mathcal{H}$, but this linkage may not have property $(*)$. We can modify it to one satisfying $(*)$ by diverting the rays along some of the paths between the new rays. But then some of the rays through which we divert may be forced to intersects the rays emanating from H^{n+1}_{n+1}, if these rays from H^{n+1}_{n+1} lie between rays from the same bundle of some H^n_m.

However, we can get around this by using the paths between the rays in $\mathcal{R}$ to jump between them *before* the linkage, so as to rearrange which bundles make use of (the tails of) which rays. More precisely, we first take a large but finite set of paths between the rays which is rich enough to allow us to rearrange which bundles end up where as much as possible. We collect these together in a *transition box*. Only then do we choose the linkage from $\mathcal{R}$ to the rays from $\mathcal{H}$, and we make sure that this linkage is after the transition box. Then, when we later see how the bundles should be arranged in order that the rays emanating from H^{n+1}_{n+1} do not appear between rays from the same bundle, we can go back and perform a suitable rearrangement within the transition box, see Figure 11.1.

This completes the sketch of the proof that locally finite graphs with a single end of finite degree are ubiquitous. Our results in this paper are for a more general class of graphs, but one

[3]An end is *thick* if there are infinitely many disjoint rays to it.

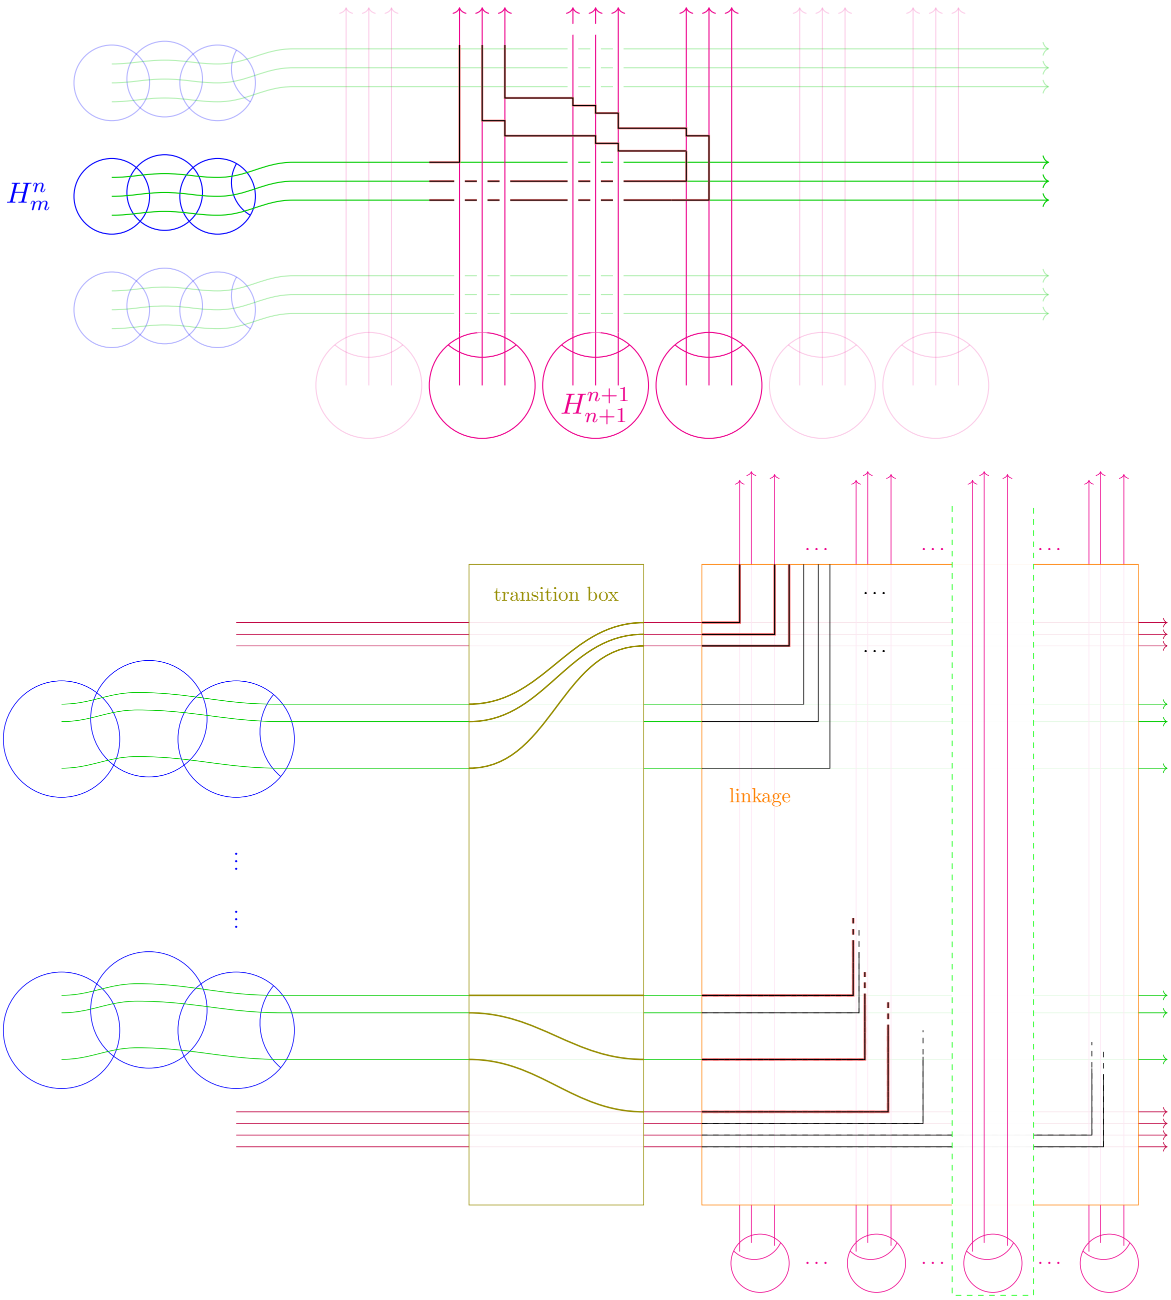

Figure 11.1: The transitioning strategy between the old and new bundles.

which is chosen to ensure that arguments of the kind outlined above will work for them. Hence we still need a tree-decomposition with properties similar to (1)-(4) from our ray-decomposition above. Tree decompositions with these properties are called *extensive*, and the details can be found in Section 11.4.

However, certain aspects of the sketch above must be modified to allow for the fact that we are now dealing with graphs G with multiple, indeed possibly infinitely many, ends. For any end δ of G and any G-minor H of Γ, all rays $H(R)$ with R in δ belong to the same end $H(\delta)$ of Γ. But for different values of δ, the ends $H(\delta)$ may well be different.

So there is no hope of finding a single end ϵ of Γ to which all rays in all G-minors converge.

Nevertheless, we can still find an end ϵ towards which the G-minors are *concentrated*, in the sense that for any finite X there are arbitrarily large families of G-minors in the same component of $G - X$ as ϵ. See Section 10.6 for details. In that section we introduce the term *tribe* for a collection of arbitrarily large families of disjoint G-minors.

The recursive construction will work pretty much as before, in that at each step n we will again have embedded G^n-minors for some large finite part G^n of G, together with a number of rays to ϵ corresponding to some canonical rays going to certain ends δ of G.

In order for this to work, we need some consistency about which $H(\delta)$ are equal to ϵ and which are not. It is clear that for any finite set Δ of ends of G there is some subset Δ' such that there is a tribe of G-minors H converging to ϵ with the property that the set of δ in Δ with $H(\delta) = \epsilon$ is Δ'. This is because there are only finitely many options for this set. But if G has infinitely many ends, there is no reason why we should be able to do this for all ends of G at once.

Our solution is to keep track of only finitely many ends of G at any stage in the construction, and to maintain at each stage a tribe concentrated towards ϵ which is consistent for all these finitely many ends. Thus in our construction consistency of questions such as which ends δ of G converge to ϵ or of the proper linear order in the ray graph of the families of canonical rays to those ends is achieved dynamically during the construction, rather than being fixed in advance. The ideas behind this dynamic process have already been used successfully in our earlier paper [24], where they appear in slightly simpler circumstances.

The paper is structured as follows. In Section 10.2 we give precise definitions of some of the basic concepts we will be using, and prove some of their fundamental properties. In Section 11.4 we introduce extensive tree decompositions and in Section 11.5 we illustrate that many locally finite graphs can be given such decompositions. Sections 11.6 and 11.7 are devoted to the possible collections of ray graphs and transition functions between them which can occur in a thick end. In Section 10.6 we introduce the notion of tribes and of their concentration towards an end and begin building some tools for the main recursive construction, which is given in Section 9.6. We conclude with a discussion of the future outlook in Section 11.10.

11.3 Preliminaries

In this paper we follow the convention that 0 is not an element of the set $\mathbb{N}$ of natural numbers.

For a graph $G = (V, E)$ and $W \subseteq V$ we write $G[W]$ for the induced subgraph. For two vertices v, w of a connected graph G, we write $\text{dist}(v, w)$ for the edge-length of a shortest $v - w$ path. A path $P = v_0 v_1 \dots v_n$ in a graph G is called a *bare path* if $\deg_G(v_i) = 2$ for all inner vertices v_i for $0 < i < n$.

11.3.1 Rays and ends

Definition 11.3.1 (Rays and initial vertices of rays)**.** *A one-way infinite path is called a* ray *and a two-way infinite path is called a* double ray*. For a ray R let* $\text{init}(R)$ *denote the* initial vertex *of R, that is the unique vertex of degree* 1 *in R. For a set $\mathcal{R}$ of rays let* $\text{init}(\mathcal{R})$ *denote the set of initial vertices of the rays in $\mathcal{R}$.*

Definition 11.3.2 (Tail of a ray)**.** *Given a ray R in a graph G and a finite set $X \subseteq V(G)$ the* tail of R after X*, $T(R, X)$, is the unique infinite component of R in $G - X$.*

Definition 11.3.3 (Concatenation of paths and rays)**.** *For a path or ray P and vertices $v, w \in V(P)$, let vPw denote the subpath of P with endvertices v and w, and $\mathring{v}P\mathring{w}$ for the subpath strictly between v and w. If P is a ray, let Pv denote the finite subpath of P between the initial* *vertex of P and v, and let vP denote the subray (or* tail*) of P with initial vertex v. Similarly, we write $P\mathring{v}$ and $\mathring{v}P$ for the corresponding paths without the vertex v.*

Given two paths or rays P and Q which which intersect in a single vertex only, which is an endvertex in both P and Q, we write PQ for the concatenation of P and Q, *that is the path, ray or double ray $P \cup Q$. Moreover, if we concatenate paths of the form vPw and wQx, then we omit writing w twice and denote the concatenation by $vPwQx$.*

For a ray $R = r_0r_1\ldots$ let R^- denote the tail r_1R of R starting at r_1. Given a set $\mathcal{R}$ of rays let $\mathcal{R}^-$ denote the set $\{R^-\colon R \in \mathcal{R}\}$

Definition 11.3.4 (Ends of a graph, cf. [43, Chapter 8])**.** *An* end *of an infinite graph Γ is an equivalence class of rays, where two rays R and S are equivalent if and only if there are infinitely many vertex disjoint paths between R and S in Γ. We denote by $\Omega(\Gamma)$ the set of ends of Γ.*

We say that a ray $R \subseteq \Gamma$ converges *(or* tends*) to an end ϵ of Γ if R is contained in ϵ. In this case we call R an* ϵ-ray.

Given an end $\epsilon \in \Omega(\Gamma)$ and a finite set $X \subseteq V(\Gamma)$ there is a unique component of $\Gamma - X$ which contains a tail of every ray in ϵ, which we denote by $C(X, \epsilon)$.

For an end $\epsilon \in \Gamma$ we define the degree *of ϵ in Γ, denoted by $\deg(\epsilon) \in \mathbb{N} \cup \{\infty\}$, as the largest cardinality of a collection of vertex disjoint ϵ-rays. An end with finite/infinite degree is called* thin/thick.

11.3.2 Inflated copies of graphs

Definition 11.3.5 (Inflated graph, branch set)**.** *Given a graph G we say that a pair (H, φ) is an* inflated copy of G *or an IG if H is a graph and $\varphi\colon V(H) \to V(G)$ is a map such that:*

- *For every $v \in V(G)$ the* branch set $\varphi^{-1}(v)$ *induces a non-empty, connected subgraph of H;*
- *There is an edge in H between $\varphi^{-1}(v)$ and $\varphi^{-1}(w)$ if and only if $(v, w) \in E(G)$ and this edge, if it exists, is unique.*

When there is no danger of confusion we will simply say that H is an IG instead of saying that (H, φ) is an IG, and denote by $H(v) = \varphi^{-1}(v)$ the branch set of v.

Definition 11.3.6 (Minor)**.** *A graph G is a minor of another graph Γ, written $G \preccurlyeq \Gamma$, if there is some subgraph $H \subseteq \Gamma$ such that H is an inflated copy of G.*

Definition 11.3.7 (Extension of inflated copies)**.** *Suppose $G \subseteq G'$ as subgraphs, and that H is an IG and H' is an IG'. We say that H'* extends *H (or that H' is an extension of H) if $H \subseteq H'$ as subgraphs and $H(v) \subseteq H'(v)$ for all $v \in V(G) \cap V(G')$.*

If H' is an extension of H and $X \subset V(G)$ is such that $H'(x) = H(x)$ for every $x \in X$ then we say H' is an extension of H fixing X.

Note that since $H \subseteq H'$, for every edge $(v, w) \in E(G)$, the unique edge between the branch sets $H'(v)$ and $H'(w)$ is also the unique edge between $H(v)$ and $H(w)$.

Definition 11.3.8 (Tidiness)**.** *An IG (H, φ) is called* tidy *if*

- *$H[\varphi^{-1}(v)]$ is a tree for all $v \in V(G)$;*
- *$H(v)$ is finite if $d_G(v)$ is finite.*

Note that every IG H contains a subgraph H' such that $(H', \varphi \restriction V(H'))$ is a tidy IG, although this choice may not be unique. In this paper we will always assume without loss of generality that each IG is tidy.

Definition 11.3.9 (Restriction). *Let G be a graph, $M \subseteq G$ a subgraph of G, and let (H, φ) be an IG. The* restriction of H to M, *denoted by $H(M)$, is the IG given by $(H(M), \varphi')$ where $\varphi'^{-1}(v) = \varphi^{-1}(v)$ for all $v \in V(M)$ and $H(M)$ consists of union of the subgraphs of H induced on each branch set $\varphi^{-1}(v)$ for each $v \in V(M)$ together with the edge between $\varphi^{-1}(u)$ and $\varphi^{-1}(v)$ for each $(u, v) \in E(M)$.*

Note that if H is tidy, then $H(M)$ will be tidy. Given a ray $R \subseteq G$ and a tidy IG H in a graph Γ, the restriction $H(R)$ is a one-ended tree, and so every ray in $H(R)$ will share a tail. Later in the paper we will want to make this correspondence between rays in G and Γ more explicit, with use of the following definition:

Definition 11.3.10 (Pullback). *Let G be a graph, $R \subseteq G$ a ray, and let H be a tidy IG. The* pullback of R to H *is the subgraph $H^{\downarrow}(R) \subseteq H$ where $H^{\downarrow}(R)$ is subgraph minimal such that $(H^{\downarrow}(R), \varphi \upharpoonright V(H^{\downarrow}(R)))$ is an IM.*

Note that, since H is tidy, $H^{\downarrow}(R)$ is well defined. As well shall see, $H^{\downarrow}(R)$ will be a ray.

Lemma 11.3.11. *Let G be a graph and let H be a tidy IG. If $R \subseteq G$ is a ray, then the pullback $H^{\downarrow}(R)$ is also a ray.*

Proof. Let $R = x_1 x_2 \ldots$. For each integer $i \geqslant 1$ there is a unique edge $(v_i, w_i) \in E(H)$ between the branch sets $H(x_i)$ and $H(x_{i+1})$. By the tidiness assumption, $H(x_{i+1})$ induces a tree in H, and so there is a unique path $P_i \subset H(x_{i+1})$ from w_i to v_{i+1} in H.

By minimality of $H^{\downarrow}(R)$, it follows that $H^{\downarrow}(R)(x_1) = \{v_1\}$ and $H^{\downarrow}(R)(x_{i+1}) = V(P_i)$ for each $i \geqslant 1$. Hence $H^{\downarrow}(R)$ is a ray. □

Definition 11.3.12. *Let G be a graph, $\mathcal{R}$ be a family of disjoint rays in G and let H be a tidy IG. We will write $H^{\downarrow}(\mathcal{R})$ for the family $(H^{\downarrow}(R) \colon R \in \mathcal{R})$.*

Definition 11.3.13. *For an end ω of G and $H \subset \Gamma$ a tidy IG, we denote by $H(\omega)$ the unique end of Γ containing all rays $H^{\downarrow}(R)$ for $R \in \omega$.*

It is an easy check that if two rays R and S in G are equivalent, then also $H^{\downarrow}(R)$ and $H^{\downarrow}(S)$ are rays (Lemma 11.3.11) which are equivalent in H, and hence also equivalent in Γ.

11.3.3 Transitional linkages and the strong linking lemma

Definition 11.3.14. *We say a linkage is* transitional *if the function which it induces between the corresponding ray graphs is a transition function.*

Lemma 11.3.15. *Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. Then for any collections $\mathcal{R} = (R_1, \ldots, R_n)$ and $\mathcal{S} = (S_1, \ldots, S_n)$ of ϵ-rays in Γ there is a finite set X such that every linkage after X is transitional.*

Proof. By definition, for every function $\sigma \colon [n] \to [n]$ which is not a transition function from $\mathcal{R}$ to $\mathcal{S}$ there is a finite set $X_\sigma \subseteq V(\Gamma)$ such that there is no linkage from $\mathcal{R}$ to $\mathcal{S}$ after X_σ which induces σ. If we let Φ be the set of σ which are not transition functions then the set $X := \bigcup_{\sigma \in \Phi} X_\sigma$ satisfies the conclusion of the lemma. □

In addition to Lemma 11.2.2 we will also need the following stronger linking lemma, which is a slight modification of [24, Lemma 4.4]:

Lemma 11.3.16 (Strong linking lemma). *Let Γ be a graph and $\omega \in \Omega(\Gamma)$. Let X be a finite set of vertices, $n \in \mathbb{N}$, and $\mathcal{R} = (R_i \colon i \in [n])$ a family of vertex disjoint rays in ω. Let $x_i = init(R_i)$ and $x_i' = init(T(R_i, X))$. Then there is a finite number $N = N(\mathcal{R}, X)$ with the following property: For every collection $(H_j \colon j \in [N])$ of vertex disjoint subgraphs of Γ, all disjoint from X and each including a specified ray S_j in ω, there is a $j \in [N]$ and a transitional linkage $\mathcal{P} = (P_i \colon i \in [n])$ from $\mathcal{R}$ to $(S_j \colon j \in [N])$ which is after X and such that the family*

$$\mathcal{T} = \left(x_i R_i x_i' P_i y_{\sigma(i)} S_{\sigma(i)} \colon i \in [n]\right)$$

avoids H_j.

Proof. Let $Y \subseteq V(\Gamma)$ be a finite set as in Lemma 11.3.15. We apply the strong linking lemma established in [24, Lemma 4.4] to the set $X \cup Y$ to obtain this version of the strong linking lemma. □

Lemma and Definition 11.3.17. *Let Γ be a graph, $\epsilon \in \Omega(\Gamma)$, $X \subseteq V(\Gamma)$ be finite, and let $\mathcal{R} = (R_i \colon i \in I_1)$, $\mathcal{S} = (S_i \colon i \in I_2)$ be two finite families of disjoint ϵ-rays with $|I_1| \leqslant |I_2|$. Then there is a finite subgraph $Y \subseteq C(X, \epsilon)$ such that for any transition function σ between $\mathcal{R}$ and $\mathcal{S}$ there is a linkage $\mathcal{P}_\sigma$ from $\mathcal{R}$ to $\mathcal{S}$ inducing σ with $\bigcup \mathcal{P}_\sigma \subseteq \Gamma[Y]$.*

We call such a set Y a transition box between $\mathcal{R}$ and $\mathcal{S}$ (after X).

Proof. Let $\sigma : I_1 \to I_2$ be a transition function between $\mathcal{R}$ and $\mathcal{S}$. By definition there is a linkage $\mathcal{P}_\sigma$ from $\mathcal{R}$ to $\mathcal{S}$ after X which induces σ. Note that, since $\mathcal{P}_\sigma$ is after X, it follows that $\bigcup \mathcal{P}_\sigma \subseteq C(X, \epsilon)$.

Let Φ be the set of all transition functions between $\mathcal{R}$ and $\mathcal{S}$ and let $Y = \bigcup_{\sigma \in \Phi} \mathcal{P}_\sigma$. Then Y is a transition box between $\mathcal{R}$ and $\mathcal{S}$ (after X). □

Remark and Definition 11.3.18. *Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. Let $\mathcal{R}_1$, $\mathcal{R}_2$, $\mathcal{R}_3$ be finite families of disjoint ϵ-rays, $\mathcal{P}_1$ a transitional linkage from $\mathcal{R}_1$ to $\mathcal{R}_2$ and $\mathcal{P}_2$ a transitional linkage from $\mathcal{R}_2$ to $\mathcal{R}_3$ after $\bigcup \mathcal{P}_2$.*

1. *$\mathcal{P}_2$ is also a transitional linkage from $(\mathcal{R}_1 \circ_{\mathcal{P}_1} \mathcal{R}_2)$ to $\mathcal{R}_3$.*

2. *The linkage from $\mathcal{R}_1$ to $\mathcal{R}_3$ yielding the rays $(\mathcal{R}_1 \circ_{\mathcal{P}_1} \mathcal{R}_2) \circ_{\mathcal{P}_2} \mathcal{R}_3$, which we call the* concatenation $\mathcal{P}_1 + \mathcal{P}_2$ of $\mathcal{P}_1$ and $\mathcal{P}_2$ *is transitional.*

The following lemmas are simple exercises.

Lemma 11.3.19. *Let $(R_i \colon i \in I)$ be a disjoint finite family of ϵ-rays, then the ray graph $RG(R_i \colon i \in I)$ is connected. Also, if R_i' is a tail of R_i for each $i \in I$, then $RG(R_i \colon i \in I) = RG(R_i' \colon i \in I)$.* □

Lemma 11.3.20 ([25, Lemma 3.4]). *Let G be a graph, $H \subseteq G$, $\mathcal{R} = (R_i \colon i \in I)$ be a finite disjoint family of rays in H and let $\mathcal{S} = (S_j \colon j \in J)$ be a finite disjoint family of rays in $G - V(H)$, where I and J are disjoint. Then $RG_H(\mathcal{R})$ is a subgraph of $RG_G(\mathcal{R} \cup \mathcal{S})\big[I\big]$.* □

11.4 Extensive tree-decompositions and self minors

The purpose of this section is to explain the extensive tree decompositions mentioned in the proof sketch. Some ideas motivating this definition are already present in Andreae's proof that locally finite trees are ubiquitous under the topological minor relation [8, Lemma 2].

11.4.1 Separations and tree-decompositions of graphs

Definition 11.4.1. *Let T be a tree with a root $v \in V(T)$. Given nodes $x, y \in V(T)$ let us denote by xTy the unique path in T between x and y, by T_x denote the component of $T - E(vTx)$ containing x, and by $\overline{T_x}$ the tree $T - T_x$.*

Given an edge $e = tt' \in E(T)$, we say that t is the lower vertex *of e, denoted by e^-, if $t \in vTt'$. In this case, t' is the* higher vertex *of e, denoted by e^+.*

If S is a subtree of a tree T let us write $\partial(S) = E(S, T \setminus S)$ for the edge cut between S and its complement in T.

Definition 11.4.2. *Let $G = (V, E)$ be a graph. A* separation *of G is a pair (A, B) of subsets of vertices such that $A \cup B = V$ and such that there is no edge between $B \setminus A$ and $A \setminus B$. Given a separation (A, B) we write $\overline{G[B]}$ for the graph obtained by deleting all edges in the* separator *$A \cap B$ from $G[B]$.*

A reader unfamiliar with tree-decompositions may also consult [43, Section 12.3].

Definition 11.4.3 (Tree-decomposition)**.** *Given a graph $G = (V, E)$ a* tree-decomposition *of G is a pair $(T, \mathcal{V})$ consisting of a rooted tree T, together with a collection of subsets of vertices $\mathcal{V} = (V_t \subseteq V(G) \colon t \in V(T))$ such that:*

- *$V(G) = \bigcup \mathcal{V}$;*
- *For every edge $e \in E(G)$ there is a $t \in V(T)$ such that e lies in $G[V_t]$;*
- *$V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever $t_2 \in V(t_1 T t_3)$.*

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

Definition 11.4.4 (Tree-width)**.** *Suppose $(T, \mathcal{V})$ is a tree-decomposition of a graph G. The* width *of $(T, \mathcal{V})$ is the number $\sup\{|V_t| - 1 \colon t \in V(T)\} \in \mathbb{N} \cup \{\infty\}$. The* tree-width *of a graph G is the least width of any tree-decomposition of G.*

Definition 11.4.5 (Separations induced by tree-decompositions)**.** *Given a tree-decomposition $(T, \mathcal{V})$ of a graph G, and an edge $e \in E(T)$, let*

- *$A(e) := \bigcup \{V_{t'} \colon t' \notin V(T_{e^+})\}$,*
- *$B(e) := \bigcup \{V_{t'} \colon t' \in V(T_{e^+})\}$, and*
- *$S(e) := A(e) \cap B(e) = V_{e^-} \cap V_{e^+}$.*

Then $(A(e), B(e))$ is a separation of G (cf. [43, 12.3.1]). We call $B(e)$ the bough of $(T, \mathcal{V})$ rooted in *e and $S(e)$ the* separator of *$B(e)$. When writing $\overline{G[B(e)]}$ it is implicitly understood that this refers to the separation $(A(e), B(e))$ (cf. Definition 11.4.2.)*

Definition 11.4.6. *Let $(T, \mathcal{V})$ be a tree-decomposition of a graph G. For a subtree $S \subseteq T$ let us write*

$$G(S) = G\left[\bigcup_{t \in V(S)} V_t\right]$$

and if H is an IG we write $H(S) = H(G(S))$ for the restriction of H to $G(S)$.

Definition 11.4.7 (Self-similar bough)**.** *Let $(T, \mathcal{V})$ be a tree-decomposition of a graph G. Given $e \in E(T)$, the bough $B(e)$ is called* self-similar *(towards an end ω of G), if there is a set $\{R_{e,s} \colon s \in S(e)\}$ of disjoint ω-rays in G such that for all $n \in \mathbb{N}$ there is an edge $e' \in E(T_{e^+})$ with $\operatorname{dist}(e, e') \geqslant n$ such that*

- *for each $s \in S(e)$ the ray $R_{e,s}$ starts in s and meets $S(e')$;*
- *there is a subgraph $W \subseteq G[B(e')]$ which is an inflated copy of $\overline{G[B(e)]}$;*
- *for each $s \in S(e)$, we have $V(R_{e,s}) \cap S(e') \subseteq W(s)$.*

Such an W is called a witness for the self-similarity of $B(e)$ of distance at least n*.*

Definition 11.4.8 (Extensive tree-decomposition)**.** *A tree-decomposition $(T, \mathcal{V})$ of G is* extensive *if*

- *T is a locally finite, rooted tree;*
- *each part of $(T, \mathcal{V})$ is finite;*
- *every vertex of G appears in only finitely many parts of $\mathcal{V}$, and*
- *for each $e \in E(T)$, the bough $B(e)$ is self-similar towards some end ω_e of G.*

The following is the main result of this paper.

Theorem 11.4.9. *Every locally finite connected graph admitting an extensive tree-decomposition is $\preccurlyeq$-ubiquitous.*

11.4.2 Self minors and push-outs

The existence of an extensive tree-decomposition of a graph G will imply the existence of many self-minors of G, which will be essential to our proof.

Throughout this subsection, let G denote a locally finite, connected graph with an extensive tree-decomposition $(T, \mathcal{V})$.

Definition 11.4.10. *Let (A, B) be a separation of G with $A \cap B = \{v_1, v_2, \dots, v_n\}$. Suppose H_1, H_2 are subgraphs of a graph Γ where H_1 is an inflated copy of $G[A]$, H_2 is an inflated copy of $\overline{G[B]}$ and for all vertices $x, y \in G$, $H_1(x) \cap H_2(y) \neq \emptyset$ only if $x = y = v_i$ for some i. Suppose further that $\mathcal{P}$ is a family of disjoint paths $(P_i : i \in [n])$ in Γ such that each P_i is a path from $H_1(v_i)$ to $H_2(v_i)$ which is otherwise disjoint from $H_1 \cup H_2$. Note that P_i may be a single vertex if $H_1(v_i) \cap H_2(v_i) \neq \emptyset$.*

We write $H_1 \oplus_{\mathcal{P}} H_2$ for the IG given by (H, ϕ) where $H = \bigcup_{i \in [n]} P_i \cup H_1 \cup H_2$ and

$$H(v) = \phi^{-1}(v) := \begin{cases} H_1(v_i) \cup V(P_i) \cup H_2(v_i) & \text{if } v = v_i \in A \cap B, \\ H_1(v) & \text{if } v \in A \setminus B, \\ H_2(v) & \text{if } v \in B \setminus A. \end{cases}$$

Definition 11.4.11 (Push-out)**.** *A self minor $G' \subseteq G$ (meaning G' is an IG) is called a* push-out of G along e to depth n *for some $e \in E(T)$ if there is an edge $e' \in R_e$ such that* $\mathrm{dist}(e^-, e'^-) \geqslant n$ *and a subgraph $W \subseteq B(e')$ which is an $IG[\overline{B(e)}]$ such that $G' = G[A(e)] \oplus_{\mathcal{P}} W$, where $\mathcal{P} = (P_s : s \in S(e))$ is defined as the family of paths where P_s is the initial segment of $R_{e,s}$ up to the first point it meets $W(s)$.*

Similarly, if H is an IG then a subgraph H' of H is a push-out of H along e to depth n *for some $e \in E(T)$ if there is an edge $e' \in R_e$ such that* $\mathrm{dist}(e^-, e'^-) \geqslant n$ *and a subgraph $W \subseteq H(B(e'))$ which is an $IG[\overline{B(e)}]$ such that*

$$H' = H(G[A(e)]) \oplus_{\mathcal{P}} W$$

where $\mathcal{P} = (P_s : s \in S(e))$ is defined as the family of paths where P_s is the initial segment of $H^{\downarrow}(R_{e,s})$ up to the first point it meets $W(s)$.

Note that, if G' is a push-out of G along e to depth n then $H(G')$ has a subgraph which is a push-out of H along e to depth n.

Lemma 11.4.12. *For each $e \in E(T)$, each $n \in \mathbb{N}$ and each witness W of the self-similarity of $B(e)$ of distance at least n there is a corresponding push-out $G_W := G[A(e)] \oplus_{\mathcal{P}} W$ of G along e to depth n, where $\mathcal{P} = (P_s : s \in S(e))$ is defined as the family of paths where P_s is the initial segment of $R_{e,s}$ up to the first point it meets $W(s)$.*

Proof. Given an edge $e \in E(T)$, by Definition 11.4.7 for every $n \in \mathbb{N}$ there is a witness W for the self-similarity of $B(e)$ of distance at least n along the ray R_e.

Explicitly there is a family of rays $(R_{e,s} : s \in S(e))$ such that for every $n \in \mathbb{N}$ there is an edge $e' \in E(T_{e^+})$ of distance at least n from e, and a subgraph $W \subseteq G[B(e')]$, such that

- for each $s \in S(e)$ the ray $R_{e,s}$ starts in s and meets $S(e')$;
- W is an inflated copy of $\overline{G[B(e)]}$;
- for each $s \in S(e)$, we have $V(R_{e,s}) \cap S(e') \subseteq W(s)$.

Since $(A(e), B(e))$ and $(A(e'), B(e'))$, and $W \subseteq B(e')$ it is clear that $W \cap G[A(e)] \subseteq S(e)$, and since

Let us define $\mathcal{P} = (P_s : s \in S(e))$ as in the statement of the lemma. It is clear that each P_s is from $G[A(e)](s)$ to $W(s)$, and is otherwise disjoint from $G[A(e)] \cup W$.

Furthermore, since $(A(e), B(e))$ and $(A(e'), B(e'))$ are nested separations of G, $A(e) \cap V(W) \subseteq S(e) \cap S(e')$. Hence if $W(s) \cap G[A(e)](s') \neq \emptyset$ it follows that $s' \in S(e) \cap S(e')$, and hence $s' \in V(R_{e,s'}) \cap S(e') \subset W(s')$, by Definition 11.4.7. In particular, $W(s) \cap G[A(e)](s') \neq \emptyset$ only if $s = s' \in S(e)$.

Hence, by Definitions 11.4.10 and 11.4.11, $G[A(e)] \oplus_{\mathcal{P}} W$ is well-defined and is indeed a push-out of G along e to depth n. □

The existence of push-out of G along e to arbitrary depths is in some sense the essence of extensive tree-decompositions, and lies at the heart of our inductive construction in Section 11.9.

11.5 Existence of extensive tree-decompositions

The purpose of this section is to examine two classes of locally finite connected graphs that have extensive tree-decompositions: Firstly, the class of graphs with finitely many ends, all of which are thin, and secondly the class of graphs of finite tree-width. We will deduce the existence of such tree-decompositions using some results about the *well-quasi-ordering* of certain classes of graphs.

A *quasi-order* is a a reflexive and transitive binary relation, such as the minor relation between graphs. A quasi-order $\preccurlyeq$ on a set X is a *well-quasi-order* if for all sequences $x_1, x_2, \ldots \in X$ there exists an $i < j$ such that $x_i \preccurlyeq x_j$. The following two alternative characterisations will be useful.

Remark. *A simple Ramsey type argument shows that if $\preccurlyeq$ is a well-quasi-order on X, then every sequence $x_1, x_2, \ldots \in X$ contains an increasing subsequence $x_{i_1}, x_{i_2}, \ldots \in X$. That is, an increasing sequence $i_1 < i_2 < \ldots$ such that $x_{i_j} \preccurlyeq x_{i_k}$ for all $j < k$.*

Also, it is simple to show that if $\preccurlyeq$ is a well-quasi-order on X and $x_1, x_2, \ldots \in X$, then there is an $i_0 \in \mathbb{N}$ such that for every $i \geqslant i_0$ there are infinitely many $j \in \mathbb{N}$ with $x_i \preccurlyeq x_j$.

A famous result of Robertson and Seymour [115], proved over a series of 20 papers, shows that finite graphs are well-quasi-ordered under the minor relation. Thomas [123] showed that for any $k \in \mathbb{N}$ the class of graphs with tree-width $\leqslant k$ is well-quasi-ordered by the minor relation.

We will use slight strengthenings of both of these result, Lemma 11.5.2 and Lemma 11.5.9, to show that our two classes of graphs admit extensive tree-decompositions.

In Section 11.10 we will discuss in more detail the connection between our proof and well-quasi-ordering, and indicate how stronger well-quasi-ordering results could be used to prove the ubiquity of larger classes of graphs.

11.5.1 Finitely many thin ends

We will consider the following strengthening of the minor relation.

Definition 11.5.1. *Given $\ell \in \mathbb{N}$ an ℓ-pointed graph is a graph G together with a point function $\pi\colon [\ell] \to V(G)$. For ℓ-pointed graphs (G_1, π_1) and (G_2, π_2), we say $(G_1, \pi_1) \preccurlyeq_p (G_2, \pi_2)$ if $G_1 \preccurlyeq G_2$ and this can be arranged in such a way that $\pi_2(i)$ is contained in the branch set of $\pi_1(i)$ for every $i \in [\ell]$.*

Lemma 11.5.2. *The set of ℓ-pointed finite graphs is well-quasi-ordered under the relation $\preccurlyeq_p$.*

Proof. This follows from a stronger statement Robertson and Seymour proved in [112, 1.7]. □

We will also need the following structural characterisation of locally finite one-ended graphs with a thin end due to Halin.

Lemma 11.5.3. *Every one-ended, locally finite connected graph G with a thin end of degree $k \in \mathbb{N}$ has a tree-decomposition $(R, \mathcal{V})$ of G such that $R = t_0 t_1 t_2 \dots$ is a ray, and for every $i \in \mathbb{N}$:*

- *$|V_{t_i}|$ is finite;*
- *$|S(t_{i-1}t_i)| = k$;*
- *$S(t_{i-1}t_i) \cap S(t_i t_{i+1}) = \emptyset$.*

Proof. See [71, Satz 3′]. □

Note that in the above lemma, for a given finite set $X \subset V(G)$, by taking the union over an initial segment of parts, one may always assume that $X \subset V_{t_0}$. Moreover, note that since $S(t_{i-1}t_i) \cap S(t_i t_{i+1}) = \emptyset$, it follows that every vertex of G is contained in at most two parts of the tree-decomposition.

Lemma 11.5.4. *Every one-ended, locally finite connected graph G with a thin end has an extensive tree decomposition $(R, \mathcal{V})$ where $R = t_0 t_1 t_2 \dots$ is a ray with root t_0.*

Proof. Let $k \in \mathbb{N}$ be the degree of the thin end of G, and let $\mathcal{R} = \{R_j \colon j \in [k]\}$ be a maximal collection of disjoint rays in G. Let $(R', \mathcal{W})$ be the tree-decomposition of G given by Lemma 11.5.3 where $R' = t'_0 t'_1 \dots$ a ray.

Without loss of generality (taking the union over the first few parts, and considering tails of rays if necessary) we may assume that each ray in $\mathcal{R}$ starts in $S(t'_0 t'_1)$. Note that each ray in $\mathcal{R}$ meets the separator $S(t'_{i-1} t'_i)$ for each $i \in \mathbb{N}$. Since $\mathcal{R}$ is a disjoint family of k rays and $|S(t'_{i-1} t'_i)| = k$ for each $i \in \mathbb{N}$, each vertex in $S(t'_{i-1} t'_i)$ is contained in a unique ray in $\mathcal{R}$.

Let $\ell = 2k$ and consider a sequence $(G_i, \pi_i)_{i\in\mathbb{N}}$ of ℓ-pointed finite graphs defined by $G_i := G[W_{t'_i}]$ and

$$\pi_i \colon [\ell] \to V(G_i),\ j \mapsto \begin{cases} \text{the unique vertex in } S(t'_{i-1}t'_i) \cap V(R_j) & \text{for } 1 \leq j \leq k, \\ \text{the unique vertex in } S(t'_i t'_{i+1}) \cap V(R_{j-k}) & \text{for } k < j \leq 2k = \ell. \end{cases}$$

By Lemma 11.5.2 and Remark there is an n_0 such that for every $n \geqslant n_0$ there are infinitely many $m > n$ with $(G_n, \pi_n) \preccurlyeq_p (G_m, \pi_m)$.

Let $V_{t_0} := \bigcup_{i=0}^{n_0} W_{t'_i}$ and $V_{t_i} := W_{t'_{n_0+i}}$ for all $i \in \mathbb{N}$. We claim that $(R, (V_{t_i} : i \in \mathbb{N}))$ is the desired extensive tree-decomposition of G where $R = t_0t_1t_2\ldots$ is a ray with root t_0. The ray R is a locally finite tree and all the parts are finite. Moreover, every vertex of G is contained in at most two parts. It remains to show that for every $i \in \mathbb{N}$, the bough $B(t_{i-1}t_i)$ is self-similar.

Let $e = t_{i-1}t_i$. Let us label $\mathcal{R} = \{R_{e,s} \colon s \in S(e)\}$ where $R_{e,s}$ is the unique ray in $\mathcal{R}$ containing s. We wish to show there is a witness W for the self-similarity of $B(e)$ of distance at least n for each $n \in Nbb$. Note that $B(e) = \bigcup_{j\geqslant 0} G_{n_0+i+j}$. By the choice of n_0 in Remark , there exists $m > i + n$ such that $(G_{n_0+i}, \pi_{n_0+i}) \preccurlyeq_p (G_{n_0+m}, \pi_{n_0+m})$. Let $e' = t_{m-1}t_m$. We will show that there exists a $W \subseteq G[B(e')]$ witnessing the self-similarity of $B(e)$.

Recursively, for each $j \geqslant 0$ we can find $m = m_0 < m_1 < m_2 < \cdots$ with

$$(G_{n_0+i+j}, \pi_{n_0+i+j}) \preccurlyeq_p (G_{n_0+m_j}, \pi_{n_0+m_j}).$$

In particular there are subgraphs $H_{m_j} \subseteq G_{n_0+m_j}$ which are inflated copies of G_{n_0+i+j}, all compatible with the point-functions. In particular, $S(t'_{n_0+m_j-1}t'_{n_0+m_j}) \cup S(t'_{n_0+m_j}t'_{n_0+m_j+1}) \subset H_{m_j}$ for each $j \geqslant 0$.

Hence for, for every $j \in \mathbb{N}$ there is a unique $H_{m_{j-1}} - H_{m_j}$ subpath $P_{p,j}$ of R_p. We claim that

$$W' := \bigcup_{j\geqslant 0} H_{m_j} \cup \bigcup_{j\in\mathbb{N}} \bigcup_{p\in[k]} P_{p,j}$$

is a subgraph of $G[B(e')]$ that is an $IG[B(e)]$.

To prove this claim, for each $j \in \mathbb{N}$ and each $s \in S(t_{j-1}t_j)$, let $R_{p(s)} \in \mathcal{R}$ be the unique ray with $s \in R_{p(s)}$. Then $W'(s) = H_{m_{j-1}}(s) \cup P_{p(s),j} \cup H_{m_j}(s)$ is a connected branch set. Indeed, by construction, every $P_{p,j}$ is a path from $\pi_{n_0+m_{j-1}}(k+p)$ to $\pi_{n_0+m_j}(p)$. And since the H_{m_j} are pointed minors of $G_{n_0+m_j}$, it follows that $\pi_{n_0+m_{j-1}}(k+p(s)) \in H_{m_{j-1}}(s)$ and $\pi_{n_0+m_j}(p(s)) \in H_{m_j}(s)$ are as desired.

Finally, since $(G_{n_0+i}, \pi_{n_0+i}) \preccurlyeq_p (G_{n_0+m}, \pi_{n_0+m})$ as witnessed by H_{m_0}, the branch set of each $s \in S(t_{i-1}t_i)$ must indeed include $V(R_{e,s}) \cap S(e')$. □

Lemma 11.5.5. *If G is a locally finite connected graph with finitely many ends, each of which is thin, then G has an extensive tree-decomposition.*

Proof. Let $\Omega(G) = \{\omega_1, \ldots, \omega_n\}$ be the set of the ends of G. Pick a finite set $X \subseteq V$ of vertices separating the ends of G, i.e. so that all $C_i = C(X, \omega_i)$ are pairwise disjoint. Without loss of generality we may assume that $V(G) = X \cup \bigcup_{i\in[n]} C_i$.

Let $G_i := G[C_i \cup S]$. Then each G_i is a locally finite connected one-ended graph, with a thin end ω_i, and hence by Lemma 11.5.4 each of the G_i admits an extensive tree-decomposition $(R^i, \mathcal{V}^i)$ with root $r^i \in V(R^i)$. Without loss of generality, $X \subset V^i_{r^i}$ for each $i \in [n]$.

Let T be the tree formed by identifying the family of rays $(R^i \colon i \in [n])$ at their roots, let r be the root of T, and let $(T, \mathcal{V})$ be the tree-decompositions whose root part is $\bigcup_{i\in[n]} V^i_{r^i}$, and which otherwise agrees with the $(R^i, \mathcal{V}^i)$. It is a simple check that $(T, \mathcal{V})$ is an extensive tree-decomposition of G. □

11.5.2 Finite tree-width

Definition 11.5.6. *A rooted tree-decomposition $(T, \mathcal{V})$ of G is* lean *if for any $k \in \mathbb{N}$, any two nodes $t_1, t_2 \in V(T)$ and any $X_{t_1} \subseteq V_{t_1}, X_{t_2} \subseteq V_{t_2}$ such that $|X_{t_1}|, |X_{t_2}| \geq k$ there are either k disjoint paths in G, between X_1 and X_2, or there is a vertex t on the path in T between t_l and t_2 such that $|V_t| < k$.*

Remark. *Kříž and Thomas [89] showed that if G has tree-width $\leq m$ for some $m \in \mathbb{N}$, then G has a lean tree-decomposition of width $\leq m$.*

Lemma 11.5.7. *If G is a connected locally finite graph and $(T, (V_t : t \in T))$ a lean tree-decomposition of G such that every V_t is finite, then there is a locally finite subtree S of T such that $(S, (V_t : t \in S))$ is also a lean tree-decomposition of G.*

Proof. Pick a arbitrary root r of T. We will build recursively finite subtrees of T whose union will be the desired locally finite tree. Let $S_0 = L_0 = \{r\}$. For each $n \in \mathbb{N}$ let L_n be the set of leaves of S_n.

Consider some $t \in L_n$. Since V_t is finite and G is locally finite, the set $\mathcal{C}_t$ of components of $G - V_t$ is finite. Then, for each edge e leaving T_n with $t = e^-$ we have, by the definition of a tree-decomposition, that there is some subset $\mathcal{C}_e \subseteq \mathcal{C}_t$ such that

$$\bigcup \mathcal{C}_e \subseteq B(e) \subseteq \bigcup \mathcal{C}_e \cup V_t.$$

For each of the finitely may sets $\mathcal{C} \subseteq \mathcal{C}_t$ appearing as some $\mathcal{C}_e$ pick an arbitrary e which witnesses this. Let $E_t \subset E(T)$ be the set of all e chosen in this way, note that E_t is finite. Let S_{n+1} be $S_n \cup E_t$.

Finally, we let $S := \bigcup_{n \in \mathbb{N}} S_n$. It is simple to check that S is a locally finite tree and that $(S, \{V_t \mid t \in S\})$ is indeed a lean tree-decomposition of G. □

Lemma 11.5.8. *Let G be a locally finite, connected graph, and let $(T, \mathcal{V})$ be a lean tree-decomposition of G with root r and width $\leq m$, with T locally finite. Then there exists a lean tree-decomposition of G with width $\leq m$ such that every bough is connected, and the decomposition tree is locally finite. Moreover, we may assume that every vertex appears in only finitely many parts.*

Proof. Let $D_0 := \{r\}$ and $(T_0, \mathcal{V}_0) := (T, \mathcal{V})$. For $i \in \mathbb{N}$ let $D_i := \{e \in E(T_{i-1}) : \mathrm{dist}_{T_i}(r, e^-) = i\}$. Construct $(T_i, \mathcal{V}_i)$ from $(T_{i-1}, \mathcal{V}_{i-1})$ by performing the following operation for each edge $e \in D_i$:

Let $t = t_e^+$ and let $C_1, \ldots, C_n$ be the connected components of $B(e)$. Replace the subtree T_t with nT_t. For each $s \in T_t$ there are k copies of s in nT_t which we will call $s_1, \ldots, s_k$. For each $s \in T_t$ and $k \in [n]$ let $V_{s_k} := C_k \cap V_s$. Finally, let $\hat{T} = \bigcup_{i \in \mathbb{N}} T_i[\{t \in T_i \mid d_{T_i}(r, t) \leq i\}]$ and $\hat{\mathcal{V}} = (V_t \mid t \in \hat{T})$.

It is simple to check that $(\hat{T}, \hat{\mathcal{V}})$ is a tree-decomposition of width $\leq m$, that $\hat{T}$ is locally finite, and by construction $B(e)$ is connected for each $e \in E(T)$. Furthermore, suppose $k \in \mathbb{N}$, $t_1, t_2 \in \hat{T}$ and $X_{t_1} \subseteq \hat{V}_{t_1}, X_{t_2} \subseteq \hat{V}_{t_2}$ are such that $|X_{t_1}|, |X_{t_2}| \geq k$. By construction, there are nodes t_1' and t_2' of T such that $X_{t_1} \subseteq \hat{V}_{t_1} \subseteq V_{t_1'}, X_{t_2} \subseteq \hat{V}_{t_2} \subseteq V_{t_2'}$. Thus, since $(T, \mathcal{V})$ is lean, either there is a vertex t' of T between t_1', t_2' such that $|V_{t'}| < k$ or there are k disjoint paths between X_{t_1} and X_{t_2} in G. However, in the first case, by construction, there also is a node t of $\hat{T}$ between t_1 and t_2 such that $\hat{V}_t \subseteq V_{t'}$. Thus, $(\hat{T}, \hat{\mathcal{V}})$ is indeed lean.

Suppose there is an edge $e = st \in \hat{T}$, such that $B(e)$ if finite, but $\hat{T}_t$ is infinite. Since $\hat{V}_x \subseteq B(e)$ for any vertex $x \in V(\hat{T}_t)$, the set $\{\hat{V}_x : x \in V(\hat{T}_t)\}$ is finite. Hence, there is a finite subtree $\overline{T}_t \subseteq \hat{T}_t$ which contains at least one node for each of these bags. Let us replace, for each minimal $e \in E(T)$ with $B(e)$ finite, the subtree $\hat{T}_t$ with $\overline{T}_t$, to give a tree $\overline{T}$, and let $\overline{\mathcal{V}} = (\hat{V}_t \colon t \in V(\overline{T}))$. Then, $(\overline{T}, \overline{\mathcal{V}})$ is a lean-tree decomposition with width $\leq m$ such that $\overline{T}$ is locally finite and every bough $B(e)$ is connected. Moreover it has the following property

$$(\dagger) \quad \text{For every } t \in V(\overline{T}), \text{ if } \overline{T}_t \text{ is infinite, then so is } B(e).$$

Finally, suppose for a contradiction that there are vertices which appear in infinitely many parts of $(\overline{T}, \overline{\mathcal{V}})$. Let X be a $\subseteq$-maximal set of vertices appearing as a subset in infinitely many parts of $(\overline{T}, \overline{\mathcal{V}})$. Note that X is finite, since every part has size at most m. Since $\overline{T}$ is locally finite and $(\overline{T}, \overline{\mathcal{V}})$ is a tree-decomposition, there is a ray R in $\overline{T}$ such that X appears as a subset in every part corresponding to a node of R. We may assume without loss of generality that $R \subseteq \overline{T}_r$ where $r = \text{init}(R)$. Since for each $t \in R$ the subtree $\overline{T}_t$ contains a tail of R, it is infinite, and hence by $(\dagger)$ $B(e)$ is infinite and $X \subset B(e)$ for every $e \in R$,. Since $B(e)$ is connected, X has a neighbour in $B(e) \setminus X$. However, since G is locally finite, X has only finitely many neighbours, and by $\subseteq$-maximality of X each neighbour appears in only finitely many parts of $(\overline{T}, \overline{\mathcal{V}})$, and so in only finitely many sets $B(e)$ with $e \in R$. This contradicts the fact that X has a neighbour in every $B(e) \setminus X$. $\square$

Lemma 11.5.9. *For all $k, \ell \in \mathbb{N}$ the class of ℓ-pointed graphs with tree-width $\leq k$ is well-quasi-ordered under the relation $\preccurlyeq_p$.*

Proof. This is a consequence of a result of Thomas [123]. $\square$

Lemma 11.5.10. *Every locally finite connected graph of finite tree-width has an extensive tree-decomposition.*

Proof. Let G be a locally finite connected graph of tree-width $m \in \mathbb{N}$. By Lemma 11.5.7 there is a lean tree-decomposition $(T, \mathcal{V})$ of G with width m, such that T is a locally finite tree with root r. By Lemma 11.5.8 we may assume that every vertex is contained in only finitely many parts of this tree-decomposition.

Let ϵ be an end of T and let R be the unique ϵ-ray starting at the root of T. Let $d_\epsilon = \liminf_{e \in R} |S(e)|$, and fix a tail $t_0^\epsilon t_1^\epsilon \ldots$ of R such that $|S(t_{i-1}^\epsilon t_i^\epsilon)| \geqslant d_\epsilon$ for all $i \in \mathbb{N}$. Note that $|S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)| = d_\epsilon$ for an infinite sequence $i_1 < i_2 < \cdots$ of indices.

Since $(T, \mathcal{V})$ is lean, there are d_ϵ disjoint paths between $S(t_{i_k-1}^\omega t_{i_k}^\omega)$ and $S(t_{i_{k+1}-1}^\omega t_{i_{k+1}}^\omega)$ for every $k \in \mathbb{N}$. Moreover, since each $S(t_{i_k-1}^\omega t_{i_k}^\omega)$ is a separator of size d_ϵ, these paths are all internally disjoint. Hence, since every vertex appears in only finitely many parts, by concatenating these paths, we get a family of d_ϵ many disjoint rays in G.

Fix one such family of rays $(R_j^\epsilon \colon j \in [d_\epsilon])$. We claim that there is an end ω of G such that $R_j^\epsilon \in \omega$ for all $j \in [d_\epsilon]$. Indeed, if not then there is a finite set X separating some pair of rays R and R'. However, since each vertex appears in only finitely many parts, there is some $k \in \mathbb{N}$ such that $X \cap V_t = \emptyset$ for all $t \in T_{t_{i_k-1}^\epsilon}$. By construction R and R' have tails in $B(t_{i_{k+1}-1}^\omega t_{i_{k+1}}^\omega))$ which is connected, and disjoint from X, contradicting the fact that X separates R and R'.

For every $k \in \mathbb{N}$ we define a point-function $\pi_{i_k}^\epsilon \colon [d_\epsilon] \to S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)$ by letting $\pi_{i_k}^\epsilon(j)$ be the unique vertex in $R_j^\epsilon \cap S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)$.

By Lemma 11.5.9 and Remark , the sequence $(G[B(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)], \pi_{i_k}^\epsilon)_{k \in \mathbb{N}_{>0}}$ has an increasing subsequence $(G[B(t_{i-1}^\epsilon t_i^\epsilon)], \pi_i^\epsilon)_{i \in I_\epsilon}$, i.e. for any $k, j \in I_\epsilon, k < j$ we have

$$(G[B(t_{k-1}^\epsilon t_k^\epsilon)], \pi_k^\epsilon) \preccurlyeq_p (G[B(t_{j-1}^\epsilon t_j^\epsilon)], \pi_j^\epsilon).$$

Let us define $F_\epsilon = \{t_{k-1}^\epsilon t_k^\epsilon \colon k \in I_\epsilon\} \subset E(T)$.

Consider $T^- = T - \bigcup_{\epsilon \in \Omega(T)} F_\epsilon$, and let us write $\mathcal{C}(T^-)$ for the components of T^-. We claim that every component $C \in \mathcal{C}(T^-)$ is a locally finite rayless tree, and hence finite. Indeed, if

C contains a ray $R \subset T$ then R is in an end ϵ of T and hence $F_\epsilon \cap R \neq \emptyset$, a contradiction. Consequently, also each set $\bigcup_{t \in C} V_t$ is finite.

Let us define a tree decomposition $(T', \mathcal{V}')$ on $T' = T/\mathcal{C}(T^-)$ where $V'_{t'} = \bigcup_{t \in t'} V_t$. We claim this is an extensive tree-decomposition.

Clearly, T' is a locally finite tree, and each part of $(T', \mathcal{V}')$ is finite, and every vertex of G in contained in only finitely many parts of the tree-decomposition. Give $e \in E(T')$ there is some $\epsilon \in \Omega(T)$ such that $e \in F_\epsilon$. Consider the family of rays $(R_{e,j} \colon j \in [d_\epsilon])$ given by $R_{e,j} = R^\epsilon_j \cap B(e)$. Let ω_e be the end of G in which the rays $R_{e,j}$ lie.

There is some $k \in \mathbb{N}$ such that $e = t^\epsilon_{k-1} t^\epsilon_k$. Given $n \in \mathbb{N}$ let $k' \in I_\epsilon$ be such that there are at least n indices $\ell \in I_\epsilon$ with $k < \ell < k'$, and let $e' = t^\epsilon_{k'-1} t^\epsilon_{k'}$. Note that $e' \in F_\epsilon$ and hence $e' \in E(T')$. Furthermore, by construction e' has distance at least n from e in T'. Since $G[B(e)] = G[B(t^\epsilon_{k-1} t^\epsilon)]$ and $G[B(e')] = G[B(t^\epsilon_{k'-1} t^\epsilon_{k'})]$ we have $(G[B(e)], \pi^\epsilon_k) \preccurlyeq_p (G[B(e')], \pi^\epsilon_{k'})$, witnessing the self-similarity of $B(e)$ towards ω_e with the rays $(R_{e,j} \colon j \in [d_\epsilon])$. $\square$

Remark. *If for every $\ell \in \mathbb{N}$ the class of ℓ-pointed locally finite graphs without thick ends is well-quasi-ordered under $\preccurlyeq_p$, then every locally finite graph without thick ends has an extensive tree-decomposition. This follows by a simple adaptation of the proof above.*

11.5.3 Special graphs

We note that, whilst Lemmas 11.5.5 and 11.5.10 show that a large class of locally finite graphs have extensive tree-decompositions, for many other graphs it is possible to construct an extensive tree-decomposition 'by hand'. In particular, the fact that no graph in these classes has a thick end is an artefact of the method of proof, rather than a necessary condition for the existence of such a tree-decomposition, as is demonstrated by the following examples:

Remark. *The grid $\mathbb{Z} \times \mathbb{Z}$ has an extensive tree-decomposition, as can be seen in Figure 11.2. More explicitly, we can take a ray decomposition of the grid given by a sequence of increasing diamond shaped regions around the origin. It is easy to check that every bough will self similar.*

A similar argument shows that the half-grid has an extensive tree-decomposition. However, we note that both of these graphs were already be shown to be ubiquitous in [25].

In fact, we do not know of any construction of a locally finite graph which does not admit an extensive tree-decomposition.

Question 11.5.11. *Do all locally finite graphs admit an extensive tree-decomposition?*

11.6 The structure of non-pebbly ends

We will need a structural understanding of how the arbitrarily large families of IGs (for some fixed graph G) can be arranged inside of some host graph Γ. In particular we are interested in how the rays of these minors occupy a given end ϵ of Γ. In [25] we established the distinction between *pebbly* and *non-pebbly* ends, cf. Definition 11.6.4. We showed that the existence of a pebbly end of Γ already guarantees the existence of a $K^{\aleph_0}$-minor in Γ, and therefore the following corollary holds:

Corollary 11.6.1 ([25, Corollary 6.4])**.** *Let Γ be a graph with a pebbly end ω and let G be a countable graph. Then $\aleph_0 G \preccurlyeq \Gamma$.*

We will now analyse the structure of non-pebbly ends and give a description of their shape. For a fixed set of start vertices we will consider the possible families of disjoint rays with these start vertices. This shall be made precise in the definition of *polypods*, cf. Definition 11.6.7 below.

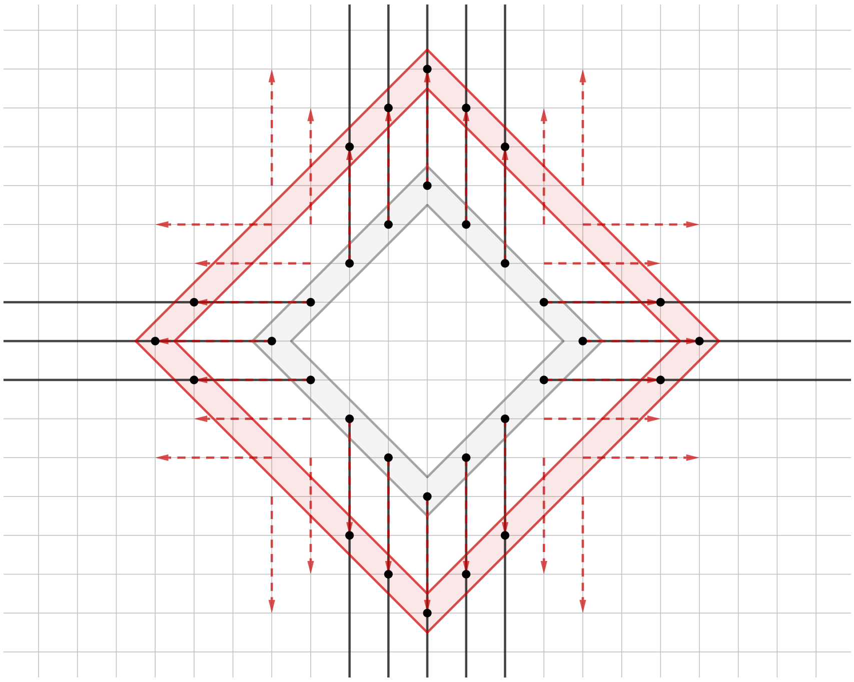

Figure 11.2: In the grid the boughs are self-similar.

We will investigate how these rays relate in terms of connecting paths between them and see that, due to the non-pebbly structure of the end, the structure of possible connections between the rays is somewhat restricted.

11.6.1 Pebble Pushing

Given a path P with end-vertices s and t we say *the orientation of P from s to t* to mean the total order on the vertices of P where $a \leq b$ if and only if a lies on sPb, in this case we say that a lies before b. Note that every path with at least one edge has precisely two orientations.

Given a cycle C a *cyclic orientation of C* is an orientation of the graph C which does not have any sink. Note that any cycle has precisely two cyclic orientations. Given a cyclic orientation and 3 distinct vertices x, y, z we say that they *appear consecutively in the order* (x, y, z) if y lies on the unique directed path from x to z. Given two cycles C, C', each with a cyclic orientation, we say that an injection $f\colon V(C) \to V(C')$ *preserves the cyclic orientation* if whenever three distinct vertices x, y and z appear on C in the order (x, y, z) then their images appear on C' in the order $(f(x), f(y), f(z))$.

A permutation of a finite set X is a bijection $\nu\colon X \to X$. A sequence $(x_1 \dots x_n)$ of distinct elements of X is called a *cycle* of ν if $\nu(x_n) = x_1$ and $\nu(x_i) = x_{i+1}$ for all $i \in \{1, \dots, n-1\}$. In this case n is called the *length* of the cycle, a cycle of length 1 is called *trivial*. The term $(x_1 \dots x_n)$ is also used to denote the bijection ν which contains the cycle $(x_1 \dots x_n)$ and otherwise is the identity on $X \setminus \{x_1, \dots, x_n\}$. It is a well-known fact that every bijection can be written as a product of (disjoint) cycles.

We utilise the following results and definitions from [25].

Definition 11.6.2 (Pebble-pushing game)**.** *Let $G = (V, E)$ be a graph. We call a tuple $(x_1, \dots, x_k) \in V^k$ a* game state *(of order k) if $x_i \neq x_j$ for all $i, j \in [k]$ with $i \neq j$.*

The pebble-pushing game *(on G) is a game played by a single player. Given a game state $Y = (y_1, \dots, y_k)$, we imagine k labelled pebbles placed on the vertices $(y_1, \dots, y_k)$. A move for a* *game state in the pebble-pushing game consists of moving a pebble from a vertex to an adjacent vertex which does not contain a pebble, or formally, a Y-*move *is a game state $Z = (z_1, \dots, z_k)$ such that there is an $\ell \in [k]$ such that $y_\ell z_\ell \in E$ and $y_i = z_i$ for all $i \in [k] \setminus \{\ell\}$.*

*Let $X = (x_1, \dots, x_k)$ be a game state. The X-*pebble-pushing game (on G) *is a pebble-pushing game where we start with k labelled pebbles placed on the vertices $(x_1, \dots, x_k)$.*

We say a game state Y is achievable *in the X-pebble-pushing game if there is a sequence $(X_i \colon i \in [n])$ of game states for some $n \in \mathbb{N}$ such that $X_1 = X$, $X_n = Y$ and X_{i+1} is a X_i-move for all $i \in [n-1]$, that is if there is a sequence of moves that pushes the pebbles from X to Y.*

A graph G is k-pebble-win *if Y is an achievable game state in the X-pebble-pushing game on G for every two game states X and Y of order k.*

Lemma 11.6.3 ([25, Lemma 4.2])**.** *Let Γ be a graph, $\omega \in \Omega(\Gamma)$, $m \geqslant k$ be positive integers and let $(S_j \colon j \in [m])$ be a family of disjoint rays in ω. For every achievable game state $Z = (z_1, z_2, \dots, z_k)$ in the $(1, 2, \dots, k)$-pebble-pushing game on $RG(S_j \colon j \in [m])$, the map σ defined via $\sigma(i) := z_i$ for every $i \in [k]$ is a transition function[4] from $(S_i \colon i \in [k])$ to $(S_j \colon j \in [m])$.*

Definition 11.6.4 (Pebbly ends)**.** *Let Γ be a graph and ω an end of Γ. We say ω is* pebbly *if for every k there is an $n \geqslant k$ and a family $\mathcal{R} = \{R_1, \dots, R_n\}$ of disjoint rays in ω such that $RG(R_i \colon i \in [n])$ is k-pebble-win. If for some k there is no such family $\mathcal{R}$ we say ω is* not k-pebble-win.

Lemma 11.6.5 ([25, Lemma 6.3])**.** *Let Γ be a graph and let $\omega \in \Omega(\Gamma)$ be a pebbly end. Then $K_{\aleph_0} \preccurlyeq \Gamma$.*

Recall that a path $P = v_0 v_1 \dots v_n$ in a graph G is called a *bare* if all its inner vertices have degree 2 in G.

Corollary 11.6.6 ([25, Corollary 5.2])**.** *Let ω be an end of Γ which is not k-pebble-win and let $\mathcal{R} = (R_i \colon i \in [m])$ be a family of $m \geqslant k + 2$ disjoint rays in ω. Then there is a bare path $P = p_1 \dots p_n$ in $RG(R_i \colon i \in [m])$ such that $n \geqslant m - k$. Furthermore, either each edge in P is a bridge, or $RG(R_i \colon i \in [m])$ is a cycle.*

11.6.2 Polypods

Definition 11.6.7. *Given an end ω of a graph Γ, a* polypod (for ω in Γ) *is a pair (X, Y) of disjoint finite sets of vertices of Γ such that there is at least one family $(R_y \colon y \in Y)$ of disjoint rays to ω, where R_y begins at y and all the R_y are disjoint from X. Such a family (R_y) is called a* family of tendrils *for (X, Y). The* order *of the polypod is $|Y|$. The* connection graph $K_{X,Y}$ *of a polypod (X, Y) is a graph with vertex set Y. It has an edge between vertices v and w if and only if there is a family $(R_y \colon y \in Y)$ of tendrils for (X, Y) such that there is an R_v–R_w-path in Γ disjoint from X and from every other R_y.*

Note that the ray graph of any family of tendrils for a polypod must be a subgraph of the connection graph of that polypod.

Definition 11.6.8. *We say that a polypod (X, Y) for ω in Γ is* tight *if its connection graph is minimal amongst connection graphs of polypods for ω in Γ with respect to the spanning isomorphic subgraph relation, i.e. for no other polypod (X', Y') for ω in Γ of order $|Y'| = |Y|$ is the graph $K_{X',Y'}$ isomorphic to a proper subgraph of $K_{X,Y}$. (Let us write $H \subsetneq G$ if H is isomorphic to a subgraph of G.) We say that a polypod* attains *its connection graph if there is some family of tendrils for that polypod whose ray graph is equal to the connection graph.*

[4] See Definition 11.2.3.

Lemma 11.6.9. *Let (X, Y) be a tight polypod, $(R_y \colon y \in Y)$ a family of tendrils and for every $y \in Y$ let v_y be a vertex on R_y. Let X' be a finite vertex set disjoint from all $v_y R_y$ and including X as well as each of the initial segments $R_y \mathring{v}_y$. Let $Y' = \{v_y \colon y \in Y\}$. Then (X', Y') is a tight polypod with the same connection graph as (X, Y).*

Proof. The family $(v_y R_y \colon y \in Y)$ witnesses that (X', Y') is a polypod. Moreover every family of tendrils for (X', Y') can be extended by the paths $R_y v_y$ to obtain a family of tendrils for (X, Y). Hence if there is an edge $v_y v_z$ in $K_{X'Y'}$ then there must also be the edge yz in $K_{X,Y}$. Thus $K_{X',Y'} \subsetneq K_{X,Y}$. But since (X, Y) is tight we must have equality. Therefore (X', Y') is tight as well. □

Lemma 11.6.10. *Any tight polypod (X, Y) attains its connection graph.*

Proof. We must construct a family of tendrils for (X, Y) whose ray graph is $K_{X,Y}$. We will recursively build larger and larger initial segments of the rays, together with disjoint paths between them.

Precisely this means that, after partitioning $\mathbb{N}$ into infinite sets A_e, one for each edge e of $K_{X,Y}$, we will construct, for each $n \in \mathbb{N}$, a family $(P^n_y \colon y \in Y)$ of paths and a path Q_n such that:

- Each P^n_y starts at y.
- Each P^n_y has length at least n.
- For $m \leqslant n$, the path P^n_y extends P^m_y.
- If $n \in A_{vw}$ then Q_n is a path from P^n_v to P^n_w.
- If $n \in A_{vw}$ then Q_n meets no P^m_y with $y \notin \{v, w\}$.
- All the Q_n are disjoint.
- All the P^n_y and all the Q_n are disjoint from X.
- For any n there is a family $(R^n_y \colon y \in Y)$ of tendrils for (X, Y) such that each P^n_y is an initial segment of the corresponding R^n_y, and the R^n_y only meet the Q_m with $m \leqslant n$ in inner vertices of the P^n_y.

It is clear that if we can do this then we will obtain a family of tendrils by letting R_y be the union of all the P^n_y. Furthermore, for any edge e of $K_{X,Y}$ the family $(Q_n \colon n \in A_e)$ will witness that e is in the ray graph of this family. So that ray graph will be the whole of $K_{X,Y}$, as required.

So it remains to explain how to carry out this recursive construction. Let vw be the edge of $K_{X,Y}$ with $1 \in A_{vw}$. By the definition of the connection graph there is a family $(R^1_y \colon y \in Y)$ of tendrils for (X, Y) such that there is a path Q_1 from R^1_v to R^1_w, disjoint from all other R^1_y and from X. For each $y \in Y$ let P^1_y be an initial segment of R^1_y of length at least 1 and containing all vertices of $Q^1 \cap R^1_y$. This choice of the P^1_y and of Q_1 clearly satisfies the conditions above.

Now suppose that we have constructed suitable P^m_y and Q_m for all $m \leqslant n$. For each $y \in Y$, let y_n be the endvertex of P^n_y. Let Y_n be $\{y_n \colon y \in Y\}$ and

$$Z_n = X \cup \bigcup_{m \leqslant n} \bigcup_{y \in Y} \left(V(P^m_y) \cup V(Q_m) \right).$$

Let X_n be $Z_n \setminus Y_n$, and note that every $V(Q_m) \subset X_n$ for every m. Then by Lemma 11.6.9 (X_n, Y_n) is a tight polypod with the same connection graph as (X, Y).

In particular, letting vw be the edge of $K_{X,Y}$ with $n+1 \in A_{vw}$, we have that $v_n w_n$ is an edge of K_{X_n,Y_n}. So there is a family $(S^{n+1}_{y_n} : y_n \in Y_n)$ of tendrils for (X_n, Y_n) together with a path Q_{n+1} from $S^{n+1}_{v_n}$ to $S^{n+1}_{w_n}$ disjoint from all other $S^{n+1}_{y_n}$ and from X_n. Now for any $y \in Y$ we let R^{n+1}_y be the ray $yP^n_y y_n S^{n+1}_{y_n}$ and let P^{n+1}_y be an initial segement of R^{n+1}_y long enough to include P^n_y, of length at least $n+1$, and containing all vertices of $Q_{n+1} \cap R^{n+1}_y$ as inner vertices. This completes the recursion step, and so the construction is complete. □

Lemma 11.6.11. *If (X, Y) is a polypod of order n for ω in Γ with connection graph $K_{X,Y}$ then for any set of n disjoint ω-rays $(R_i : i \in [n])$ in Γ, $RG(R_i : i \in [n]) \subsetneq K_{X,Y}$.*

Proof. If we apply the Weak Linking Lemma 11.2.2 to the rays $(R_i : i \in [n])$ and a family of tendrils for (X, Y), together with the finite set X, we obtain a family of tendrils for (X, Y) whose tails coincide with that of $(R_i : i \in [n])$. Hence, the ray graph of these tendrils is $RG(R_i : i \in [n])$ and so $RG(R_i : i \in [n]) \subsetneq K_{X,Y}$. □

Corollary and Definition 11.6.12. *Any two polypods for ω in Γ of the same order which attain their connection graphs have isomorphic connection graphs.*

We will refer to the graph arising in this way for polypods of order n for ω in Γ as the n^{th} shape graph *of the end ω.* □

11.6.3 Frames

Akin to the transition boxes defined in Lemma 11.3.17 we want to consider *frames*, finite subgraphs which are just large enough to include a linkage which, say, induces a transition function of the family of tendrils of some polypod. This will allow us to reason about transition functions in terms of graph automorphisms.

Definition 11.6.13. *Let Y be a finite set. A Y-*frame *(L, α, β) consists of a finite graph L together with two injections α and β from Y to $V(L)$. The set $A = \alpha(Y)$ is called the* source set *and the set $B = \beta(Y)$ is called the* target set*. A* weave *of the Y-frame is a family $\mathcal{Q} = (Q_y : y \in Y)$ of disjoint paths in L from A to B, where the initial vertex of Q_y is $\alpha(y)$ for each $y \in Y$. The* weave pattern *$\pi_{\mathcal{Q}}$ of $\mathcal{Q}$ is the bijection from Y to itself sending y to the inverse image under β of the endvertex of Q_y. In order words, $\pi_{\mathcal{Q}}$ is the function so that every Q_y is an $\alpha(y) - \beta(\pi_{\mathcal{Q}}(y))$ path. The* weave graph *$K_{\mathcal{Q}}$ of $\mathcal{Q}$ has vertex set Y and an edge joining distinct vertices u and v of Y precisely when there is a path from Q_u to Q_v in L disjoint from all other Q_y. We call the Y-frame* strait *if it has at most one weave graph and at most one weave pattern. For a graph K with vertex set Y, we say that the Y-frame is K-*spartan *if all its weave graphs are subgraphs of K and all its weave patterns are automorphisms of K.*

Connection graphs of polypods and weave graphs of frames are closely connected:

Lemma 11.6.14. *Let (X, Y) be a polypod for ω in Γ attaining its connection graph $K_{X,Y}$ and let $\mathcal{R} = (R_y : y \in Y)$ be a family of tendrils for (X, Y). Let L be any finite subgraph of Γ disjoint from X but meeting all the R_y. For each $y \in Y$ let $\alpha(y)$ be the first vertex of R_y in L and $\beta(y)$ the last vertex of R_y in L. Then the Y-frame (L, α, β) is $K_{X,Y}$-spartan.*

Proof. Since there is some family of tendrils $(S_y : y \in Y)$ attaining $K_{X,Y}$ and there is by Lemma 11.2.2 a linkage from $(R_y : y \in Y)$ to $(S_y : y \in Y)$ after X and $V(L)$, we may assume without loss of generality that $RG(R_y : y \in Y)$ is isomorphic to $K_{X,Y}$.

For a given weave $\mathcal{Q} = (Q_y : y \in Y)$, applying the definition of the connection graph to the rays $R_y \alpha(y) Q_y \beta(\pi_{\mathcal{Q}}(y)) R_{\pi_{\mathcal{Q}}(y)}$ shows that $K_{\mathcal{Q}}$ is a subgraph of $K_{X,Y}$ and that the inverse image of any edge of $K_{X,Y}$ under π_Q is again an edge of $K_{X,Y}$, from which it follows that π_Q is an automorphism of $K_{X,Y}$. $\square$

Corollary 11.6.15. *Let (X,Y) be a polypod for ω in Γ attaining its connection graph $K_{X,Y}$ and let $\mathcal{R} = (R_y \colon y \in Y)$ be a family of tendrils for (X,Y). Then for any transition function σ from $\mathcal{R}$ to itself there is a $K_{X,Y}$-spartan Y-frame for which both σ and the identity are weave patterns.*

Proof. Pick a linkage $(P_y \colon y \in Y)$ from $\mathcal{R}$ to itself after X inducing σ. Let L be a finite subgraph graph of Γ containing all P_y as well as a finite segment of each R_y, such that each P_y is a path between two such segments. Then the frame on L which exists by Lemma 11.6.14 has the desired properties. $\square$

Lemma 11.6.16. *Let (X,Y) be a polypod for ω in Γ attaining its connection graph $K_{X,Y}$ and let $\mathcal{R} = (R_y \colon y \in Y)$ be a family of tendrils for (X,Y). Then there is a $K_{X,Y}$-spartan Y-frame for which both $K_{X,Y}$ and $RG(R_y \colon y \in Y)$ are weave graphs.*

Proof. By adding finitely many vertices and edges to X if necessary, we may obtain a superset X' of X such that for any two of the R_y if there is any path between them disjoint from all the other rays and X', then there are infinitely many such paths. Let $(S_y \colon y \in Y)$ be any family of tendrils for (X,Y) with connection graph $K_{X,Y}$.

For each edge $e = uv$ of $RG(R_y \colon y \in Y)$ let P_e be a path from R_u to R_v disjoint from all the other R_y and from X'. Similarly for each edge $f = uv$ of $K_{X,Y}$ let Q_f be a path from S_u to S_v disjoint from all the other S_y and from X'. Let $(P'_y \colon y \in Y)$ be a linkage from the S_y to the R_y after

$$X' \cup \bigcup_{e \in E(RG(R_y \colon y \in Y))} P_e \cup \bigcup_{f \in E(K_{X,Y})} Q_f.$$

Let the initial vertex of P'_y be $\gamma(y)$ and the end vertex be $\beta(y)$. Let $\pi(y)$ be the element of Y with $\beta(y)$ on $R_{\pi(y)}$. Let L be the subgraph of Γ containing all paths of the forms $S_y\gamma(y), R_{\pi(Y)}\beta(y), P'_y, P_e$ and Q_f.

Letting α be the identity function on Y, it follows from Lemma 11.6.14 that (L, α, β) is a $K_{X,Y}$-spartan Y-frame. The paths Q_f witness that the weave graph for the paths $S_y\gamma(y)P'_y$ includes $K_{X,Y}$ and so, by $K_{X,Y}$-spartanness, must be equal to $K_{X,Y}$. The paths P_e witness that the weave graph for the paths $R_y\beta(y)$ includes the ray graph of the R_y. The two must be equal since whenever for two of the R_y there is any path between them, disjoint from all the other R_y and from X', then there are infinitely many disjoint such paths. $\square$

Hence to understand ray graphs and the transition functions between them it is useful to understand the possible weave graphs and weave patterns of spartan frames. Their structure can be captured in terms of automorphisms and cycles:

Definition 11.6.17. *Let K be a finite graph. An automorphism σ of K is called* local *if it is a cycle $(z_1 \dots z_t)$ where, for any $i \leqslant t$, there is an edge from z_i to $\sigma(z_i)$ in K. If $t \geqslant 3$ this means that $z_1 \dots z_t z_1$ is a cycle of K, and we call such cycles* turnable*. If $t = 2$ then we call the edge $z_1 z_2$ of K* flippable*. We say that an automorphism of K is* locally generated *if it is a product of local automorphisms.*

Remark. *A cycle C in K is turnable if and only if all its vertices have the same neighbourhood in $K - C$, and whenever a chord of length $\ell \in \mathbb{N}$ is present in $K[C]$, then all chords of length ℓ are present. Similarly an edge e of K is flippable if and only if its two endvertices have the same neighbourhood in $K - e$. Thus, if K contains at least 3 vertices, no vertex of degree one* *or cutvertex of K can lie on a turnable cycle or a flippable edge. So vertices of degree one and cutvertices are preserved by locally generated automorphisms.*

Lemma 11.6.18. *Let (L, α, β) be a Y-frame which is K-spartan but not strait. Then each of its weave graphs includes a turnable cycle or a flippable edge of K and for any two of its weave patterns π and π' the automorphism $\pi^{-1} \cdot \pi'$ of K is locally generated.*

Proof. Suppose not for a contradiction, and let (L, α, β) be a counterexample in which $|E(L)|$ is minimal. Note that, as L is not strait, there are either at least two weave patterns for L or there are at least two weave graphs for L. Thus, we can find weaves $\mathcal{P} = (P_y \colon y \in Y)$ and $\mathcal{Q} = (Q_y \colon y \in Y)$ such that either $K_\mathcal{P} \neq K_\mathcal{Q}$ or $\pi_\mathcal{P} \neq \pi_\mathcal{Q}$ and such that either $K_\mathcal{Q}$ includes no turnable cycle or flippable edge or $\pi_\mathcal{P}^{-1} \cdot \pi_\mathcal{Q}$ is not locally generated. Furthermore, by exchanging $\mathcal{P}$ and $\mathcal{Q}$ if necessary, we may assume that $K_\mathcal{P}$ is not a proper subgraph of $K_\mathcal{Q}$.

Each edge of L is in one of $\mathcal{P}$ or $\mathcal{Q}$ since otherwise we could simply delete it. Similarly no edge appears in both $\mathcal{P}$ and $\mathcal{Q}$ since otherwise we could simply contract it. No vertex appears on just one of P_y or Q_y since otherwise we could contract one of the two incident edges. Vertices appearing in neither $\mathcal{P}$ nor $\mathcal{Q}$ are isolated and so may be ignored. Thus we may suppose that each edge appears in precisely one of $\mathcal{P}$ or $\mathcal{Q}$, and that each vertex appears in both.

Let Z be the set of those $y \in Y$ such that $\alpha(y) \neq \beta(y)$. For any $z \in Z$ let $\gamma(z)$ be the second vertex of P_z and let $f(z) \in Y$ be chosen such that $\gamma(z)$ lies on $Q_{f(z)}$. Then since $\gamma(z) \neq \alpha(f(z))$ we have $f(z) \in Z$ for all $z \in Z$. Furthermore, Z is nonempty as $\mathcal{P}$ and $\mathcal{Q}$ are distinct. Let z be any element of Z. Then since Z is finite there must be $i < j$ with $f^i(z) = f^j(z)$, which means that $f^i(z) = f^{j-i}(f^i(z))$. Let $t > 0$ be minimal such that there is some $z_1 \in Z$ with $z_1 = f^t(z_1)$.

If $t = 1$ then we may delete the edge $\alpha(z_1)\gamma(z_1)$ and replace the path P_{z_1} with $\alpha(z_1)Q_{z_1}\gamma(z_1)P_{z_1}$. This preserves all of $\pi_\mathcal{P}$, $\pi_\mathcal{Q}$ and $K_\mathcal{Q}$ and can only make $K_\mathcal{P}$ bigger, contradicting the minimality of our counterexample. So we must have $t \geqslant 2$.

For each $i \leqslant t$ let z_i be $f^{i-1}(z_1)$ and let σ be the bijection $(z_1 z_2 \ldots z_t)$ on Y. Let L' be the graph obtained from L by deleting all vertices of the form $\alpha(z_i)$. Let α' be the injection from Y to $V(L')$ sending z_i to $\gamma(z_i)$ for $i \leqslant n$ and sending any other $y \in Y$ to $\alpha(y)$. Then (L', α', β) is a Y-frame. For any weave $(\hat{P}_y \colon y \in Y)$ in this Y-frame, $(\alpha(y)\gamma(y)\hat{P}_y)_{y \in Y}$ is a weave in (L, α, β) with the same weave pattern and whose weave graph includes that of $(\hat{P}_y \colon y \in Y)$. Thus (L', α', β) is K-spartan.

Let P'_y be $\alpha'(y)P_y$ and Q'_{y_i} be $\alpha'(y_i)Q_{\sigma(y_i)}$ for any $y \in Y$. Then we have $\pi_{\mathcal{Q}'} = \pi_\mathcal{Q} \cdot \sigma$ and so $\sigma = \pi_\mathcal{Q}^{-1} \cdot \pi_{\mathcal{Q}'}$ is an automorphism of K. For any $i \leqslant t$ the edge $\alpha(z_i)\gamma(z_i)$ witnesses that $z_i\sigma(z_i)$ is an edge of $K_\mathcal{Q}$ and so σ is local. Hence $K_\mathcal{Q}$ includes a turnable cycle or a flippable edge. By the minimality of $|E(L)|$ we know that $\pi_{\mathcal{P}'}^{-1} \cdot \pi_{\mathcal{Q}'}$ is locally generated and hence so is $\pi_\mathcal{P}^{-1} \cdot \pi_\mathcal{Q} = \pi_{\mathcal{P}'}^{-1} \cdot \pi_{\mathcal{Q}'} \cdot \sigma^{-1}$. This is the desired contradiction. □

Finally, the following two lemmas are the main outcomes of this section:

Lemma 11.6.19. *Let (X, Y) be a polypod attaining its connection graph $K_{X,Y}$ such that $K_{X,Y}$ is a cycle of length at least 4. Then for any family of tendrils $\mathcal{R}$ for this polypod the ray graph is $K_{X,Y}$. Furthermore, any transition function from $\mathcal{R}$ to itself preserves each of the cyclic orientations of $K_{X,Y}$.*

Proof. By Lemma 11.6.16 there is some $K_{X,Y}$-spartan Y-frame for which both $K_{X,Y}$ and the ray graph are weave graphs. Since $K_{X,Y}$ is a cycle of length at least 4 and hence has no flippable edges, the ray graph must include a cycle by Lemma 11.6.18 and so since it is a subgraph of $K_{X,Y}$ it must be the whole of $K_{X,Y}$. Similarly Lemma 11.6.18 together with Corollary 11.6.15 shows that all transition functions must be locally generated and so must preserve the orientation. □

Lemma 11.6.20. *Let (X, Y) be a polypod attaining its connection graph $K_{X,Y}$ such that $K_{X,Y}$ includes a bare path P whose edges are bridges. Let $\mathcal{R}$ be a family of tendrils for (X, Y) whose ray graph is $K_{X,Y}$. Then for any transition function σ from $\mathcal{R}$ to itself, the restriction of σ to P is the identity.*

Proof. By Lemmas 11.6.15 and 11.6.18 any transition function must be a locally generated automorphism of $K_{X,Y}$, and so by Remark it cannot move the vertices of the bare path, which are vertices of degree one or cutvertices. □

11.7 Grid-like and half-grid-like ends

We are now in a position to analyse the different kinds of thick ends which can arise in a graph in terms of the possible ray graphs and the transition functions between them. We fix a graph Γ together with a thick end ω of Γ. If ω is pebbly then $K_{\aleph_0} \preccurlyeq \Gamma$ by Lemma 11.6.5, and every locally finite graph G satisfies $\aleph_0 G \preccurlyeq K_{\aleph_0} \preccurlyeq \Gamma$.

So in the following we further restrict ourselves to the case that ω is not pebbly; for this section we fix a number N such that there is no family $(R_i \colon i \in [n])$ of disjoint rays with $n \geqslant N$ such that $RG(R_i \colon i \in [n])$ is N-pebble win. Under these circumstances we get nontrivial restrictions on the ray graphs and the transition functions between them. There are two essentially different cases, corresponding to the two cases in Corollary 11.6.6: The grid-like and the half-grid-like case.

11.7.1 Grid-like ends

The first case is ends which behave like that of the infinite grid. In this case, all large enough ray graphs are cycles and all transition functions between them preserve the cyclic order.

Formally, we say that the end ω is *grid-like* if the $(N+2)^{\text{nd}}$ shape graph for ω is a cycle. For the rest of this subsection we will assume that ω is grid-like. Let us fix some polypod (X, Y) of order $N + 2$ attaining its connection graph. Let $(S_y \colon y \in Y)$ be a family of tendrils for (X, Y) whose ray graph is the cycle $C_{N+2} = K_{X,Y}$.

Lemma 11.7.1. *Any ray graph K for a set $(R_i \colon i \in I)$ of ω-rays in Γ with $|I| \geqslant N + 2$ is a cycle.*

Proof. Let $(T_y \colon y \in Y)$ be a family of tendrils for (X, Y) obtained by transitioning from the S_y to the R_i after X along a linkage, and let $\sigma \colon Y \to I$ be the function induced by this linkage. Then by Lemma 11.6.19 the ray graph of the T_y is the cycle $K_{X,Y}$. We know by Corollary 11.6.6 that K includes a bare path P such that $|V(P)| \geqslant |V(K)| - N$. Thus there are distinct vertices $y_1, y_2 \in Y$ with $\sigma(y_1), \sigma(y_2) \in P$ and no other vertex in the image of σ between them on P. Then for any other vertex y of Y there are paths from y to y_1 avoiding y_2 and from y to y_2 avoiding y_1 in $K_{X,Y}$. Hence there are paths from $\sigma(y)$ to each of $\sigma(y_1)$ and $\sigma(y_2)$ avoiding $\sigma(y_1)P\sigma(y_2)$. Thus none of the edges of $\sigma(y_1)P\sigma(y_2)$ is a bridge, so by Corollary 11.6.6 again K is a cycle. □

We will now choose cyclic orientations of all these cycles such that the transition functions preserve the cyclic orders corresponding to those orientations. To that end, we fix a cyclic orientation of $K_{X,Y}$. We say that a cyclic orientation of the ray graph for a family $(R_i \colon i \in I)$ of at least $N + 3$ disjoint ω-rays is *correct* if there is a transition function σ from the S_y to the R_i which preserves the cyclic orientation of $K_{X,Y}$.

Lemma 11.7.2. *For any such family $(R_i \colon i \in I)$ of at least $N + 3$ disjoint ω-rays there is precisely one correct cyclic orientation of its ray graph.*

Proof. That there is at least one is clear by Lemma 11.2.2. Suppose for a contradiction that there are two, and let σ and σ' be transition functions witnessing that both orientations of the ray graph are correct. By Lemma 11.6.3 we may assume without loss of generality that the images of σ and σ' are the same. Call this common image I'. Since the ray graphs of $(R_i \colon i \in I)$ and $(R_i \colon i \in I')$ are both cycles, the former is obtained from the latter by subdivision of edges. Since this doesn't affect the cyclic order, we may assume without loss of generality that $I' = I$. By Lemma 11.2.2 again, there is some transition function τ from the R_i to the S_y. By Lemma 11.6.19 both $\tau \cdot \sigma$ and $\tau \cdot \sigma'$ must preserve the cyclic order, which is the desired contradiction. □

It therefore makes sense to refer to *the* correct orientation of a ray graph.

Corollary 11.7.3. *Any transition function between ray graphs on at least $N+3$ rays preserves the correct orientations of the cycles.* □

11.7.2 Half-grid-like ends

In this subsection we suppose that ω is thick but neither pebbly nor grid-like. We shall call such ends *half-grid-like*, since as we shall shortly see in this case the ray graphs and the transition functions between them behave similarly to those for the unique end of the half grid.

We will need to carefully consider how the ray graphs are divided up by their cutvertices. In particular, for a graph K and vertices x and y of K we will denote by $C^{xy}(K)$ the union of all components of $K - x$ which do not contain y, and we will denote by K^{xy} the graph $K - C^{xy}(K) - C^{yx}(K)$. We will refer to K^{xy} as the part of K *between* x and y.

As in the last subsection, let (X, Y) be a polypod of order $N+2$ attaining its connection graph and let $(S_y \colon y \in Y)$ be a family of tendrils for (X, Y) with ray graph $K_{X,Y}$, which by assumption is not a cycle. By Corollary 11.6.6 there is a bare path of length at least 2 in $K_{X,Y}$ of which all edges are bridges. Let y_1y_2 be any edge of that path. Without loss of generality we have $C^{y_1y_2}(K_{X,Y}) \neq \emptyset$.

Let $(R_i \colon i \in I)$ be a family of disjoint rays with $|I| \geqslant N+3$ and let K be its ray graph.

Remark. *For any transition function σ from the S_y to the R_i we have $\sigma[C^{y_1y_2}(K_{X,Y})] \subseteq C^{\sigma(y_1)\sigma(y_2)}(K)$ and $\sigma[C^{y_2y_1}(K_{X,Y})] \subseteq C^{\sigma(y_2)\sigma(y_1)}(K)$. Thus $\sigma[K_{X,Y}]$ and $K^{\sigma(y_1)\sigma(y_2)}$ meet precisely in $\sigma(y_1)$ and $\sigma(y_2)$.*

Lemma 11.7.4. *For any transition function σ from the S_y to the R_i the graph $K^{\sigma(y_1)\sigma(y_2)}$ is a path from $\sigma(y_1)$ to $\sigma(y_2)$. This path is a bare path in K and all of its edges are bridges.*

Proof. Since K is connected, $K^{\sigma(y_1)\sigma(y_2)}$ must include a path P from $\sigma(y_1)$ to $\sigma(y_2)$. If it is not equal to that path then it follows from Lemma 11.6.3 that the function σ', which we define to be just like σ except for $\sigma'(y_1) = \sigma(y_2)$ and $\sigma'(y_2) = \sigma(y_1)$, is a transition function from the S_y to the R_i. But then by Remark we have $\sigma[C^{y_1y_2}(K_{X,Y})] \subseteq C^{\sigma(y_1)\sigma(y_2)}(K) \cap C^{\sigma'(y_1)\sigma'(y_2)}(K) = C^{\sigma(y_1)\sigma(y_2)}(K) \cap C^{\sigma(y_2)\sigma(y_1)}(K) = \emptyset$. So this is impossible, and $K^{\sigma(y_1)\sigma(y_2)} = P$. The last sentence of the lemma now follows from the definition of $K^{\sigma(y_1)\sigma(y_2)}$. □

Now we fix a transition function $\sigma_{\max}$ so that the path $P := K^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}$ is as long as possible. If $\sigma_{\max}[C^{y_1y_2}(K_{X,Y})]$ were a proper subset of $C^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}(K)$ then we would be able to use Lemma 11.6.3 to produce a transition function in which this path is longer. So we must have $\sigma_{\max}[C^{y_1y_2}(K_{X,Y})] = C^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}(K)$ and similarly $\sigma_{\max}[C^{y_2y_1}(K_{X,Y})] = C^{\sigma_{\max}(y_2)\sigma_{\max}(y_1)}(K)$.

We call P the *central* path of K and the orientation from $\sigma_{\max}(y_1)$ to $\sigma_{\max}(y_2)$ the *correct* orientation. In the following lemma we use this orientation to determine which vertices appear before which along P.

Lemma 11.7.5. *For any two vertices v_1 and v_2 of K, there is a transition function $\sigma\colon K_{X,Y} \to K$ with $\sigma(y_1) = v_1$ and $\sigma(y_2) = v_2$ if and only if v_1 and v_2 both lie on P, with v_1 before v_2.*

Proof. The 'if' direction is clear by applying Lemma 11.6.3 to $\sigma_{\max}$. For the 'only if' direction, we begin by setting $c_1 = |C^{y_1 y_2}(K_{X,Y})|$ and $c_2 = |C^{y_2 y_1}(K_{X,Y})|$. We enumerate $C^{y_1 y_2}(K_{X,Y})$ as $y_3 \dots y_{c_1+2}$ and $C^{y_2 y_1}(K_{X,Y})$ as $y_{c_1+3} \dots y_{c_1+c_2+2}$. Then for any $N+2$-tuple $(x_1 \dots x_{N+2})$ of distinct vertices achievable in the $(\sigma_{\max}(y_1), \dots, \sigma_{\max}(y_{N+2}))$ pebble pushing game must have the following 3 properties, since they are preserved by any single move:

- x_1 and x_2 lie on P, with x_1 before x_2.
- $\{x_3, \dots, x_{c_1+2}\} \subseteq C^{x_1 x_2}(K)$.
- $\{x_{c_1+3}, \dots, x_{c_1+c_2+2}\} \subseteq C^{x_2 x_1}(K)$.

Now let σ be any transition function from the S_y to the R_i. Let $(x_1, \dots, x_{N+2})$ be an $N+2$-tuple achievable in the $(\sigma_{\max}(y_1), \dots, \sigma_{\max}(y_{N+2}))$ pebble pushing game such that $\{x_1, \dots, x_{N+1}\} = \sigma[Y]$. By Lemma 11.6.3 the function σ' sending y_i to x_i for each $i \leqslant N+2$ is also a transition function and $\sigma'[Y] = \sigma[Y]$. Let τ be a transition function from $(R_i\colon i \in \sigma[Y])$ to the S_y. Then by Lemma 11.6.20 both $\tau \cdot \sigma$ and $\tau \cdot \sigma'$ keep both y_1 and y_2 fixed. Thus $\sigma(y_1) = \sigma'(y_1) = x_1$ and $\sigma(y_2) = \sigma'(y_2) = x_2$. As noted above, this means that $\sigma(y_1)$ and $\sigma(y_2)$ both lie on P with $\sigma(y_1)$ before $\sigma(y_2)$, as desired. □

Thus the central path and the correct orientation depend only on our choice of y_1 and y_2. Hence we get

Corollary 11.7.6. *Each ray graph contains a unique central path with a correct orientation and all transition functions between ray graphs send vertices of the central path to vertices of the central path and preserve the correct orientation.*

We note that, in principle, this trichotomy that an end of a graph is either pebbly, grid-like or half-grid-like, and the information that this implies about its finite rays graphs, could be derived from earlier work of Diestel and Thomas [53], who gave a structural characterisation of graphs without a $K_{\aleph_0}$-minor. However, to introduce their result and derive what we needed from it would have been at least as hard, if not more complicated, and so we have opted for a straightforward and self-contained presentation.

11.7.3 Core rays in the half-grid-like case

Definition 11.7.7. *Given a graph G, an end ω and three rays R, S, T in ω such that R, S, T have disjoint tails, we say that S* separates *R* from *T if the tails of R and T disjoint from S belong to different ends of $G - S$.*

For the following, recall the definition of *ray graph* in Definition 11.2.4.

Lemma 11.7.8. *Let G be a graph, ω an end of G and $(R_i)_{i \in I}$ a finite family of disjoint ω-rays. If, for some $i_1, i_2, j \in I$, the vertices i_1 and i_2 belong to different components of $RG((R_i)_{i \in I}) - j$, then R_j separates R_{i_1} from R_{i_2}.*

Proof. If R_{i_1} and R_{i_2} belong to the same end of $G - V(R_j)$, there are infinitely many disjoint paths between R_{i_1} and R_{i_2} in $G - V(R_j)$. Hence, by the pigeonhole principle there are indices j_1 and j_2 belonging to different components of $RG((R_i)_{i \in I}) - j$, such that these disjoint paths induce infinitely many disjoint paths from R_{j_1} to R_{j_2} all disjoint from all other R_i. Thus there is an edge from j_1 to j_2 in $RG((R_i)_{i \in I})$ contradicting the assumption that i disconnects j_1 from j_2. □

Lemma 11.7.9. *Consider three rays R, S, T belonging to the same end ω of some graph G. If S separates R from T, then T does not separate R from S and R does not separate S from T.*

Proof. As R and T both belong to ω, there are infinitely many disjoint paths between them. As S separates R from T, S must meet infinitely many of these paths. Hence, there are infinitely many disjoint paths from S to R, all disjoint from T. Similarly, there are infinitely many disjoint paths from S to T, all disjoint from R. Hence T does not separate R from S and R does not separate S from T. $\square$

Definition 11.7.10. *Given a graph G and two (possibly infinite) vertex-sets X and Y, we say that an end ω of $G - X$ is a* sub-end *of an end ω' of $G - Y$ if every ray in ω has a tail in ω'.*

Definition 11.7.11. *Let ω be a half-grid-like end, let R be an ω-ray. We say R is a* core ray *(*of ω*) if there is a finite family $\mathcal{R} = (R_i : i \in I)$ of disjoint ω-rays with $R = R_c$ for some $c \in I$ such that c lies on, but is not an endpoint, of the central path of $\mathcal{R}$.*

Lemma 11.7.12. *Let R be a core ray of ω. Then in $G - R$ the end ω splits into precisely two different ends. (That is, there are two ends ω' and ω'' of $G - R$ such that every ω-ray in $G \setminus V(R)$ is in ω' or ω''.)*

Proof. Let $\mathcal{R} = (R_i : i \in I)$ be a family witnessing that $R = R_c$ for some $c \in I$ is a core ray. Then there are exactly two ends in $G \setminus V(R)$ which contain rays in $\mathcal{R}$, since connected components of $RG(\mathcal{R})$ when we delete the vertex corresponding to R are equivalent sets of rays in $G \setminus V(R)$ and more over, no two of these connected components can belong to the same end of $G \setminus V(R)$ by Lemma 11.7.8.

Suppose there is a third end in $G \setminus V(R)$ that contains an ω-ray S. We first claim that there is a tail of S which is disjoint from $\mathcal{R}$. Indeed, clearly S is disjoint from R, and if S met $\bigcup \mathcal{R}$ infinitely often then it would meet some $R_i \in \mathcal{R}$ infinitely often, and hence lie in the same end of $G \setminus V(R)$ as R_i. Let us assume then that S is disjoint from $\mathcal{R}$.

Let us consider the ray graph $RG(\mathcal{R} \cup \{S\})$. Again, if S is adjacent to any ray except R in the ray graph, it would lie in the same end as some ray in $\mathcal{R}_J$ in $G \setminus V(R)$.

Since S is an ω-ray the ray graph is connected, and hence S is adjacent to R, and R is still connected to its neighbours in $RG(\mathcal{R})$. However, $\mathcal{R} \cup \{S\}$ is also a family that witnesses that $R = R_c$ is a core ray and hence c has degree two in $RG(\mathcal{R} \cup \{S\})$, a contradiction. $\square$

Given a family of rays $(R_i)_{i \in I}$ witnessing that $R = R_c$ is a core ray, we denote by $\top(R, (R_i)_{i \in I})$ the end of $G - V(R)$ containing rays R_i satisfying $i < c$ and with $\perp(R, (R_i)_{i \in I})$ the end containing rays R_i satisfying $i > c$.

Lemma 11.7.13. *Let R and S be disjoint core rays of ω. Let us suppose that ω splits in $G - S$ in ω'_S and ω''_S and in $G - R$ in ω'_R and ω''_R. If R belongs to ω'_S and S belongs to ω'_R, then ω''_S is a sub-end of ω'_R and ω''_R is a sub-end of ω'_S.*

Proof. Let T be a ray in ω''_S. As R belongs to a different end of $G - S$ than T, there is a tail of T which is disjoint from R. Thus, we may assume that T and R are disjoint. As S separates R from T, by Lemma 11.7.9, R does not separate S from T, hence T belongs to ω'_R. $\square$

Lemma and Definition 11.7.14. *Let $\mathcal{R}_1 = (R_i : i \in I_1)$, $\mathcal{R}_2 = (R_i : i \in I_2)$ be two finite families of disjoint ω-rays both witnessing that for some $c \in I_1 \cap I_2$ the ray R_c is a core ray in ω. Then $\top(R, (R_i)_{i \in I_1}) = \top(R, (R_i)_{i \in I_2})$ and $\perp(R, (R_i)_{i \in I_1}) = \perp(R, (R_i)_{i \in I_2})$.*

We therefore write $\top(\omega, R)$ for the end $\top(R, (R_i)_{i \in I_1})$ and $\perp(\omega, R)$ respectively, i.e $\top(\omega, R)$ is the end of $G - R$ containing rays that appear on the central path of some ray graph before R according to the correct orientation and $\perp(\omega, R)$ is the end of $G - R$ containing rays that appear *on the central path of some ray graph after R according to the correct orientation. Note that $\top(\omega, R) \cap \bot(\omega, R) = \emptyset$.*

Proof. Suppose, this is not the case, hence $\omega_1 := \top(R_c, (R_i)_{i\in I_1}) = \bot(R_c, (R_i)_{i\in I_2})$ and $\omega_2 := \bot(R_c, (R_i)_{i\in I_1}) = \top(R_c, (R_i)_{i\in I_2})$. Let $\mathcal{R}_{I_1\omega_1}$ be the set of rays in $\mathcal{R}_1$ belonging to ω_1. Let $\mathcal{R}_{I_1\omega_2}$,$\mathcal{R}_{I_2\omega_1}$ and $\mathcal{R}_{I_2\omega_2}$ be defined accordingly. If $|\mathcal{R}_{I_1\omega_1}| > |\mathcal{R}_{I_2\omega_1}|$ we define $\mathcal{R}_{\omega_1}$ to be $\mathcal{R}_{I_1\omega_1}$, otherwise $\mathcal{R}_{\omega_1} = \mathcal{R}_{I_2\omega_1}$. Let $\mathcal{R}_{\omega_2}$ be defined similarly.

Let us consider $\mathcal{R} := \mathcal{R}_{\omega_1} \cup \mathcal{R}_{\omega_2} \cup \{R_c\}$. After replacing some of the rays with tails, this is a collection of disjoint rays, so let us assume that $\mathcal{R}$ itself is a family of disjoint rays. There is a transition function from $\mathcal{R}_{I_1}$ to $\mathcal{R}$ mapping R_c to itself, every ray in $\mathcal{R}_{I_1\omega_1}$ to a ray in $\mathcal{R}_{\omega_1}$ and every ray in $\mathcal{R}_{I_1\omega_2}$ to a ray in $\mathcal{R}_{\omega_2}$:

Consider a finite separator X separating ω_1 from ω_2 in $G - V(R_c)$. Consider linkages after X in $G - V(R_c)$ from $\mathcal{R}_{\omega_1}$ to $\mathcal{R}_{\omega_1}$ and from $\mathcal{R}_{\omega_2}$ to $\mathcal{R}_{\omega_2}$. Pairs of such linkages can be combined to suitable linkages on G, inducing a transition function which is as desired.

Similarly there is a transition function from $\mathcal{R}_{I_2}$ to $\mathcal{R}$ mapping R_c to itself, every ray in $\mathcal{R}_{I_2\omega_1}$ to a ray in $\mathcal{R}_{\omega_1}$ and every ray in $\mathcal{R}_{I_2\omega_2}$ to a ray in $\mathcal{R}_{\omega_2}$.

These transition functions preserve the central path, thus c lies on the central path of $RG(\mathcal{R})$. Moreover, $\mathcal{R}$ also witness that R_c is a core ray. However, the first transition function shows that $\omega_1 = \top(R_c, \mathcal{R})$ whereas the second one shows that $\omega_2 = \top(R_c, \mathcal{R})$, contradicting the assumption that $\omega_1 \neq \omega_2$. □

Lemma and Definition 11.7.15. *Let* $\operatorname{core}(\omega)$ *denote the set of core rays in ω. We define a partial order $\leqslant_\omega$ on* $\operatorname{core}(\omega)$ *by*

$$R \leqslant_\omega S \textit{ if and only if either } R = S, \quad \textit{or } R \textit{ and } S \textit{ have disjoint tails } xR \textit{ and } yS \textit{ and } xR \in \top(\omega, yS)$$

for $R, S \in$ $\operatorname{core}(\omega)$.

Proof. For the anti-symmetry let us suppose that R and S are disjoint rays such that $R \leqslant_\omega S$ and $S \leqslant_\omega R$. Hence, $R \in \top(\omega, S)$ as well as $S \in \top(\omega, R)$. Let $\mathcal{R}_S$ be a family of rays witnessing that S is a core ray and $\mathcal{R}_R$ a family witnessing that R is a core ray. By Lemma 11.7.13, $\bot(\omega, S)$ is a sub-end of $\top(\omega, R)$ and $\bot(\omega, R)$ is a sub-end of $\top(\omega, S)$. Let $\mathcal{R}_{\bot(S)}$ be the subset of $\mathcal{R}_S$ of rays, which belong to $\bot(\omega, S)$. Let $\mathcal{R}_{\bot(R)}$ be defined accordingly. After replacing rays with tails all rays in $\mathcal{R} := \mathcal{R}_{\bot(S)} \cup \mathcal{R}_{\bot(R)} \cup \{R\} \cup \{S\}$ are pairwise disjoint. More over, R and S both lie on the central path of $RG(\mathcal{R})$ and are both not endpoints of this central path. Thus either $S \in \bot(\omega, R)$ or $R \in \bot(\omega, S)$ contradicting Lemma 11.7.14.

For the transitivity, let us suppose that R, S, T are rays, such that $R \leqslant_\omega S$ and $S \leqslant_\omega T$. We may assume that R and S, and S and T are disjoint. As $\leqslant_\omega$ is anti-symmetric, it is $T \nleqslant_\omega S$, hence $T \in \bot(\omega, S)$. Thus, R and T belong to different ends of $G - S$, thus we may assume that they are also disjoint. As S therefore separates R from T, by Lemma 11.7.9, T does not separate S from R. Thus, R and S belong to the same end of $G - T$. Hence $R \in \top(\omega, T)$. □

Remark. *Let $R, S \in$* $\operatorname{core}(\omega)$ *and let $\mathcal{R}$ be a finite family of disjoint ω-rays.*

1. *Any ray which shares a tail with R is also a core ray of ω.*
2. *If R and S are disjoint, then R and S are comparable under $\leqslant_\omega$.*
3. *If R and S are on the central path of $\mathcal{R}$, then $R \leqslant_\omega S$ if and only if R appears before S in the correct orientation of $RG(\mathcal{R})$.*

4. The maximum number of disjoint rays in $\omega \setminus \mathrm{core}(\omega)$ is bounded by $2 \cdot (p_\omega + 1)$.

Lemma 11.7.16. *Let $R, S \in \mathrm{core}(\omega)$. Let $Z \subseteq V(G)$ be a finite set such that $\top(\omega, S)$ and $\bot(\omega, S)$ are separated by Z in $G - V(S)$. Let $H \subseteq G - Z$ be a connected subgraph which is disjoint to S and contains R, and let $T \subseteq H$ be some core ω-ray. Then S is in the same relative $\leqslant_\omega$-order to T as to R.*

Proof. Assume $S \leqslant_\omega R$ and hence $R \in \top(\omega, S)$. Since H is connected, we obtain that $T \in \top(\omega, S)$ as well and hence $S \leqslant_\omega T$. The other case is analogous. □

Lemma and Definition 11.7.17. *Let $\mathcal{R}$ be a finite family of disjoint core ω-rays. Then there exists a family $\mathcal{R}'$ of disjoint ω-rays such that $RG(\mathcal{R})$ is precisely the inner vertices of the central path of $RG(\overline{\mathcal{R}})$. Even though such a family is not unique, we denote by $\overline{\mathcal{R}}$ an arbitrary such family.*

Definition 11.7.18. *If $\mathcal{P}$ is a linkage from $\mathcal{R}$ to $\mathcal{S}$ then a* sub-linkage *of $\mathcal{P}$ is just a subset of $\mathcal{P}$, considered as a linkage from the corresponding subset of $\mathcal{R}$ to $\mathcal{S}$.*

Remark. *A sub-linkage of a transitional linkage is transitional.*

Proof. By Remark 2 the rays in $\mathcal{R}$ are linearly ordered by $\leqslant_\omega$. Let R denote the $\leqslant_\omega$-smallest and S denote the $\leqslant_\omega$-greatest element of $\mathcal{R}$. As in the proof of Lemma 11.7.15, consider the sets $\mathcal{R}_{\bot(R)}$ and $\mathcal{R}_{\top(S)}$, which are without loss of generality minimal with respect to their defining property. Now $\mathcal{R}_{\bot(R)} \subseteq \bot(\omega, R)$ and $R' \in \top(\omega, R)$ for every $R' \in \mathcal{R} \setminus \{R\}$ and hence tails of $\mathcal{R}_{\bot(R)}$ are disjoint to $\bigcup \mathcal{R}$. Analogously, $\mathcal{R}_{\top(S)} \subseteq \top(\omega, S)$ and $R' \in \bot(\omega, S)$ for every $R' \in \mathcal{R} \setminus \{S\}$ and hence tails of $\mathcal{R}_{\bot(R)}$ are disjoint to $\bigcup \mathcal{R}$. Finally, $\mathcal{R}_{\top(S)} \subseteq \top(\omega, R)$ and $\mathcal{R}_{\bot(R)} \subseteq \bot(\omega, S)$ by Lemma 11.7.13, yielding that tails of $\mathcal{R}_{\top(S)}$ are necessarily disjoint from tails in $\mathcal{R}_{\bot(R)}$. Their the union of those tails with $\mathcal{R}$ yields a set $\overline{\mathcal{R}}$ as desired. □

Definition 11.7.19. *Let $\mathcal{R}$, $\mathcal{S}$ be finite families of disjoint ω-rays and let $\mathcal{R}'$ be a subfamily of $\mathcal{R}$ consisting of core rays. A linkage $\mathcal{P}$ between $\mathcal{R}$ and $\mathcal{S}$ is* preserving on $\mathcal{R}'$ *if $\mathcal{P}$ links $\mathcal{R}'$ to core rays and preserves the order $\leqslant_\epsilon$.*

The following remarks are a direct consequence of the definitions and Corollary 11.7.6.

Remark. *Let $\mathcal{R}$, $\mathcal{S}$, $\mathcal{T}$ be finite families of disjoint ω-rays, let $\mathcal{R}' \subseteq \mathcal{R}$ be a subfamily of core rays, and let $\mathcal{P}_1$, $\mathcal{P}_2$ be a linkages from $\mathcal{R}$ to $\mathcal{S}$ and from $(\mathcal{R} \circ_{\mathcal{P}_1} \mathcal{S})$ to $\mathcal{T}$ respectively.*

1. *If $\mathcal{P}_1$ is transitional and $\mathcal{R}'$ is on the central path of $\mathcal{R}$, then it is preserving on $\mathcal{R}'$.*
2. *If $\mathcal{P}_1$ is preserving on $\mathcal{R}'$, then the sub-linkage of $\mathcal{P}_1$ from $\mathcal{R}'$ to the respective subfamily of $\mathcal{S}$ is transitional.*
3. *If $\mathcal{P}_1$ is preserving on $\mathcal{R}'$, then any $\mathcal{P}_1' \subseteq \mathcal{P}_1$ as a linkage between the respective subfamilies is preserving on the respective subfamily of $\mathcal{R}'$.*
4. *If $\mathcal{P}_1$ is preserving on $\mathcal{R}'$ and $\mathcal{P}_2$ is preserving on $\mathcal{R}' \circ_{\mathcal{P}_1} \mathcal{S}$, then the concatenation $\mathcal{P}_1 + \mathcal{P}_2$ is preserving on $\mathcal{R}'$.*

Lemma 11.7.20. *Let $\mathcal{R}$ and $\mathcal{S}$ be finite families of disjoint core rays of ω, and let $\mathcal{S}' \subseteq \mathcal{S}$ be a subfamily of $\mathcal{S}$ with $|\mathcal{R}| = |\mathcal{S}'|$. Then there is a transitional linkage from $\overline{\mathcal{R}}$ to $\overline{\mathcal{S}}$ which is preserving on $\mathcal{R}$ and links the rays in $\mathcal{R}$ to rays in $\mathcal{S}'$.*

Proof. Consider $\mathcal{T} := (\overline{\mathcal{S}} \setminus \mathcal{S}) \cup \mathcal{S}' \subseteq \overline{\mathcal{S}}$. Take a transitional linkage from $\overline{\mathcal{R}}$ to $\mathcal{T}$. This linkage can be viewed as a linkage from $\overline{\mathcal{R}}$ to $\overline{\mathcal{S}}$, is preserving on $\mathcal{R}$ by Remark 1, and hence the sub-linkage from $\mathcal{R}$ to $\mathcal{S}'$ is also preserving on $\mathcal{R}$ by Remark 3 as well as transitional by Remark . □

11.8 G-tribes and concentration of G-tribes towards an end

To show that a given graph G is $\preccurlyeq$-ubiquitous, we shall assume that $nG \preccurlyeq \Gamma$ for every $n \in \mathbb{N}$ and need to show that this implies $\aleph_0 G \preccurlyeq \Gamma$. To this end we use the following notation for such collections of nG in Γ which is established in [24] and [25] .

Definition 11.8.1 (G-tribes)**.** *Let G and Γ be graphs.*

- *A G-*tribe in Γ (with respect to the minor relation) *is a family $\mathcal{F}$ of finite collections F of disjoint subgraphs H of Γ such that each* member *H of $\mathcal{F}$ is an IG.*
- *A G-tribe $\mathcal{F}$ in Γ is called* thick*, if for each $n \in \mathbb{N}$ there is a* layer *$F \in \mathcal{F}$ with $|F| \geqslant n$; otherwise, it is called* thin*.*
- *A G-tribe $\mathcal{F}$ is* connected *if every member H of $\mathcal{F}$ is connected. Note that this is the case precisely if G is connected.*
- *A G-tribe $\mathcal{F}'$ in Γ is a G-*subtribe [5] *of a G-tribe $\mathcal{F}$ in Γ, denoted by $\mathcal{F}' \preccurlyeq \mathcal{F}$, if there is an injection $\Psi\colon \mathcal{F}' \to \mathcal{F}$ such that for each $F' \in \mathcal{F}'$ there is an injection $\varphi_{F'}\colon F' \to \Psi(F')$ with $V(H') \subseteq V(\varphi_{F'}(H'))$ for every $H' \in F'$. The G-subtribe $\mathcal{F}'$ is called* flat*, denoted by $\mathcal{F}' \subseteq \mathcal{F}$, if there is such an injection Ψ satisfying $F' \subseteq \Psi(F')$.*
- *A thick G-tribe $\mathcal{F}$ in Γ is* concentrated at an end ϵ *of Γ, if for every finite vertex set X of Γ, the G-tribe $\mathcal{F}_X = \{F_X \colon F \in \mathcal{F}\}$ consisting of the layers $F_X = \{H \in F \colon H \not\subseteq C(X, \epsilon)\} \subseteq F$ is a thin subtribe of $\mathcal{F}$.*

We note that, if G is connected, every thick G-tribe $\mathcal{F}$ contains a thick subtribe $\mathcal{F}'$ such that every $H \in \bigcup \mathcal{F}$ is a tidy IG. We will use the following lemmas from [24].

Lemma 11.8.2 (Removing a thin subtribe, [24, 5.2])**.** *Let $\mathcal{F}$ be a thick G-tribe in Γ and let $\mathcal{F}'$ be a thin subtribe of $\mathcal{F}$, witnessed by $\Psi\colon \mathcal{F}' \to \mathcal{F}$ and $(\varphi_{F'}\colon F' \in \mathcal{F}')$. For $F \in \mathcal{F}$, if $F \in \Psi(\mathcal{F}')$, let $\Psi^{-1}(F) = \{F'_F\}$ and set $\hat{F} = \varphi_{F'_F}(F'_F)$. If $F \notin \Psi(\mathcal{F}')$, set $\hat{F} = \emptyset$. Then*

$$\mathcal{F}'' := \{F \setminus \hat{F} \colon F \in \mathcal{F}\}$$

is a thick flat G-subtribe of $\mathcal{F}$.

Lemma 11.8.3 (Pigeon hole principle for thick G-tribes, [24, 5.3])**.** *Suppose for some $k \in \mathbb{N}$, we have a k-colouring $c\colon \bigcup \mathcal{F} \to [k]$ of the members of some thick G-tribe $\mathcal{F}$ in Γ. Then there is a monochromatic, thick, flat G-subtribe $\mathcal{F}'$ of $\mathcal{F}$.*

Lemma 11.8.4 ([24, 5.4])**.** *Let G be a connected graph and Γ a graph containing a thick connected G-tribe $\mathcal{F}$. Then either $\aleph_0 G \preccurlyeq \Gamma$, or there is a thick flat subtribe $\mathcal{F}'$ of $\mathcal{F}$ and an end ϵ of Γ such that $\mathcal{F}'$ is concentrated at ϵ.*

Lemma 11.8.5 ([24, 5.5])**.** *Let G be a connected graph and Γ a graph containing a thick connected G-tribe $\mathcal{F}$ concentrated at an end ϵ of Γ. Then the following assertions hold:*

1. *For every finite set X, the component $C(X, \epsilon)$ contains a thick flat G-subtribe of $\mathcal{F}$.*
2. *Every thick subtribe $\mathcal{F}'$ of $\mathcal{F}$ is concentrated at ϵ, too.*

The following lemma from [25] shows that we can restrict ourself to thick G-tribes which are concentrated at thick ends.

[5] When G is clear from the context we will often refer to a G-subtribe as simply a subtribe.

Lemma 11.8.6 ([25, 6.7])**.** *Let G be a connected graph and Γ a graph containing a thick G-tribe $\mathcal{F}$ concentrated at an end $\epsilon \in \Omega(\Gamma)$ which is thin. Then $\aleph_0 G \preccurlyeq \Gamma$.*

Given an extensive tree decomposition $(T, \mathcal{V})$ of G, broadly, our strategy will be to obtain a family of disjoint IGs by choosing a sequence of trees $T_0 \subseteq T_1 \subseteq \ldots$ such that $\bigcup T_i = T$ and to construct inductively a family of finitely many $IG[T_{k+1}]$s which extend the $IG[T_k]$s built previously (cf. Definition 11.4.6). The extensiveness of the tree-decomposition ensures that, at each stage, there will be some edges in $\partial(T_i) = E(T_i, T \setminus T_i)$, each of which has in G a family of rays $\mathcal{R}_e$ along which the graph displays self-similarity.

In order to extend our $IG[T_k]$ at each step, we will want to assume that the IGs in $\mathcal{F}$ lie in a 'uniform' manner in the graph Γ in terms of these rays $\mathcal{R}_e$.

More specifically, for each edge $e \in \partial(T_i)$ the rays $\mathcal{R}_e$ tend to a common end ω_e in G, and for each $H \in \bigcup \mathcal{F}$, the corresponding rays in H converge to an end $H(\omega_e) \in \Omega(\Gamma)$ (cf. Definition 11.3.13) which might either be ϵ, or another end of Γ. We would like that our G-tribe $\mathcal{F}$ makes a consistent choice of whether $H(\omega_e)$ is ϵ, for each $e \in \partial(T_i)$.

Furthermore, if $H(\omega_e) = \epsilon$ for every $H \in \bigcup \mathcal{F}$ then this imposes some structure on the end ω_e of G. More precisely with [25, Lemma 9.1] we may assume that $RG_H(H^{\downarrow}(\mathcal{R}_e))$ is a path for each H in the G-tribe $\mathcal{F}$.

By moving to a thick subtribe, we may assume that every ray in every $H \in \bigcup \mathcal{F}$ is core, in which case $\leqslant_\epsilon$ imposes a linear order on every family of rays $H^{\downarrow}(\mathcal{R}_e)$, which induces one of the two distinct orientations of the path $RG_H(H^{\downarrow}(\mathcal{R}_e))$ (reference to make this clear/precise). We will also want that our tribe $\mathcal{F}$ induces this orientation in a consistent manner.

Let us make the preceding discussion precise with the following definitions:

Definition 11.8.7. *Let G be a connected locally finite graph with a extensive tree-decomposition $(T, \mathcal{V})$, S be an initial subtree of T. Let $H \subseteq \Gamma$ be an IG, $\mathcal{H}$ be a set of tidy IGs in Γ and ϵ an end of Γ.*

- *Given an end ω of G, we say that ω* converges to ϵ according to H *if for every ray $R \in \omega$ we have $H^{\downarrow}(R) \in \epsilon$. The end ω* converges to ϵ according to $\mathcal{H}$ *if it converges to ϵ according to every element of $\mathcal{H}$.*

  We say that ω is cut from ϵ according to H *if for every ray $R \in \omega$ we have $H^{\downarrow}(R) \notin \epsilon$. The end ω is* cut from ϵ according to $\mathcal{H}$ *if it is cut from ϵ according to every element of $\mathcal{H}$.*

  Finally we say that $\mathcal{H}$ determines whether ω converges to ϵ *if either ω converges to ϵ according to $\mathcal{H}$ or ω is cut from ϵ according to $\mathcal{H}$.*

- *Given $E \subseteq E(T)$, we say $\mathcal{H}$* weakly agrees about E *if for each $e \in E$, $\mathcal{H}$ determines whether ω_e converges to ϵ. If $\mathcal{H}$ weakly agrees about $\partial(S)$ we let*

  $$\partial_\epsilon(S) := \{e \in \partial(S) \colon \omega_e \textit{ converges to } \epsilon \textit{ according to } \mathcal{H}\},$$
  $$\partial_{\neg\epsilon}(S) := \{e \in \partial(S) \colon \omega_e \textit{ is cut from } \epsilon \textit{ according to } \mathcal{H}\},$$

  and write

  $$S^{\neg\epsilon} \textit{ for the component of the forest } T - \partial_\epsilon(S) \textit{ containing the root of } T,$$
  $$S^{\epsilon} \textit{ for the component of the forest } T - \partial_{\neg\epsilon}(S) \textit{ containing the root of } T.$$

  Note that $S = S^{\neg\epsilon} \cap S^{\epsilon}$.

- *We say that $\mathcal{H}$ is* well-separated from ϵ at S, *if $\mathcal{H}$ weakly agrees about $\partial(S)$ and $H(S^{\neg\epsilon})$ can be separated from ϵ in Γ for all elements $H \in \mathcal{H}$, i.e. for every H there is a finite $X \subseteq V(\Gamma)$ such that $H(S^{\neg\epsilon}) \cap C_\Gamma(X, \epsilon) = \emptyset$.*

In the case that ϵ is half-grid-like, we say that $\mathcal{H}$ strongly agrees *about $\partial(S)$ if*

- *it weakly agrees about $\partial(S)$;*
- *for each $H \in \mathcal{H}$ every ϵ-ray $R \subseteq H$ is in* core(ϵ)*; and*
- *for every $e \in \partial_\epsilon(S)$ there is a linear order $\leqslant_{\mathcal{F},e}$ on $S(e)$ such that the order induced on $H^\downarrow(\mathcal{R}_e)$ by $\leqslant_{\mathcal{F},e}$) agrees with $\leqslant_\epsilon$ on $H^\downarrow(\mathcal{R}_e)$ for all $H \in \mathcal{H}$.*

If $\mathcal{F}$ is a thick G-tribe concentrated at an end ϵ, we use these terms in the following way:

- *Given $E \subseteq E(T)$, we say that $\mathcal{F}$* weakly agrees about *E if $\bigcup \mathcal{F}$ weakly agrees about E w.r.t. ϵ.*
- *We say that $\mathcal{F}$ is* well-separated from ϵ at S *if $\bigcup \mathcal{F}$ is.*
- *We say that $\mathcal{F}$* strongly agrees *about $\partial(S)$ if $\bigcup \mathcal{F}$ does.*

Remark. *We note that the properties of weakly agreeing about E, being well separated from ϵ and strongly agreeing about $\partial(S)$ are all preserved under taking subsets, and hence under taking flat subtribes.*

Note that by the pigeon hole principle for G-tribes, given a finite edge set $E \subset E(T)$, any thick G-tribe $\mathcal{F}$ concentrated at ϵ has a thick (flat) subtribe which weakly agrees about E.

The next few lemmas show that, with some slight modification, we may restrict to a further subtribe which strongly agrees about E and is also well-separated from ϵ.

Definition 11.8.8 ([25]). *Let ω be an end of a graph G. We say ω is* linear *if $RG(\mathcal{R})$ is a path for every finite family $\mathcal{R}$ of disjoint ω-rays.*

Lemma 11.8.9 ([25, 8.1]). *Let ϵ be a non-pebbly end of Γ and let $\mathcal{F}$ be a G-tribe such that for every $H \in \bigcup \mathcal{F}$ there is an end $\omega_H \in \Omega(G)$ such that $H(\omega_H) = \epsilon$. Then there is a thick flat subtribe $\mathcal{F}'$ such that ω_H is linear for every $H \in \bigcup \mathcal{F}'$.*

Corollary 11.8.10. *Let G be a connected locally finite graph with an extensive tree-decomposition $(T, \mathcal{V})$, S be an initial subtree of T, and let $\mathcal{F}$ be a thick G-tribe which is concentrated at a non-pebbbly end ϵ of a graph Γ and weakly agrees about S. Then ω_e is linear for every $e \in \partial_\epsilon(S)$.*

Proof. For any $e \in \partial_\epsilon(S)$ apply Lemma 11.8.9 to $\mathcal{F}$ with $\omega_H = \omega_e$ for each $H \in \bigcup \mathcal{F}$. □

Lemma 11.8.11. *Let G be a connected locally-finite graph with a tree-decomposition $(T, \mathcal{V})$. Let $\mathcal{F}$ be a thick G-tribe in Γ concentrated at ϵ which weakly agrees about some finite $\partial(S) \subset E(T)$. Then $\mathcal{F}$ has a flat thick subtribe $\mathcal{F}'$ so that $\mathcal{F}'$ strongly agrees about $\partial(S)$.*

Proof. □

Lemma 11.8.12. *Let G be a connected locally-finite graph with an extensive tree-decomposition $(T, \mathcal{V})$. Let $H \subseteq \Gamma$ be an IG and ϵ an end of Γ. Let e be an edge of T, such that $H(\omega_e) \neq \epsilon$. There is a finite set $X \subseteq V(G)$ such that for every finite $X' \supseteq X$ there exists a push-out H_e of H along e so that $C_\Gamma(X', G(\omega_e)) \neq C_\Gamma(X', \epsilon)$ and*

1. *$H_e(G[B(e)]) \subseteq C_\Gamma(X', G(\omega_e))$,*

2. $H_e(G[B(e)]) \setminus X \subseteq H(G[B(e')])$ *for an edge* e' *on* R_e*, and*

3. $H_e(G[A(e)])$ *extends* $H(G[A(e)])$ *fixing* $A(e) \setminus S(e)$*.*

Proof. Let $X_1 \subseteq V(\Gamma)$ be a finite vertex set such that $C_\Gamma(X, G(\omega_e)) \neq C_\Gamma(X, \epsilon)$, then given any finite $X' \supseteq X_1$, surely $C_\Gamma(X', G(\omega_e)) \neq C_\Gamma(X', \epsilon)$. Since X_1 is finite, there are only finitely many $v \in G$ whose branch sets $H(v)$ meet X_1. By extensiveness, every vertex of G is contained in only finitely many parts of the tree-decomposition, and so there exists an edge e_1 on R_e with

$$H(G[B(e_1)]) \cap X_1 = \emptyset.$$

For each $s \in S(e)$ let P_s be the initial segment of $R_{e,s}$ up to the first time it meets $S(e_1)$. Let

$$X = X_1 \cup \bigcup_{v \in V(P_s), s \in S(e)} H(v).$$

Then, given any $X' \supseteq X$, as before there is an edge e' on R_e such that

$$H(G[B(e')]) \cap X' = \emptyset.$$

Since $(T, \mathcal{V})$ is an extensive tree-decomposition there is a witness W of the self-similarity of $B(e)$ at distance at least $\max\{\text{dist}(e^-, e_1^-), \text{dist}(e^-, e'^-)\} := n$. Then by Definition 11.4.11 and Lemma 11.4.12 there is a push-out H_e of H along e to depth n.

By Definition 11.4.11 $V(H_e(G[B(e)]) \subseteq V(H_e(W)) \cup X$ and hence (1) and (2) hold, and also $H_e([A(e)])$ extends $H(G[A(e)])$ fixing $A(e) \setminus S(e)$. □

Lemma 11.8.13. *Let G be a connected locally finite graph with an extensive tree-decomposition $(T, \mathcal{V})$ with root r. Let Γ be a graph and $\mathcal{F}$ a thick G-tribe concentrated at a half-grid-like end ϵ of Γ. Then there is a thick sub-tribe $\mathcal{F}'$ of $\mathcal{F}$ such that*

(1) $\mathcal{F}'$ is concentrated at a half-grid-like end ϵ.

(2) $\mathcal{F}'$ strongly agrees about $\partial(\{r\})$.

(3) $\mathcal{F}'$ is well-separated from ϵ at $\{r\}$.

Proof. Since $d(r)$ is finite, by choosing a thick flat subtribe of $\mathcal{F}$, we may assume that $\mathcal{F}$ weakly agrees about $\partial(\{r\})$. Moreover, by Lemma 11.8.11, we may even assume that $\mathcal{F}$ strongly agrees about $\partial(\{r\})$.

For every member H of $\mathcal{F}$, and for every $e \in \partial_{\neg\epsilon}(\{r\})$ there exists by Lemma 11.8.12 a finite set X_e such that for every finite $X' \supseteq X_e$ there is a push-out H_e of H along e so that $C_\Gamma(X', G(\omega_e)) \neq C_\Gamma(X', \epsilon)$ and

1. $H_e(G[B(e)]) \subseteq C_\Gamma(X', G(\omega_e))$,

2. $H_e(G[B(e)]) \setminus X_e \subseteq H(G[B(e')])$ for an edge e' on R_e, and

3. $H_e(G[A(e)])$ extends $H(G[A(e)])$ fixing $A(e) \setminus S(e)$.

Let X be the union of all these X_e together with $H(\{r\})$. For each $e \in \partial_{\neg\epsilon}(\{r\})$ let H_e be the push-out whose existence is guaranteed by the above with respect to this set X.

Let us define an IG

$$H' := \bigcup_{e \in \partial_{\neg\epsilon}(\{r\})} H_e\left(\{r\}^\epsilon \cup T_{e^+}\right).$$

It is straightforward, although not quick, to check that this is indeed an IG and so we will not do this in detail. Briefly, this can be deduced from multiple applications of Defintion 11.4.10 and by (3) all that we need to check is that the extra vertices added to the branch sets of vertices in $S(e)$ are distinct for each edge e. However, this follows from Definition 11.4.11, since these vertices come from $H(\mathcal{R}_e)$ and the rays $R_{e,s}$ and $R_{e',s'}$ are disjoint except in their initial vertex when $s = s'$. Let $\mathcal{F}'$ be the tribe given by $\{F' \colon F \in \mathcal{F}\}$ where $F' = \{H' \colon H \in F\}$ for each $F \in \mathcal{F}$. We claim that $\mathcal{F}'$ satisfies the conclusion of the lemma.

Firstly, we claim that H strongly agrees with H' about $\partial(\{r\})$ for every member H of $\mathcal{F}$. Indeed, by construction for each $e \in \partial_{\neg\epsilon}(\{r\})$, $H'(G[B(e)]) \subseteq C_\Gamma(X', G(\omega_e))$, and hence ω_e is cut from ϵ according to H'. Furthermore, by construction $H(\{r\}^\epsilon) \setminus X = H'(\{r\}^\epsilon) \setminus X$ and so ω_e is converges to ϵ according to H' for every $e \in \partial_{\neg\epsilon}(\{r\})$. In fact, $H^{\downarrow}(\mathcal{R}_e) = H'^{\downarrow}(\mathcal{R}_e)$ for every $e \in \partial_{\neg\epsilon}(\{r\})$. Finally, since $H' \subset H$, and $\mathcal{F}$ strongly agrees about $\partial(\{r\})$ it follows that every ϵ-ray in H' is in $\mathrm{core}(\epsilon)$.

Then, since $\mathcal{F}$ is strongly concentrated at ϵ and strongly agrees about $\partial(\{r\})$ it follows that (1) and (2) hold for $\mathcal{F}'$. It remains to show that $\mathcal{F}'$ is well-separateed from ϵ at $\{r\}$.

However, we claim that for each member H of $\mathcal{F}$ the set X defined above separates $H'(\{r\}^{\neg\epsilon})$ from ϵ in Γ. Indeed,

$$H'(\{r\}^{\neg\epsilon}) = H'(\{r\}) \cup \bigcup_{e \in \partial_{\neg\epsilon}(\{r\})} H'(G[B(e)]),$$

and so $H'(\{r\}^{\neg\epsilon}) \cap C_\Gamma(X, \epsilon) = \emptyset$. It follows that $\mathcal{F}'$ satisfies the conclusion of the lemma.

□

Lemma 11.8.14 (Well-separated push-out)**.** *Let G be a connected locally-finite graph with an extensive tree-decomposition $(T, \mathcal{V})$. Let $H \subseteq \Gamma$ be an IG and ϵ an end of Γ. Let S be a finite subtree of T such that $\{H\}$ is well-separated from ϵ at S and let $f \in \partial_\epsilon(S)$. Then there exists exists a push-out H' of H along f to depth* 0 *(see Definition 11.4.11) such that $\{H'\}$ is well-separated from ϵ at $\tilde{S} = S \cup \{f\}$.*

Proof. Let $X' \subseteq V(\Gamma)$ be a finite set with $H(S^{\neg\epsilon}) \cap C_\Gamma(X', \epsilon) = \emptyset$. If $\partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S) = \emptyset$ then $H' = H$ satisfies the conclusion of the lemma, hence we may assume that $\partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)$ is non-empty.

By applying Lemma 11.8.12 to every $e \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)$, we obtain a finite set $X \supseteq X'$ and a family $(H_e \colon e \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S))$ where each H_e is a push out of H along e such that

1. $H_e(G[B(e)]) \subseteq C_\Gamma(X, H(\omega_e))$,
2. $H_e(G[B(e)]) \subseteq H(G[B(e')])$ for some edge e' on R_e, and
3. $H_e(G[A(e)])$ extends $H(G[A(e)])$ fixing $A(e) \setminus S(e)$.

Let

$$H' := \bigcup_{e \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)} H_e\left(S^\epsilon \cup T_{e^+}\right).$$

As before it is straightforward to check that H' is an IG, and that H' is a push out of H along f to depth 0. We claim that H' is well-separated from ϵ at $\tilde{S}$. Since H is well-separated from ϵ at S there is a finite set X such that $H(S^{\neg\epsilon}) \cap C_\Gamma(X, \epsilon) = \emptyset$. Let

$$\overline{X} = X \cup \bigcup_{e \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)} V\left(H_e(S(e))\right),$$

note that $\overline{X}$ is finite.

It is sufficient to show that $\overline{X}$ separates $H'(G[B(e)])$ from ϵ in Γ for each $e \in \partial_{\neg\epsilon}(\tilde{S})$, since then $\overline{X}$ together with $H'(S)$ separates $H'(S^{\neg\epsilon})$ from ϵ in Γ. Given an edge $e \in \partial_{\neg\epsilon}(\tilde{S})$ either $e \in \partial_{\neg\epsilon}(S)$ or $e \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)$. In the first case, since

$$H'(G[B(e)]) \subseteq \bigcup_{e' \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)} H_{e'}(G[B(e)]) \subseteq H(G[B(e)]) \cup \bigcup_{e' \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)} H_{e'}(S(e')),$$

by (3), it follows that $H'(G[B(e)]) \cap C_\Gamma(\overline{X}, \epsilon) = \emptyset$.

In the second case $e \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)$, and so again it follows from (3) that

$$H'(G[B(e)]) \subseteq H_e(G[B(e)]) \cup \bigcup_{e \neq e' \in \partial_{\neg\epsilon}(\tilde{S}) \setminus \partial(S)} H_{e'}(S(e)).$$

Hence, $H'(G[B(e)]) \cap C_\Gamma(\overline{X}, \epsilon) = \emptyset$.

□

The following lemma contains a large part of the work needed for our inductive construction. The idea behind the statement is the following: At step n in our construction we will have a G-tribe $\mathcal{F}_n$ which agrees about $\partial(T_n)$, which will allows us to extend our $IG[T_n]$s to $IG[T_{n+1}]$s. In order to perform the next stage of our construction we will need to 'refine' $\mathcal{F}_n$ to a G-tribe $\mathcal{F}_{n+1}$ which agrees about the boundary of T_{n+1}.

This would be a relatively simple application of the pigeon hole principle for G-tribes, Lemma 11.8.3, except that in our construction we cannot extend by a member of $\mathcal{F}_{n+1}$ naively. Indeed, suppose we wish to use an IG, say H, to extend an $IG[T_n]$ to an $IG[T_{n+1}]$. There is some subgraph, $H(T_{n+1} \setminus T_n)$, of H which is an $IG[T_{n+1} \setminus T_n]$, however in order to use this to extend the $IG[T_n]$ we first have to link the branch sets of the boundary vertices to this subgraph, and there may be no way to do so without using other vertices of $H(T_{n+1} \setminus T_n)$.

For this reason we ensure the existence of an 'intermediate G-tribe' $\mathcal{F}^*$, which has the property that for each member H of $\mathcal{F}^*$, there are push-outs at arbitrary depth of H which are members of $\mathcal{F}_{n+1}$. This allows us to first link our $IG[T_n]$ to some $H \in \mathcal{F}^*$ and then choose a push-out $H' \in \mathcal{F}_{n+1}$ of H such that $H'(T_{n+1} \setminus T_n)$ avoids the vertices we used to link.

Lemma 11.8.15 (G-tribe refinement lemma)**.** *Let G be a connected locally finite graph with an extensive tree-decomposition $(T, \mathcal{V})$, let S be a subtree of T with $\partial(S)$ finite, and let $\mathcal{F}$ be a thick G-tribe of a graph Γ such that*

(1) $\mathcal{F}$ is concentrated at a half-grid-like end ϵ.

(2) $\mathcal{F}$ strongly agrees about $\partial(S)$.

(3) $\mathcal{F}$ is well-separated from ϵ at S.

Suppose $f \in \partial_\epsilon(S)$ and let $\tilde{S} = S \cup \{f\}$. Then there is a thick flat subtribe $\mathcal{F}^*$ of $\mathcal{F}$ and a thick G-tribe $\mathcal{F}'$ in Γ with the following properties:

(i) $\mathcal{F}'$ is concentrated at ϵ.

(ii) $\mathcal{F}'$ strongly agrees about $\partial(\tilde{S})$.

(iii) $\mathcal{F}'$ is well-separated from ϵ at $\tilde{S}$.

(iv) $\mathcal{F}' \cup \mathcal{F}$ strongly agrees about $\partial(S) \setminus \{f\}$.

(v) $S^{\neg\epsilon}$ w.r.t. $\mathcal{F}$ is a subtree of $\tilde{S}^{\neg\epsilon}$ w.r.t. $\mathcal{F}'$.

(vi) For every $F \in \mathcal{F}^*$ and every $m \in \mathbb{N}$, there is $F' \in \mathcal{F}'$ such that for all $H \in F$ there is an $H' \in F'$ which is a push-out of H to depth m along f.

Proof. For every member H of $\mathcal{F}$ consider a sequence $(H^{(i)} \colon i \in \mathbb{N})$ where $H^{(i)}$ is a push-out of H along f to depth at least i. After choosing a subsequence of $(H^{(i)} \colon i \in \mathbb{N})$ and relabelling (monotonically), we may assume that for each H, the set $\{H^{(i)} \colon i \in \mathbb{N}\}$ weakly agrees on $\partial(\tilde{S})$, i.e. for every $e \in \partial(\tilde{S})$ either $H^{(i)}(R) \in \epsilon$ for every $R \in \omega_e$ and all i or $H^{(i)}(R) \notin \epsilon$ for every $R \in \omega_e$ and all i. Note that a monotone relabelling preserves the property of $H^{(i)}$ being a push-out of H along f to depth at least i.

This uniform behaviour of $(H^{(i)} \colon i \in \mathbb{N})$ on $\partial(\tilde{S})$ for each member H of $\mathcal{F}$ gives rise to a finite colouring $c \colon \bigcup \mathcal{F} \to 2^{\partial(\tilde{S})}$. By Lemma 11.8.3 we may choose a thick flat subtribe $\mathcal{F}_1 \subseteq \mathcal{F}$ such that c is constant on $\bigcup \mathcal{F}_1$.

Recall that by Corollary 11.8.10 for every $e \in \partial_\epsilon(\tilde{S})$ (w.r.t. $\mathcal{F}_1$) the ray graph $RG_G(\mathcal{R}_e)$ is a path. We pick an arbitrary orientation of this path and denote by $\leq_e$ the corresponding linear order on $\mathcal{R}_e$.

Again for every member $H \in \bigcup \mathcal{F}_1$ define

$$d_H \colon \{H^{(i)} \colon i \in \mathbb{N}\} \to \{-1, 0, 1\}^{\partial_\epsilon(\tilde{S})}$$

where

$$d_H(H^{(i)})_e = \begin{cases} 0 & \text{if } H^{(i)}(\mathcal{R}_e) \text{ are not all core rays,} \\ +1 & \text{if } H^{(i)}(\mathcal{R}_e) \text{ are all core rays and } \leqslant_\epsilon \text{ agrees with } \leqslant_e, \\ -1 & \text{if } H^{(i)}(\mathcal{R}_e) \text{ are all core rays and } \leqslant_\epsilon \text{ agrees with } \geqslant_e. \end{cases}$$

Since d_H has finite range we may assume as above, after choosing a subsequence and relabelling, that d_H is constant on $\{H^{(i)} \colon i \in \mathbb{N}\}$ and that $H^{(i)}$ is still a push-out of H along f to depth at least i.

Now consider $d \colon \bigcup \mathcal{F}_1 \to \{-1, 0, 1\}^{\partial_\epsilon(\tilde{S})}$ with $d(H) = d_H(H^{(1)})$ $(= d_H(H^{(i)})$ for all $i)$. Again, we may choose a thick flat subtribe $\mathcal{F}_2 \subseteq \mathcal{F}_1$ such that d is constant on $\mathcal{F}_2$.

Note that no coordinate of d takes the value 0. Indeed, for $e \in \partial_\epsilon(\tilde{S})$ and every layer $F \in \mathcal{F}_2$ the rays in $(H^{(1)}(\mathcal{R}_e) \colon H \in F)$ are disjoint, and for large enough F it cannot be the case that there is a non-core ray in every $H^{(1)}(\mathcal{R}_e)$.

We can now apply Lemma 11.8.14 to each $H^{(i)}$ to obtain $H'^{(i)}$, the collection of which is well-separated from ϵ at $\tilde{S}$. Note that $H'^{(i)}$ is still a push-out of H along f to depth i.

Now let $\mathcal{F}^* = \mathcal{F}_2$ and $\mathcal{F}' = \{\{H'^{(i)} \colon H \in F\} \colon i \in \mathbb{N}, F \in \mathcal{F}^*\}$. Let us verify that these satisfy (i)–(vi). $\mathcal{F}^*$ is concentrated at ϵ because it is a thick flat subtribe of $\mathcal{F}$ by Lemma 11.8.5. By a comparison, layer by layer, since all members of $\mathcal{F}'$ are push-outs of members of $\mathcal{F}^*$ along f, the tribe $\mathcal{F}'$ is also concentrated at ϵ, satisfying (i).

(ii) is satisfied: Since c and d are constant on $\bigcup \mathcal{F}_2$ the collection of the $H^{(i)}$ (for $H \in \bigcup \mathcal{F}_2$) strongly agrees on $\partial(\tilde{S})$, since we have chosen an appropriate subsequence in which $d_H(H^{(i)})$ is constant. The $H'^{(i)}$ are constructed such that this property is preserved. Property (iii) is immediate from the choice of $H'^{(i)}$. Properties (iv) & (v) follow from (2) and the fact that every member of $\mathcal{F}'$ is a push-out of a member of $\mathcal{F}$ along f. Property (vi) is immediate from the construction of $\mathcal{F}'$. $\square$

11.9 The inductive argument

In this section we prove Theorem 11.4.9. Given a connected, locally finite graph G which admits an extensive tree-decomposition $(T, \mathcal{V})$ and a graph Γ which contains a thick G-tribe $\mathcal{F}$, our aim is to construct an infinite family $(Q_i \colon i \in \mathbb{N})$ of disjoint G-minors in Γ inductively.

Our work so far will allow us to make certain assumptions about $\mathcal{F}$. For example, by Lemma 11.8.4 we may assume that $\mathcal{F}$ is concentrated at some end ϵ of Γ, which by Lemma 11.8.6 we may assume is a thick end, and by Lemma 11.6.5 we may assume is not pebbly. Hence, by the work of Section 11.7 we may assume that ϵ is either half-grid-like or grid-like.

At this point our proof will split into two different cases, depending on the nature of ϵ. However, the two cases are very similar, with the grid-like case being significantly simple. Therefore we will first prove Theorem 11.4.9 in the case where ϵ is half-grid-like, and then in Section 11.9.2 we will briefly sketch the differences for the grid-like case.

So, to briefly recap, in the following section we will be working under the standing assumptions that there is a thick G-tribe $\mathcal{F}$ in Γ and an end ϵ of Γ such that

- $\mathcal{F}$ is concentrated at ϵ;
- ϵ is thick;
- ϵ is not pebbly;
- ϵ is half-grid-like.

11.9.1 The half-grid-like case

As explained in Section 11.2, our strategy will be to take some sequence of subtrees $S_1 \subseteq S_2 \subseteq S_3 \dots$ of T, such that $\bigcup_i S_i = T$, and to inductively build a collection of n inflated copies of $G(S_n)$, at each stage extending the previous copies. However, in order to ensure that we can continue the construction at each stage, we will require the existence of additional structure.

Let us pick an enumeration $\{t_i : i \geqslant 0\}$ of $V(T)$ such that t_0 is the root of T and $T_n := T[\{t_i : 0 \leqslant i \leqslant n\}]$ is connected for every $n \in \mathbb{N}$. We will not take the S_n above to be the subtrees T_n, but instead the subtrees $T_n^{\neg\epsilon}$ with respect to some tribe $\mathcal{F}_n$ which weakly agrees about $\partial(T_n)$. This will ensure that every edge in the boundary $\partial(S_n)$ will be in $\partial_\epsilon(T_n)$. For every edge $e \in E(T)$ let us fix a family $\mathcal{R}_e = (R_{e,s} : s \in S(e))$ of disjoint rays witnessing the self-similarity of the bough $B(e)$ towards an end ω_e of G where $\mathrm{init}(R_{e,s}) = s$. By taking $S_n = T_n^{\neg\epsilon}$ we guarantee that for each edge in $e \in \partial(S_n)$, $s \in S(e)$ and every $H \in \bigcup \mathcal{F}_n$ the ray $H^{\downarrow}(R_{e,s})$ is an ϵ-ray.

Furthermore, since $\partial(T_n)$ is finite, we may assume by Lemma 11.8.11 that $\mathcal{F}_n$ strongly agrees about $\partial(T_n)$. We can now describe the additional structure that we require for the induction hypothesis.

At each stage of our construction we will have built some inflated copies of $G(S_n)$, which we wish to extend in the next stage. However, S_n will not in general be a finite subtree, and so we will need some control over where these copies lie in Γ to ensure we have not 'used up' all of Γ. The control we will want is that there is a finite set of vertices X, which we call a *bounder* which separates all we have built so far from the end ϵ. This will guarantee, since $\mathcal{F}$ is concentrated at ϵ, that we can find arbitrarily large layers of $\mathcal{F}$ which are disjoint from what we've built so far.

Furthermore, in order to extend these copies in the next set we will need to be able to link the boundary of our inflated copies of $G(S_n)$ to this large layer of $\mathcal{F}$. To this end we will also want to keep track of some structure which allows us to do this, which we call an *extender*. Let us make the preceding discussion precise.

Definition 11.9.1 (Bounder, extender)**.** *Let $\mathcal{F}$ be a thick G-tribe which is concentrated at ϵ and strongly agrees about $\partial(S)$ for some subtree S of T, and let $k \in \mathbb{N}$. Let $\mathcal{Q} = (Q_i : i \in [k])$ be a family of disjoint inflated copies of $G(S^{\neg\epsilon})$ in Γ (note, $S^{\neg\epsilon}$ depends on $\mathcal{F}$).*

- *A* bounder *for $\mathcal{Q}$ is a finite set X of vertices in Γ separating each Q_i in $\mathcal{Q}$ from ϵ, i.e. such that*

$$C(X,\epsilon) \cap \bigcup_{i=1}^{k} Q_i = \emptyset.$$

- *For $A \subseteq E(T)$,let $I(A,k)$ denote the set $\{(e,s,i) \colon e \in A, s \in S(e), i \in [k]\}$.*

- *An* extender *for $\mathcal{Q}$ is a family $\mathcal{E} = (E_{e,s,i} \colon (e,s,i) \in I(\partial_\epsilon(S),k))$ of ϵ-rays in Γ such that the graphs in $\mathcal{E}^- \cup \mathcal{Q}$ are pairwise disjoint and such that* $\text{init}(E_{e,s,i}) \in Q_i(s)$.

- *Given an extender $\mathcal{E}$, an edge $e \in \partial_\epsilon(S)$ and $i \in [k]$ we let*

$$\mathcal{E}_{e,i} := (E_{e,s,i} \colon s \in S(e)).$$

Recall that, since ϵ is half-grid like, there is a partial order $\leqslant_\epsilon$ defined on the core rays of ϵ, see Lemma 11.7.15. Furthermore, if $\mathcal{F}$ strongly agrees about $\partial(S)$ then, as in Definition 11.8.7, for each $e \in \partial_\epsilon(S)$ there is a linear order $\leqslant_{\mathcal{F},e}$ on $S(e)$.

Definition 11.9.2 (Extension scheme)**.** *Under the conditions above, we call a tuple $(X,\mathcal{E})$ an* extension scheme *for $\mathcal{Q}$ if the following holds:*

(ES1) X is a bounder for $\mathcal{Q}$ and $\mathcal{E}$ is an extender for $\mathcal{Q}$;

(ES2) $\mathcal{E}$ is a family of core rays;

(ES3) the order $\leqslant_\epsilon$ on $\mathcal{E}_{e,i}$ (and thus on $\mathcal{E}^-_{e,i}$) agrees with the order induced by $\leqslant_{\mathcal{F},e}$ on $\mathcal{E}^-_{e,i}$ for all $e \in \partial_\epsilon(S)$ and $i \in [k]$;

(ES4) the sets $\mathcal{E}^-_{e,i}$ are intervals with respect to $\leqslant_\epsilon$ on $\mathcal{E}^-$ for all $e \in \partial_\epsilon(S)$ and $i \in [k]$.

We will in fact split our inductive construction into two types of extensions, which we will do on odd and even steps respectively.

In an even step $n = 2k$, starting with a G-tribe $\mathcal{F}_k$, k disjoint inflated copies of $G(T_k^{\neg\epsilon})$ and an appropriate extension scheme, we will construct Q^n_{k+1}, a further disjoint inflated copy of $G(T_k^{\neg\epsilon})$, and an appropriate extension scheme for everything we built so far.

In an odd step $n = 2k-1$ (for $k \geqslant 1$), starting with the same G-tribe $\mathcal{F}_{k-1}$ from the previous step, k disjoint inflated copies of $G(T_{k-1}^{\neg\epsilon})$ and an appropriate extension scheme, we will refine to a new G-tribe $\mathcal{F}_k$ which strongly agrees on $\partial(T_k)$, extend each copy Q^n_i of $G(T_{k-1}^{\neg\epsilon})$ to a copy Q^{n+1}_i of $G(T_k^{\neg\epsilon})$ for $i \in [k]$, and construct an appropriate extension scheme for everything we built so far.

So, we will assume inductively that for some $n \in \mathbb{N}$, with $r := \lfloor n/2 \rfloor$ and $s := \lceil n/2 \rceil$ we have:

(I1) a thick G-tribe $\mathcal{F}_r$ in Γ which

- is concentrated at ϵ;
- strongly agrees about $\partial(T_r)$;
- is well-separated from ϵ at T_r; and
- whenever $l < k \le r$, $T_k^{\neg\epsilon}$ with respect to $\mathcal{F}_k$ is a sub-tree of $T_l^{\neg\epsilon}$ with respect to $\mathcal{F}_l$.

(I2) a family $\mathcal{Q}_n = (Q^n_i \colon i \in [s])$ of s pairwise disjoint inflated copies of $G(T_r^{\neg\epsilon})$ (where $T_r^{\neg\epsilon}$ is considered with respect to $\mathcal{F}_r$) in Γ;
if $n \geqslant 1$, we additionally require that Q^n_i extends Q^{n-1}_i for all $i \leqslant s-1$;

(I3) an extension scheme $(X_n, \mathcal{E}_n)$ for $\mathcal{Q}_n$;

(I4) if n is even and $\partial_\epsilon(T_r) \neq \emptyset$, we require that there is a set $\mathcal{J}_r$ of disjoint core ϵ-rays disjoint to $\mathcal{E}_n$ with $|\mathcal{J}_r| \geqslant (|\partial_\epsilon(T_r)| + 1) \cdot |\mathcal{E}_n|$.

Suppose we have inductively constructed $\mathcal{Q}_n$ for all $n \in \mathbb{N}$. Let us define $H_i := \bigcup_{n \geqslant 2i-1} Q_i^n$. Since $T_k^{\neg\epsilon}$ with respect to $\mathcal{F}_k$ is a sub-tree of $T_l^{\neg\epsilon}$ with respect to $\mathcal{F}_l$ for all $k < l$, we have $\bigcup_{n\in\mathbb{N}} T_n^{\neg\epsilon} = T$ (where we considered $T_n^{\neg\epsilon}$ w.r.t. $\mathcal{F}_n$), and due to the extension property (I2), the collection $(H_i : i \in \mathbb{N})$ is an infinite family of disjoint G-minors, as required.

So let us start the construction. To see that our assumptions for the case $n = 0$ we first note that since $T_0 = t_0$, by Lemma 11.8.13 there is a thick subtribe $\mathcal{F}_0$ of $\mathcal{F}$ which satisfies (I1). Let us further take

- $\mathcal{Q}_0 = \mathcal{E}_0 = X_0 = \emptyset$;
- $\mathcal{J}_0$ be any suitably large set of disjoint core rays of ϵ.

The following notation will be useful throughout the construction. Given $e \in E(T)$ and some inflated copy H of G, recall that $H^\downarrow(\mathcal{R}_e)$ denotes the family $(H^\downarrow(R_{e,s}) : s \in S(e))$. Given a G-tribe $\mathcal{F}$, a layer $F \in \mathcal{F}$ and a family of rays $\mathcal{R}$ in G we will write $F^\downarrow(\mathcal{R}) = (H^\downarrow(R) : H \in F, R \in \mathcal{R})$.

Construction part 1: $n = 2k$ is even

Case 1: $\partial_\epsilon(T_k) = \emptyset$.

In this case $T_k^{\neg\epsilon} = T$ and so picking any member $H \in \mathcal{F}_k$ with $H \subseteq C(X_n, \epsilon)$ and setting $Q_{k+1}^{n+1} = H(T_k^{\neg\epsilon})$ gives us a further inflated copy of $G(T_k^{\neg\epsilon})$ disjoint from all the previous ones. We set $Q_i^{n+1} = Q_i^n$ for all $i \in [k]$ and $\mathcal{Q}_{n+1} = (Q_i^{n+1} : i \in [k+1])$. Using that $\mathcal{F}_k$ is well-separated from ϵ at T_k, there is a suitable bounder $X_{n+1} \supseteq X_n$ for $\mathcal{Q}_{n+1}$. Then $(X_{n+1}, \emptyset)$ is an extension scheme for $\mathcal{Q}_{n+1}$ while $\mathcal{F}_k$ remains unchanged.

Case 2: $\partial_\epsilon(T_k) \neq \emptyset$. (See Figure 11.9.1)

Consider the family $\mathcal{R}^- := \bigcup\{\mathcal{R}_e^- : e \in \partial_\epsilon(T_k)\}$. Moreover, set $\mathcal{C} := \mathcal{E}_n^- \cup \mathcal{J}_k$ and consider $\overline{\mathcal{C}}$ as in Definition 11.7.17. Let $Y \subseteq C(X_n, \epsilon)$ be a finite set which is a transition box between $\overline{\mathcal{E}_n^-}$ and $\overline{\mathcal{C}}$ as in Lemma 11.3.17. Let $\mathcal{F}'$ be a flat thick G-subtribe of $\mathcal{F}_k$ such that each member of $\mathcal{F}'$ is contained in $C(X_n \cup Y, \epsilon)$, which exists by Lemma 11.8.5 since both X_n and Y are finite.

Let R be an arbitrary element of $\mathcal{R}$. Let $F \in \mathcal{F}'$ be large enough such that we may apply Lemma 11.3.16 to find a transitional linkage $\mathcal{P} \subseteq C(X_n \cup Y, \epsilon)$ from $\overline{\mathcal{C}}$ to $F^\downarrow(\mathcal{R}^-)$ after $X_n \cup Y$ avoiding some member $H \in F$. Note that, since X_n is a bounder and $\mathcal{P} \subseteq C(X_n \cup Y, \epsilon)$, $\mathcal{P}$ is disjoint from all $\mathcal{Q}_n$ and Y.

Let
$$Q_{k+1}^{n+1} := H(T_k^{\neg\epsilon}).$$
Note that Q_{k+1}^{n+1} is an inflated copy of $G(T_k^{\neg\epsilon})$. Moreover let $Q_i^{n+1} := Q_i^n$ for all $i \in [k]$ and $\mathcal{Q}_{n+1} := (Q_i^{n+1} : i \in [k+1])$, yielding property (I2).

Since $\mathcal{F}_k$ is well-separated from ϵ at T_k, and $H \in \bigcup \mathcal{F}_k$, there is a finite set $X_{n+1} \subseteq \Gamma$ containing $X_n \cup Y$ such that $C(X_{n+1}, \epsilon) \cap Q_{k+1}^{n+1} = \emptyset$. This set X_{n+1} is a bounder for $\mathcal{Q}_{n+1}$.

Since $\mathcal{P}$ is transitional, Remark 1 implies that the linkage is preserving on $\mathcal{C}$. Since all rays in $F^\downarrow(\mathcal{R}^-)$ are core rays, $\leq$ is a linear order on $F^\downarrow(\mathcal{R}^-)$. Moreover, for each $e \in \partial_\epsilon(T_k)$, the rays in $H^\downarrow(\mathcal{R}_e)$ correspond to an interval in this order. Thus, deleting these intervals from $F^\downarrow(\mathcal{R}^-)$ leaves behind at most $|\partial_\epsilon(T_k)| + 1$ intervals in $F^\downarrow(\mathcal{R}^-)$ (with respect to $\leq$) which do not contain any rays in $H^{\downarrow}(\mathcal{R})$. Since $|\mathcal{J}_k| \geqslant (|\partial_\epsilon(T_k)|+1)\cdot|\mathcal{E}_n|$, by the pigeonhole principle there is such an interval on $F^{\downarrow}(\mathcal{R}^-)$ that

– does not contain rays in $H^{\downarrow}(\mathcal{R})$; and

– where a subset $\mathcal{P}' \subseteq \mathcal{P}$ of size $|\mathcal{E}_n^-|$ links a corresponding subset $\mathcal{A}' \subseteq \mathcal{A}$ of $\mathcal{C}$ to rays $\mathcal{B}$ in that interval.

By Lemma 11.7.20 and Remark (1 and 3), and Lemma 11.3.17 there is a linkage $\mathcal{P}''$ from $\mathcal{E}_n^-$ to $\mathcal{A}$ contained in $\Gamma[Y]$ which is preserving on $\mathcal{E}_n^-$.

For $e \in \partial_\epsilon(T_k)$ and $s \in S(e)$ define

$$E^{n+1}_{e,s,k+1} = H^{\downarrow}(R_{e,s}) \text{ for the corresponding ray } R_{e,s} \in \mathcal{R}_e.$$

and moreover for each $i \in [k]$, we define

$$E^{n+1}_{e,s,i} = \operatorname{init}(E^n_{e,s,i})(E^-_{e,s,i} \circ_{\mathcal{P}''} \mathcal{A}) \circ_{\mathcal{P}'} \mathcal{B}$$

By construction, all these rays are, except for their first vertex, disjoint from $\mathcal{Q}_{n+1}$. Moreover, $\mathcal{E}_{n+1} := (E^{n+1}_{e,s,i} \colon (e,s,i) \in I(\partial_\epsilon(T_k), k+1))$ is an extender for $\mathcal{Q}_{n+1}$. Note that each ray in $\mathcal{E}_{n+1}$ shares a tail with a ray in $F^{\downarrow}(\mathcal{R}^-)$.

We claim that $(X_{n+1}, \mathcal{E}_{n+1})$ is an extension scheme for $\mathcal{Q}_{n+1}$ and hence property (I3) is satisfied. Since every ray in $\mathcal{E}_{n+1}$ has a tail which is also a tail of a ray in $F^{\downarrow}(\mathcal{R}^-)$, property (ES2) is satisfied by Remark 1. Since $\mathcal{P}'$ is preserving on $\mathcal{A}'$ and $\mathcal{P}''$ is preserving on $\mathcal{E}_n^-$, Remark 4 implies that the linkage $\mathcal{P}'' + \mathcal{P}'$ is preserving on $\mathcal{E}_n^-$. Hence property (ES3) holds for each $i \in [k]$. Furthermore, since $E^{n+1}_{e,s,k+1} = H^{\downarrow}(R_{e,s})$ for each $e \in \partial_\epsilon(T_k)$ and $s \in S(e)$, it is clear that property (ES3) holds for $i = k+1$. Finally, property (ES4) holds for $i = k+1$ since for each $e \in \partial_\epsilon(T_k)$, the rays in $H^{\downarrow}(\mathcal{R}_e)$ are an interval with respect to $\leqslant_\epsilon$ on $F^{\downarrow}(\mathcal{R}^-)$, and it holds for $i \in [k]$ by the fact that $\mathcal{P}'' + \mathcal{P}'$ is preserving on $\mathcal{E}_n^-$ together with the fact that $\mathcal{P}'' + \mathcal{P}'$ is preserving on $\mathcal{E}_n^-$ links $\mathcal{E}_n^-$ to an interval of $F^{\downarrow}(\mathcal{R}^-)$ containing no ray in $H^{\downarrow}(\mathcal{R})$.

Finally note that (I1) is still satisfied by $\mathcal{F}_k$ and T_k, and (I4) is vacuously satisfied.

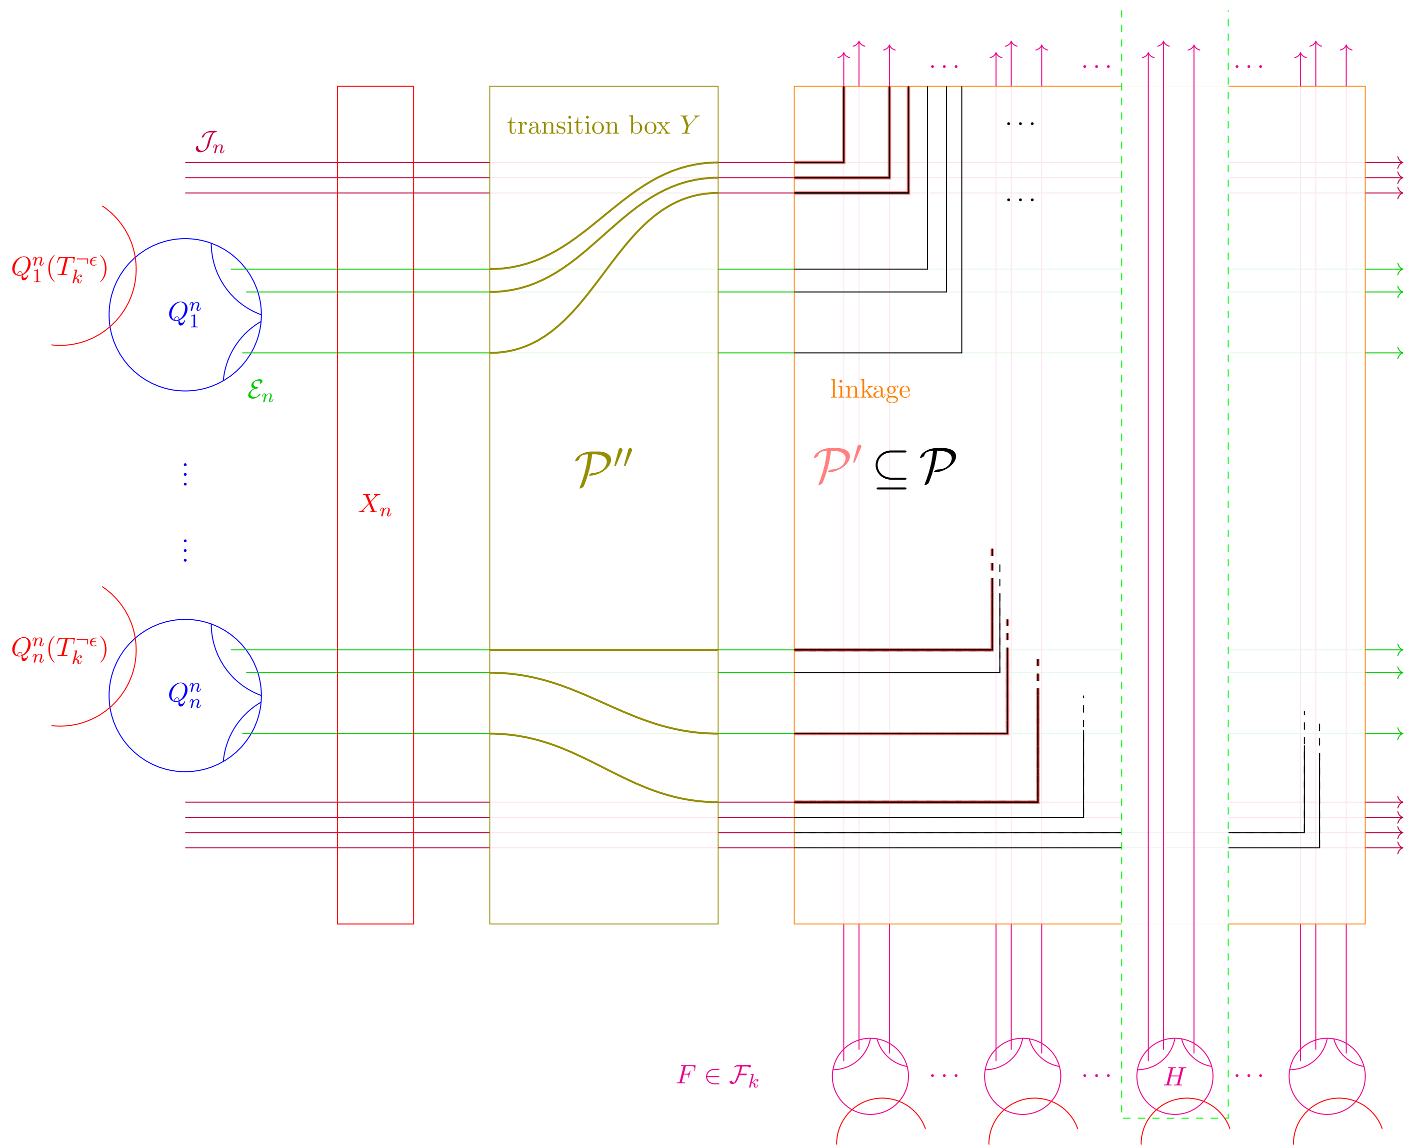

Figure 11.3: Adding a new copy when $n = 2k$ is even.

Construction part 2: $n = 2k - 1$ **is odd (for** $k \geqslant 1$**).**

Let f denote the unique edge of T between T_{k-1} and $T_k \setminus T_{k-1}$.

Case 1: $f \notin \partial_\epsilon(T_{k-1})$.

Let $\mathcal{F}_k := \mathcal{F}_{k-1}$. Since $\mathcal{F}_{k-1}$ is well separated from ϵ at T_k it follows that $e \in \partial_{\neg\epsilon}(T_k)$ for every $e \in \partial(T_k) \setminus \partial(T_{k-1})$. Hence $T_k^{\neg\epsilon} = T_{k-1}^{\neg\epsilon}$ and $\partial_\epsilon(T_{k-1}) = \partial_\epsilon(T_k)$, and so we can simply take $\mathcal{Q}_{n+1} := \mathcal{Q}_n$, $\mathcal{E}_{n+1} := \mathcal{E}_n$, $\mathcal{J}_k := \mathcal{J}_{k-1}$ and $X_{n+1} := X_n$ to satisfy (I1), (I2), (I3) and (I4).

Case 2: $f \in \partial_\epsilon(T_{k-1})$. (See Figure 11.9.1)

By (I1) we can apply Lemma 11.8.15 to $\mathcal{F}_{k-1}$ in order to find a thick G-tribe $\mathcal{F}_k$ and a thick flat sub-tribe $\mathcal{F}^*$ of $\mathcal{F}_{k-1}$, both concentrated at ϵ, satisfying properties (i)–(vi) from that lemma. It follows that $\mathcal{F}_k$ satisfies (I1) for the next step.

Let $F \in \mathcal{F}^*$ be a layer of $\mathcal{F}^*$ such that

$$|F| \geqslant (\partial_\epsilon(T_k) + 2) \cdot |I(\partial_\epsilon(T_k), k)|$$

and consider the rays $F^{\downarrow}(\mathcal{R}_f)$. Consider the rays in the extender corresponding to the edge f, that is $\mathcal{E}_f := (E^n_{f,s,i} \colon i \in [k], s \in S(f))$. By Lemma 11.7.20, there is, for every subset $\mathcal{S}$ of $F^{\downarrow}(\mathcal{R}_f)$ of size $|\mathcal{E}_f^-|$ a transitional linkage $\mathcal{P} \subseteq C(X_n, \epsilon)$ from $\overline{\mathcal{E}_n^-}$ to $\overline{F^{\downarrow}(\mathcal{R}_f)}$ after $X_n \cup \mathrm{init}(\mathcal{E}_n)$ such that $\mathcal{P}$ links $\mathcal{E}_f$ to $\mathcal{S}$, if we view it as a linkage from $\overline{\mathcal{E}_n}$ to $\overline{F^{\downarrow}(\mathcal{R}_f)}$. Since all rays in $\mathcal{E}_f$ and in $F^{\downarrow}(\mathcal{R}_f)$ are core rays, any such linkage is preserving on $\mathcal{E}_f$.

Let us choose $H_1, H_2, \dots, H_k \in F$ and let $\mathcal{S} = \left(H_i^{\downarrow}(R_{f,s}) \colon i \in [k], s \in S(f)\right)$. Let $\mathcal{P}$ be the linkage given by the previous paragraph, which we recall is preserving on $\mathcal{E}_f$. Since for every $i \leq k$ the family $\left(E^n_{f,s,i} \colon s \in S(f)\right)$ forms an interval in $\mathcal{E}_n$ and the set $H^{\downarrow}(\mathcal{R}_f)$ forms an interval in $F^{\downarrow}(\mathcal{R}_f)$ it follows that, after perhaps relabelling the H_i, for every $i \in [k]$ and $s \in S(f)$, $\mathcal{P}$ links $E^n_{f,s,i}$ to $H_i^{\downarrow}(R_{f,s})$.

Let $Z \subseteq V(\Gamma)$ be a finite set such that $\top(\omega, R)$ and $\bot(\omega, R)$ are separated by Z in $\Gamma - V(R)$ for all $R \in F^{\downarrow}(\mathcal{R}_f)$ (cf. Lemma 11.7.16).

Since $|F|$ is finite and $(T, \mathcal{V})$ is an extensive tree-decomposition there exists an $m \in \mathbb{N}$ such that if $e \in R_f$ with $\mathrm{dist}(f^-, e^-) = m$ then $H(B(e)) \cap (X_n \cup Z \cup V(\bigcup \mathcal{P})) = \emptyset$. Let $\vec{F} \in \mathcal{F}_k$ be as in Lemma 11.8.15(vi) for F with such an m.

Hence, by definition, for each $H_i \in F$ there is some subgraph $W_i \subseteq H(B(e))$ which is an $\overline{IG[B(f)]}$ such that for each $s \in S(f)$, $W_i(s)$ contains the first point of W_i on $H_i^{\downarrow}(R_{f,s})$.

For each $i \in [k]$ we construct Q_i^{n+1} from Q_i^n as follows. Consider the part of G that we want to add $G(T_{k-1}^{\neg\epsilon})$ to obtain $G(T_k^{\neg\epsilon})$, namely

$$D := \overline{G[B(f)]}\left[V_{f^+} \cup \bigcup_{e \in \partial_{\neg\epsilon}(T_k) \setminus \partial_{\neg\epsilon}(T_{k-1})} B(e)\right].$$

Let $K_i := W_i(D)$. Note that, this is an inflated copy of D and for each $s \in S(f)$ and each $i \in [k]$ the branch set $K_i(s)$ contains the first point of K_i on $H_i^{\downarrow}(R_{f,s})$.

Note further that by the choice of m, all the K_i are disjoint to $\mathcal{Q}_n$. Let $x_{f,s,i}$ denote the first vertex on the ray $H_i^{\downarrow}(R_{f,s})$ in K_i, and let

$$O_{s,i} := (E^n_{f,s,i} \circ_{\mathcal{P}} F(\mathcal{R}_f))x_{f,s,i}.$$

Then, if we let $\mathcal{O}_i := (O_{s,i} \colon s \in S(f))$ and $\mathcal{O} = (O_{s,i} \colon s \in S(f), i \in [k])$, we see that

$$Q_i^{n+1} := Q_i^n \oplus_{\mathcal{O}_i} K_i$$

(see Definition 11.4.10) is an inflated copy of $G(T_k^{\neg\epsilon})$ extending Q_i^n. Hence,

$$\mathcal{Q}^{n+1} := (Q_i^{n+1} \colon i \in [k])$$

is a family satisfying (I2).

Since $\mathcal{F}_k$ is well-separated from ϵ at T_k, and each K_i is a subgraph of the restriction of $\vec{H}_i$ to D, for each K_i there is a finite set $\hat{X}_i$ separating K_i from ϵ, and hence the set

$$X_{n+1} := X_n \cup \bigcup_{i\in[k]} \hat{X}_i \cup V\left(\bigcup \mathcal{O}\right)$$

is a bounder for $\mathcal{Q}^{n+1}$.

For $e \in \partial_\epsilon(T_{k-1}) \setminus \{f\}$, $s \in S(e)$ and $i \in [k]$ we set

$$E_{e,s,i}^{n+1} = E_{e,s,i}^n \circ_{\mathcal{P}} F^{\downarrow}(\mathcal{R}_f),$$

and set

$$\mathcal{E}' := \left(E_{e,s,i}^{n+1} \colon (e,s,i) \in I\left(\partial_\epsilon(T_{k-1}) \setminus \{f\}, k\right)\right)$$

Moreover, for $e \in \partial_\epsilon(T_k) \setminus \partial_\epsilon(T_{k-1})$, $s \in S(e)$ and $i \in [k]$ we set

$$E_{e,s,i}^{n+1} = \vec{H}_i^{\downarrow}(R_{e,s}),$$

and set

$$\mathcal{E}'' := \left(E_{e,s,i}^{n+1} \colon (e,s,i) \in I\left(\partial_\epsilon(T_k) \setminus \partial_\epsilon(T_{k-1}), k\right)\right).$$

Note that, by construction, such a ray has its initial vertex in the branch set $Q_i^{n+1}(s)$ and is otherwise disjoint to $\bigcup \mathcal{Q}_{n+1}$. We set $\mathcal{E}_{n+1} := \mathcal{E}' \cup \mathcal{E}''$. It is easy to check that this is an extender for $\mathcal{Q}_{n+1}$.

We claim that $(X_{n+1}, \mathcal{E}_{n+1})$ is an extension scheme. Property (ES1) is apparent. Since the G-tribes $\mathcal{F}_k$ and $\mathcal{F}^*$ both strongly agree about $\partial(T_k)$, and every ray in $\mathcal{E}_{n+1}$ shares a tail with a ray in a member of $\mathcal{F}_k$ or $\mathcal{F}^*$ it follows that all rays in $\mathcal{E}_{n+1}$ are core rays, and so (ES2) holds.

For any $e \in \partial_\epsilon(T_{k-1}) \setminus \{f\}$ and $i \in [k]$ the rays $(\mathcal{E}_{n+1})_{e,i}$ are a subfamily of $\mathcal{E}'$, obtained by transitioning from the family $(\mathcal{E}_n)_{e,i}$ to $F^{\downarrow}(\mathcal{R}_f)$ along linkage $\mathcal{P}$. By the induction hypothesis $\leqslant_\epsilon$ agreed with the order induced by $\leqslant_{\mathcal{F}_{k-1},e}$ on $(\mathcal{E}_n)_{e,i}$, and since $\mathcal{F}_k \cup \mathcal{F}_{k-1}$ strongly agrees about $\partial_\epsilon(T_{k-1}) \setminus \{f\}$, this is also the order induced by $\leqslant_{\mathcal{F}_k,e}$. Hence, since $\mathcal{P}$ is preserving, by Remark 1 it follows that the order induced by $\leqslant_{\mathcal{F}_k,e}$ on $(\mathcal{E}_{n+1})_{e,i}$ agrees with $\leqslant_\epsilon$.

For for $e \in \partial_\epsilon(T_k) \setminus \partial_\epsilon(T_{k-1})$ and $i \in [k]$ the rays $(\mathcal{E}_{n+1})_{e,i}$ are $(\vec{H}_i^{\downarrow}(R_{e,s}) \colon s \in S(e))$. Since $\vec{H}_i \in \vec{F} \in \mathcal{F}_k$ and $\mathcal{F}_k$ strongly agrees about $\partial(T_k)$, it follows that the order induced by $\leqslant_{\mathcal{F}_k,e}$ on $(\mathcal{E}_{n+1})_{e,i}$ agrees with $\leqslant_\epsilon$. Hence Property (ES3) holds.

Finally, by Lemma 11.3.20 it is clear that for any $e \in \partial_\epsilon(T_{k-1}) \setminus \{f\}$ and $i \in [k]$ the rays $(\mathcal{E}_{n+1}^-)_{e,i}$ form an interval with respect to $\leqslant_\epsilon$ on $\mathcal{E}_{n+1}^-$, since they are each contained in a connected subgraph $\vec{H}_i$ to which the tails of the rest of $\mathcal{E}_{n+1}^-$ are disjoint. Furthermore, by choice of Z and Lemma 11.7.16 it it clear that, since $\mathcal{P}$ is preserving on $\mathcal{E}_n^-$, for each $e \in \partial_\epsilon(T_k) \setminus \partial_\epsilon(T_{k-1})$ and $i \in [k]$ the rays $(\mathcal{E}_{n+1}^-)_{e,i}$ also form an interval with respect to $\leqslant_\epsilon$ on $\mathcal{E}_{n+1}^-$. Hence property (ES4) holds and therefore (I3) is satisfied for the next step.

For property (I4) we note that every ray in $\mathcal{E}_{n+1}$ has a tail in some $H \in F \in \mathcal{F}^*$. Since there is at least one core ϵ-ray in each $H \in F \in \mathcal{F}^*$, we can find family of at least $|F| - |\mathcal{E}_{n+1}|$ such rays. However since

$$|F| \geqslant (\partial_\epsilon(T_k) + 2) \cdot |\mathcal{E}_{n+1}|$$

it follows that we can find a suitable family $|\mathcal{J}_k|$.

This concludes the induction step. $\square$

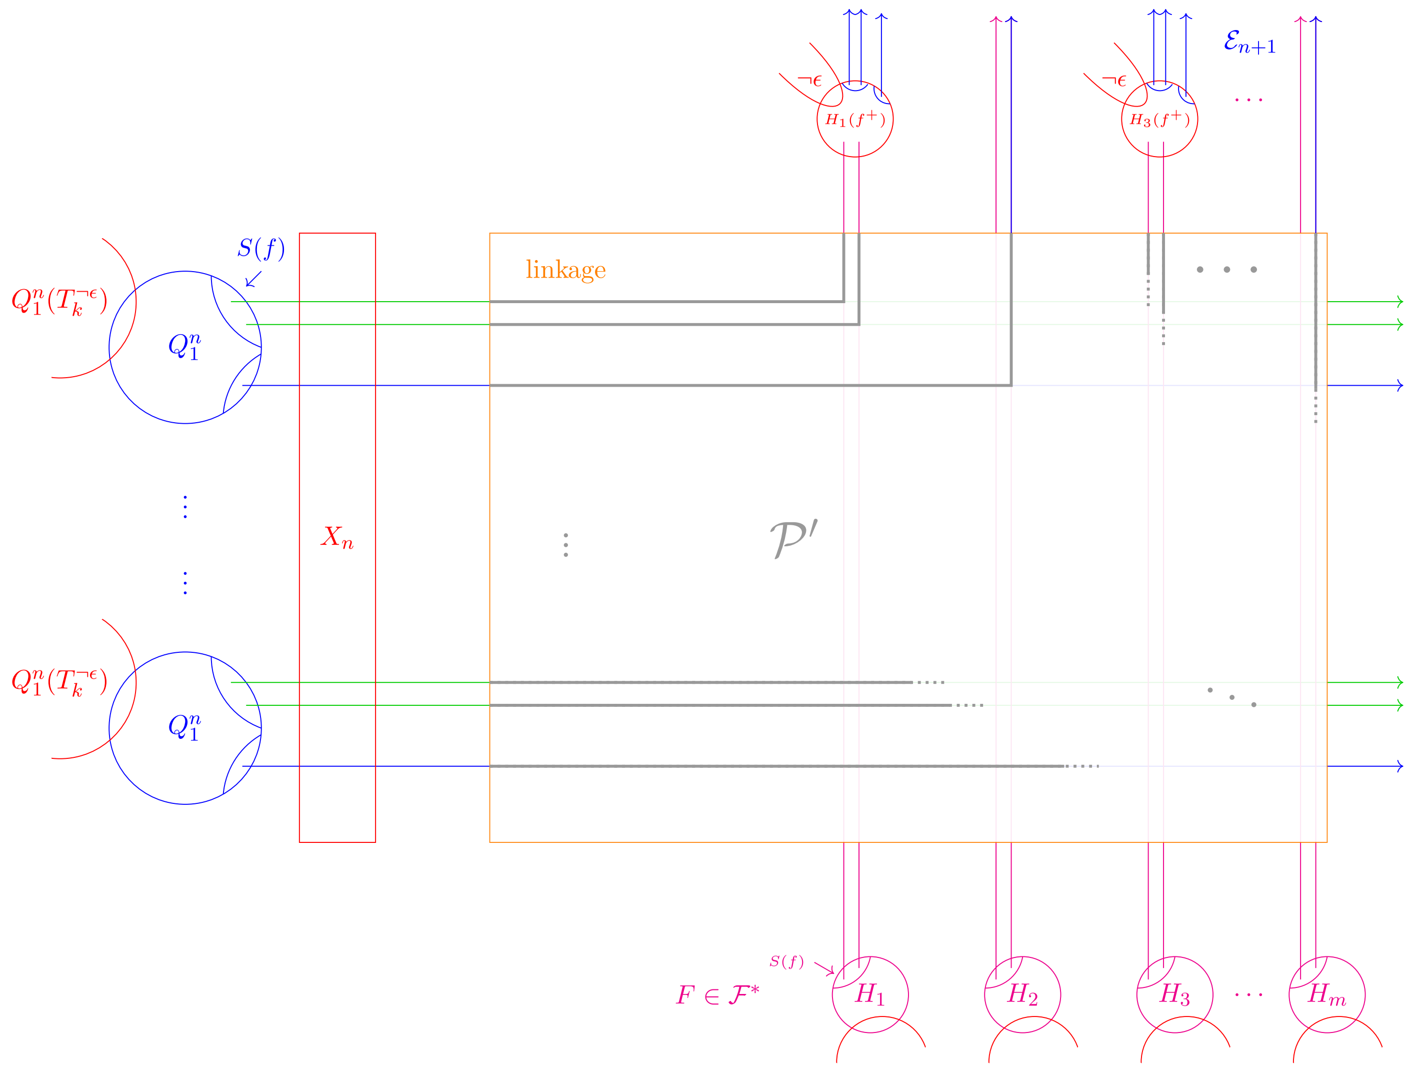

Figure 11.4: Extending the copies when $n = 2k - 1$ is odd.

11.9.2 The grid-like case

In this section we will give a brief sketch of how the argument differs in the case where the end ϵ, towards which we may assume our G-tribe $\mathcal{F}$ is concentrated, is grid-like.

In the case where ϵ is half-grid-like we showed that the end ϵ had a roughly linear structure, in the sense that there is a global partial order $\leqslant_\epsilon$ which is defined on almost all of the ϵ-rays, namely the core ones, such that every pair of disjoint core rays are comparable, and that this order determines the relative structure of any finite family of disjoint core rays, since it determines the ray graph.

Since, by Corollary 11.8.10, $RG_G(\mathcal{R}_e)$ is a path whenever $e \in \partial_\epsilon(T_k)$, there are only two ways that $\leqslant_\epsilon$ can order $H^\downarrow(\mathcal{R}_e)$, and, since $\partial_\epsilon(T_k)$ is finite, by various pigeon-hole type arguments we can assume that it does so consistently for each $H \in \bigcup \mathcal{F}_k$ and each $\mathcal{E}_{e,i}$.

We use this fact crucially in part 2 of the construction, where we wish to extend the graphs $(Q_i^n : i \in [k])$ from inflated copies of $G(T_{k-1}^{\neg\epsilon})$ to inflated copies of $G(T_k^{\neg\epsilon})$ along an edge $e \in \partial(T_{k-1})$. We wish to do so by constructing a linkage from the extender $\mathcal{E}_n$ to some layer $F \in \mathcal{F}_k$, using the self-similarity of G to find an inflated copy of $G(e^+)$ which is 'rooted' on the rays $H^\downarrow(\mathcal{R}_e)$ and extending each Q_i^n by such a subgraph.

However, for this step to work it is necessary that the linkage from $\mathcal{E}_n$ to F is such that for each $i \in [k]$ there is some $H \in F$ such that ray $E_{e,s,i}$ is linked to $H^\downarrow(R_{e,s})$ for each $s \in S(e)$. However, since any transitional linkage we construct between $\mathcal{E}$ and a layer $F \in \mathcal{F}_n$ will respect $\leqslant_\epsilon$, we can use a transition box to 're-route' our linkage such that the above property holds.

In the case where ϵ is grid-like we would like to say that the end has a roughly cyclic structure, in the sense that there is a global 'partial cyclic order' C_ϵ, defined again on almost all of the ϵ-rays which will again determine the relative structure of any finite family of disjoint 'core' rays.

As before, since $RG_G(\mathcal{R}_e)$ is a path whenever $e \in \partial_\epsilon(T_n)$, there are only two ways that C_ϵ can order $H^\downarrow(\mathcal{R}_e)$ ('clockwise' or 'anti-clockwise') and so we can use similar arguments to assume that it does so consistently for each $H \in \bigcup \mathcal{F}_k$ and each $\mathcal{E}_{e,i}$, which allows us as before to control the linkages we build.

To this end, suppose ϵ is a grid-like end, and that N is a number such that no family of disjoint ϵ-rays has a ray graph which is N-pebble win. We say that an ϵ-ray R is a *core ray (of ϵ)* if there is some finite family $(R_i : i \in [n])$ of $n \geqslant N+3$ disjoint ϵ-rays such that $R = R_i$ for some $i \in [n]$[6].

Every large enough ray graph is a cycle, which has a correct orientation by Lemma 11.7.2 and we would like to say that this orientation is induced by a global 'partial cyclic order' defined on the core rays of ϵ.

By a similar argument as in Section 11.7.3 one can show the following:

Lemma 11.9.3. *Let R and R' be disjoint core rays of ϵ. Then in $G - (V(R) \cup V(R'))$ the end ϵ splits into precisely two different ends.*

Definition 11.9.4. *Let R and R' be a core ray of ϵ. We denote by $\top(\epsilon, R, R')$ the end of $G - (V(R) \cup V(R'))$ containing rays which appear between R and R' according to the correct orientation of some ray graph and by $\bot(\epsilon, R, R')$ the end of $G - (V(R) \cup V(R'))$ containing rays which appear between R' and R in the correct orientation of some ray graph.*

We will model our global 'partial cyclic order' as a ternary relation on the set of core rays of ϵ. That is, a *partial cyclic order* on a set X is a relation $C \subset X^3$ written $[a,b,c]$ satisfying the following axioms:

- If $[a,b,c]$ then $[b,c,a]$.

[6] We note that it is possible to show that, if ϵ is grid-like, then in fact $N = 3$.

- If $[a, b, c]$ then not $[c, b, a]$.
- If $[a, b, c]$ and $[a, c, d]$ then $[a, b, d]$.

Lemma and Definition 11.9.5. *Let* $\mathrm{core}(\epsilon)$ *denote the set of core rays of* ϵ*. We define a partial cyclic order* C_ϵ *on* $\mathrm{core}(\epsilon)$ *as follows:*

$$[R, S, T] \text{ if and only if } R, S, T \text{ have disjoint tails } xR, yS, zT \text{ and } yS \in \top(\epsilon, xR, zT).$$

Then, for any disjoint family of at least $N + 3$ ϵ*-rays* $(R_i : i \in [n])$ *the cyclic order induced on* $(R_i : i \in [n])$ *by* C_ϵ *agrees with the correct orientation.*

Again by a similar argument as in Section 11.7.3 on can show that this relation is in fact a partial cyclic order and that it always agrees with the correction orientation of large enough ray graphs. Furthermore, by Lemma 11.7.3, given two families $\mathcal{R}$ and $\mathcal{S}$ of at least $N + 3$ disjoint ϵ-rays, every transitional linkage between $\mathcal{R}$ and $\mathcal{S}$ *preserves* C_ϵ, for the obvious definition of preserving.

Given a disjoint family of ω-rays $\mathcal{R} = (R_i : i \in [n])$ with a linear order $\leqslant$ on $\mathcal{R}$ we say that $\leqslant$ *agrees* with C_ϵ if $[R_i, R_j, R_k]$ whenever $R_i < R_j < R_k$.

Recall that, given a family $F = (f_i : i \in I)$ and a linear order $\leqslant$ on I we denote by $F(\leqslant)$ the linear order on F induced by $\leqslant$, i.e. the order defined by $f_i F(\leqslant) f_j$ if and only if $i \leqslant j$.

We say $\mathcal{F}$ *strongly agrees about* $\partial(T_n)$ if

- it weakly agrees about $\partial(T_n)$;
- for each $H \in \bigcup \mathcal{F}$ every ϵ-ray $R \subseteq H$ is in $\mathrm{core}(\epsilon)$; and
- for every $e \in \partial_\epsilon(T_n)$ there is a linear order $\leqslant_{\mathcal{F},e}$ on $S(e)$ such that $H^\downarrow(\mathcal{R}_e)(\leqslant_{\mathcal{F},e})$ agrees with C_ϵ on $H^\downarrow(\mathcal{R}_e)$ for all $H \in \bigcup F$.

Using this definition the G-tribe refinement lemma (Lemma 11.8.15) can also be shown to hold in the case where ω is a grid-like-end.

Furthermore we modify the definition of an extension scheme for a family of disjoint inflated copies of $G(T_n^{\neg\epsilon})$.

Definition 11.9.6 (Extension scheme)**.** *Let* $\mathcal{Q} = (Q_i : i \in [k])$ *be a family of disjoint inflated copies of* $G(S^{\neg\epsilon})$ *and* $\mathcal{F}$ *be a* G*-tribe which strongly agrees about* $\partial(S)$*. We call a tuple* $(X, \mathcal{E})$ *an* extension scheme *for* $\mathcal{Q}$ *if the following holds:*

(ES1) X *is a bounder for* $\mathcal{Q}$ *and* $\mathcal{E}$ *is an extender for* $\mathcal{Q}$*;*

(ES2) $\mathcal{E}$ *is a family of core rays;*

(ES3) the order C_ϵ *agrees with* $\mathcal{E}^-_{e,i}(\leqslant_{\mathcal{F},e})$ *for every* $e \in \partial_\epsilon(S)$*;*

(ES4) the sets $\mathcal{E}^-_{s,i}$ *are intervals of* C_ϵ *on* $\mathcal{E}^-$ *for all* $e \in \partial_\epsilon(S)$ *and* $i \in [k]$*.*

The we can then proceed by induction as before, with the same induction hypotheses. For the most part the proof will follow verbatim, apart from one slight technical issue.

Recall that, in the case where n is even, we use the existence of the family of rays $\overline{\mathcal{C}}$ to find a linkage from $\mathcal{C}$ to $\underline{F}^\downarrow(\mathcal{R}^-)$ which is preserving on $\mathcal{C}$ and similarly, in the case where n is odd, we do the same for $\mathcal{E}_n^-$. In the grid-like case we don't have to be so careful, since every transitional linkage from $\mathcal{C}$ to $F^\downarrow(\mathcal{R}^-)$ will preserve C_ϵ, as long as $|\mathcal{C}|$ is large enough.

However, in order to ensure that $|\mathcal{C}|$ and $|\mathcal{E}_n^-|$ are large enough in each step, we should start by building $N + 3$ inflated copies of $G(T_0^{\neg\epsilon})$ in the first step, which can be done relatively straightforwardly. Indeed, in the case $n = 0$ most of the argument in the construction is unnecessary, since a large part of the construction is constructing a new copy whilst re-routing the the rays $\mathcal{E}_n$ to avoid this new copy, but $\mathcal{E}_0$ is empty. Therefore it is enough to choose a layer $F \in \mathcal{F}_0$ with $|F| \geqslant N+3$, with say $H_1, \dots, H_N \in F$ and to take

$$Q_i^1 =: H(T_k^{\neg\epsilon})$$

for each $i \in [N+3]$ and to take $E_{e,s,i}^1 = H_i^{\downarrow}(R_{e,s})$ for each $e \in \partial_\epsilon(T_0)$, $s \in S(e)$ and $i \in [N+3]$. One can then proceed as before, extending the copies in odd steps and adding a new copy in even steps.

11.10 Outlook: connections with well-quasi-ordering and better-quasi-ordering

Our aim in this section is to sketch what we believe to be the limitations of the techniques of this paper. We will often omit or ignore technical details in order to give a simpler account of the relationship of the ideas involved.

Our strategy for proving ubiquity is heavily reliant on well-quasi-ordering results. The reason is that they are the only known tool for finding extensive tree-decompositions for broad classes of graphs.

To more fully understand this, let's recall how well-quasi-ordering was used in the proofs of Lemmas 11.5.5 and 11.5.10. Lemma 11.5.5 states that any locally finite connected graph with only finitely many ends, all of them thin, has an extensive tree decomposition. The key idea of the proof was as follows: for each end, there is a sequence of separators converging towards that end. The graphs between these separators are finite, and so are well-quasi-ordered by the Graph Minor Theorem. This well-quasi-ordering guarantees the necessary self-similarity.

Lemma 11.5.10, where infinitely many ends are allowed but the graph must have finite tree-width, is similar: once more, for each end there is a sequence of separators converging towards that end. The graphs between these separators are not necessarily finite, but they have bounded tree-width and so they are again well-quasi-ordered.

Note that the Graph Minor Theorem is not needed for this latter result. Instead, the reason it works can be expressed in the following slogan, which will motivate the considerations in the rest of this section:

> Trees of wombats are well-quasi-ordered precisely when wombats themselves are better-quasi-ordered.

Here better-quasi-ordering is a strengthening of well-quasi-ordering introduced by Nash-Williams in [101] essentially in order to make this slogan be true. Since graphs of bounded tree-width can be encoded as trees of graphs of bounded size, what is used here is that graphs of bounded size are better-quasi-ordered.

What if we wanted to go a little further, for example by allowing infinite tree-width but requiring that all ends should be thin? In that case, all we would know about the graphs between the separators would be that all their ends are thin. Such graphs are essentially trees of finite graphs. So, by the slogan above, to show that such trees are well-quasi-ordered we would need the statement that finite graphs are better-quasi-ordered.

Indeed, this problem arises even if we restrict our attention to the following natural common strengthening of Theorems 11.1.1 and 11.1.2:

Conjecture 11.10.1. *Any locally finite connected graph in which all blocks are finite is ubiquitous.*

In order to attack this conjecture with our current techniques we would need better-quasi-ordering of finite graphs.

Thomas has conjectured that countable graphs are well-quasi-ordered with respect to the minor relation. If this were true, it could allow us to resolve problems like those discussed above for countable graphs at least, since all the graphs appearing between the separators are countable. But this approach does not allow us to avoid the issue of better-quasi-ordering of finite graphs. Indeed, since countable trees of finite graphs can be coded as countable graphs, well-quasi-ordering of countable graphs would imply better-quasi-ordering of finite graphs.

Thus until better-quasi-ordering of finite graphs has been established, the best that we can hope for – using our current techniques – is to drop the condition of local finiteness from the main results of this paper, something which we hope to do in the next paper in this series [27].

Chapter 12

Hamilton decompositions of one-ended Cayley graphs

12.1 Introduction

A *Hamiltonian cycle* of a finite graph is a cycle which includes every vertex of the graph. A finite graph $G = (V, E)$ is said to have a *Hamilton decomposition* if its edge set can be partitioned into disjoint sets $E = E_1 \dot\cup E_2 \dot\cup \cdots \dot\cup E_r$ such that each E_i is a Hamiltonian cycle in G.

The starting point for the theory of Hamilton decompositions is an old result by Walecki from 1890 according to which every finite complete graph of odd order has a Hamilton decomposition (see [3] for a description of his construction). Since then, this result has been extended in various different ways, and we refer the reader to the survey of Alspach, Bermond and Sotteau [4] for more information.

Hamiltonicity problems have also been considered for infinite graphs, see for example the survey by Gallian and Witte [129]. While it is sometimes not obvious which objects should be considered the correct generalisations of a Hamiltonian cycle in the setting of infinite graphs, for one-ended graphs the undisputed solution is to consider *double-rays*, i.e. infinite, connected, 2-regular subgraphs. Thus, for us a *Hamiltonian double-ray* is then a double-ray which includes every vertex of the graph, and we say that an infinite graph $G = (V, E)$ has a *Hamilton decomposition* if we can partition its edge set into edge-disjoint Hamiltonian double-rays.

In this paper we will consider infinite variants of two long-standing conjectures on the existence of Hamilton decompositions for finite graphs. The first conjecture concerns Cayley graphs: Given a finitely generated abelian group $(\Gamma, +)$ and a finite generating set S of Γ, the *Cayley graph* $G(\Gamma, S)$ is the multi-graph with vertex set Γ and edge multi-set

$$\{(x, x + g) \,:\, x \in \Gamma, g \in S\}.$$

Conjecture 12.1.1 (Alspach [1, 2])**.** *If Γ is an abelian group and S generates G, then the simplification of $G(\Gamma, S)$ has a Hamilton decomposition, provided that it is $2k$-regular for some k.*

Note that if $S \cap -S = \emptyset$, then $G(\Gamma, S)$ is automatically a $2|S|$-regular simple graph. If $G(\Gamma, S)$ is finite and 2-regular, then the conjecture is trivially true. Bermond, Favaron and Maheo [19] showed that the conjecture holds in the case $k = 2$. Liu [95] proved certain cases of the conjecture for finite 6-regular Cayley graphs, and his result was further extended by Westlund [128]. Liu [96, 97] also gave some sufficient conditions on the generating set S for such a decomposition to exist.

Our main theorem in this paper is the following affirmative result towards the corresponding infinite analogue of Conjecture 12.1.1:

Theorem 12.1.2. *Let Γ be an infinite, finitely generated abelian group, and let S be a generating set such that every element of S has infinite order. If the Cayley graph $G = G(\Gamma, S)$ is one-ended, then it has a Hamilton decomposition.*

We remark that under the assumption that elements of S are non-torsion, the simplification of $G(\Gamma, S)$ is always isomorphic to a Cayley graph $G(\Gamma, S')$ with $S' \subseteq S$ and $S' \cap -S' = \emptyset$, and so our theorem implies the corresponding version of Conjecture 12.1.1 for non-torsion generators, in particular for Cayley graphs of $\mathbb{Z}^n$ with arbitrary generators.

In the case when $G = G(\Gamma, S)$ is two-ended, there are additional technical difficulties when trying to construct a decomposition into Hamiltonian double-rays. In particular, since each Hamiltonian double-ray must meet every finite edge cut an odd number of times, there can be parity reasons why no decomposition exists. One particular two-ended case, namely where $\Gamma \cong \mathbb{Z}$, has been considered by Bryant, Herke, Maenhaut and Webb [30], who showed that when $G(\mathbb{Z}, S)$ is 4-regular, then G has a Hamilton decomposition unless there is an odd cut separating the two ends.

The second conjecture about Hamiltonicity that we consider concerns Cartesian products of graphs: Given two graphs G and H the *Cartesian product* (or product) $G \square H$ is the graph with vertex set $V(G) \times V(H)$ in which two vertices (g, h) and (g', h') are adjacent if and only if either

- $g = g'$ and h is adjacent to h' in H, or
- $h = h'$ and g is adjacent to g' in G.

Kotzig [88] showed that the Cartesian product of two cycles has a Hamilton decomposition, and conjectured that this should be true for the product of three cycles. Bermond extended this conjecture to the following:

Conjecture 12.1.3 (Bermond [18])**.** *If G_1 and G_2 are finite graphs which both have Hamilton decompositions, then so does $G_1 \square G_2$.*

Alspach and Godsil [5] showed that the product of any finite number of cycles has a Hamilton decomposition, and Stong [120] proved certain cases of Conjecture 12.1.3 under additional assumptions on the number of Hamilton cycles in the decomposition of G_1 and G_2 respectively.

Applying techniques we developed to prove Theorem 12.1.2, we show as our second main result of this paper that Conjecture 12.1.3 holds for countably infinite multi-graphs.

Theorem 12.1.4. *If G and H are countable multi-graphs which both have Hamilton decompositions, then so does their product $G \square H$.*

Note that the restriction to countable multi-graphs, i.e multi-graphs with countably many vertices and edges, is necessary. Indeed the existence of a spanning double ray implies that G and H have countable vertex sets. But then if G contains a countable edge cut, then so does $G \square H$. However, if H has uncountably many edges, then any Hamilton decomposition of $G \square H$ must consist of uncountably many edge-disjoint double-rays, contradicting the existence of a countable edge cut.

The paper is structured as follows: In Section 12.2 we mention some group theoretic results and definitions we will need. In Section 12.3 we state our main lemma, the *Covering Lemma*, and show that it implies Theorem 12.1.2. The proof of the Covering Lemma will be the content of Section 12.4. In Section 12.5 we apply our techniques to prove Theorem 12.1.4. Finally, in Section 12.6 we list open problems and possible directions for further work.

12.2 Notation and preliminaries

If $G = (V, E)$ is a graph, and $A, B \subseteq V$, we denote by $E(A, B)$ the set of edges between A and B, i.e. $E(A, B) = \{(x, y) \in E \colon x \in A, y \in B\}$. For $A \subseteq V$ or $F \subseteq E$ we write $G[A]$ and $G[F]$ for the subgraph of G induced by A and F respectively.

For $A, B \subseteq \Gamma$ subsets of an abelian group Γ we write $-A := \{-a \colon a \in A\}$ and $A + B := \{a + b \colon a \in A, b \in B\} \subseteq \Gamma$. If Δ is a subgroup of Γ, and $A \subset \Gamma$ a subset, then $A^{\Delta} = \{a + \Delta \colon a \in A\}$ denotes the family of corresponding cosets. If $g \in \Gamma$ we say that the *order* of g is the smallest $k \in \mathbb{N}$ such that $k \cdot g = 0$. If such a k exists, then g is a *torsion element.* Otherwise, we say the order of g is infinite and g is a *non-torsion* element. For $k \in \mathbb{N}$ we write $[k] = \{1, 2, \ldots, k\}$.

The following terminology will be used throughout.

Definition 12.2.1. *Given a graph G, an edge-colouring $c \colon E(G) \to [s]$ and a colour $i \in [s]$, the* i-subgraph *is the subgraph of G induced by the edge set $c^{-1}(i)$, and the* i-components *are the components of the i-subgraph.*

Definition 12.2.2 (Standard and almost-standard colourings of Cayley graphs)**.** *Let Γ be an infinite abelian group, $S = \{g_1, g_2, \ldots, g_s\}$ a finite generating set for Γ such that every $g_i \in S$ has infinite order, and let G be the Cayley graph $G(\Gamma, S)$.*

- *The* standard colouring *of G is the edge colouring $c_{\text{std}} \colon E(G) \to [s]$ such that $c_{\text{std}}\big((x, x + g_i)\big) = i$ for each $x \in \Gamma, g_i \in S$.*
- *Given a subset $X \subseteq V(G)$ we say that a colouring c is* standard on X *if c agrees with c_{std} on $G[X]$. Similarly if $F \subset E(G)$ we say that c is* standard on F *if c agrees with c_{std} on F.*
- *A colouring $c \colon E(G) \to [s]$ is* almost-standard *if the following are satisfied:*
  - *there is a finite subset $F \subseteq E(G)$ such that c is standard on $E(G) \setminus F$;*
  - *for each $i \in [s]$ the i-subgraph is spanning, and each i-component is a double-ray.*

Definition 12.2.3 (Standard squares and double-rays)**.** *Let Γ and S be as above. Given $x \in \Gamma$ and $g_i \neq g_j \in S$, we call*

$$\blacksquare(x, g_i, g_j) := \{(x, x + g_i), (x, x + g_j), (x + g_i, x + g_i + g_j), (x + g_j, x + g_i + g_j)\}$$

an (i, j)-square with base point x*, and*

$$\leftrightsquigarrow(x, g_i) := \{(x + ng_i, x + (n + 1)g_i) \colon n \in \mathbb{Z}\}$$

an i-double-ray with base point x.

Moreover, given a colouring $c \colon E(G(\Gamma, S)) \to [s]$ we call $\blacksquare(x, g_i, g_j)$ and $\leftrightsquigarrow(x, g_i)$ an (i, j)-standard square *and* i-standard double-ray *if c is standard on $\blacksquare(x, g_i, g_j)$ and $\leftrightsquigarrow(x, g_i)$ respectively.*

Since Γ is an abelian group, every $\blacksquare(x, g_i, g_j)$ is a 4-cycle in $G(\Gamma, S)$ (provided $g_i \neq -g_j$), and since S contains no torsion elements of Γ, $\leftrightsquigarrow(x, g_k)$ really is a double-ray in the Cayley graph $G(\Gamma, S)$.

Let Γ be a finitely generated abelian group. By the Classification Theorem for finitely generated abelian groups (see e.g. [64]), there are integers $n, q_1, \ldots, q_r$ such that $\Gamma \cong \mathbb{Z}^n \oplus \bigoplus_{i=1}^{r} \mathbb{Z}_{q_i}$, where $\mathbb{Z}_q$ is the additive group of the integers modulo q. In particular, for each Γ there is an integer n and a finite abelian group Γ_{fin} such that $\Gamma \cong \mathbb{Z}^n \oplus \Gamma_{\text{fin}}$.

The following structural theorem for the ends of finitely generated abelian groups is well-known:

Theorem 12.2.4. *For a finitely generated group $\Gamma \cong \mathbb{Z}^n \oplus \Gamma_{fin}$, the following are equivalent:*

- $n \geqslant 2$,
- *there exists a finite generating set S such that $G(\Gamma, S)$ is one-ended, and*
- *for all finite generating sets S, the Cayley graph $G(\Gamma, S)$ is one-ended.*

Proof. See e.g. [117, Proposition 5.2] for the fact the number of ends of $G(\Gamma, S)$ is independent of the choice of the generating set S, and [117, Theorem 5.12] for the equivalence with the first item. $\square$

A group Γ satisfying one of the conditions from Theorem 12.2.4 is called *one-ended*.

Corollary 12.2.5. *Let Γ be an abelian group, $S = \{g_1, \dots, g_s\}$ be a finite generating set such that the Cayley graph $G(\Gamma, S)$ is one-ended. Then, for every $g_i \in S$ of infinite order, there is some $g_j \in S$ such that $\langle g_i, g_j \rangle \cong (\mathbb{Z}^2, +)$.*

Proof. Suppose not. It follows that in $\Gamma / \langle g_i \rangle$ every element has finite order, and since it is also finitely generated, it is some finite group Γ_f such that $\Gamma \cong \mathbb{Z} \oplus \Gamma_f$. Thus, by Theorem 12.2.4, G is not one-ended, a contradiction. $\square$

12.3 The covering lemma and a high-level proof of Theorem 12.1.2

Every Cayley graph $G(\Gamma, S)$ comes with a natural edge colouring c_{std}, where we colour an edge $(x, x + g_i)$ with $x \in \Gamma$ and $g_i \in S$ with the index i of the corresponding generating element g_i. If every element of S has infinite order, then every i-subgraph of $G(\Gamma, S)$ consists of a spanning collection of edge-disjoint double-rays, see Definitions 12.2.1 and 12.2.2. So, it is perhaps a natural strategy to try to build a Hamiltonian decomposition by combining each of these monochromatic collections of double-rays into a single monochromatic spanning double-ray.

Rather than trying to do this directly, we shall do it in a series of steps: given any colour $i \in [s]$ where $s = |S|$ and any finite set $X \subset V(G)$, we will show that one can change the standard colouring at finitely many edges, in particular only edges outside of X, so that there is a single double-ray in the colour i which covers X. Moreover, we can ensure that the resulting colouring maintains enough of the structure of the standard colouring that we can repeat this process inductively: it should remain *almost-standard*, i.e. all monochromatic components are still double-rays, see Definition 12.2.2. By taking an appropriate sequence of sets $X_1 \subseteq X_2 \subseteq \cdots$ exhausting the vertex set of G, and varying which colour i we consider, we can ensure that in the limit, each colour class consists of a single spanning double-ray, giving us the desired Hamilton decomposition.

In this section, we formulate our key lemma, namely the Covering Lemma 12.3.1, which allows us to do each of these steps. We will then show how Theorem 12.1.2 follows from the Covering Lemma. The proof of the Covering Lemma is given in Section 12.4.

Lemma 12.3.1 (Covering lemma)**.** *Let Γ be an infinite, one-ended abelian group, $S = \{g_1, g_2, \dots, g_s\}$ a finite generating set such that every $g_i \in S$ has infinite order, and $G = G(\Gamma, S)$ the corresponding Cayley graph.*

Then for every almost-standard colouring c of G, every colour i and every finite subset $X \subseteq V(G)$, there exists an almost-standard colouring $\hat{c}$ of G such that

- $\hat{c} = c$ *on* $E(G[X])$*, and*

- *some* i*-component in* $\hat{c}$ *covers* X*.*

Proof of Theorem 12.1.2 given Lemma 12.3.1. Fix an enumeration $V(G) = \{v_n : n \in \mathbb{N}\}$. Let $X_0 = D'_0 = D'_{-1} = \ldots = D'_{-s+1} = \{v_0\}$ and $c_0 = c_{\text{std}}$. For each $n \geqslant 1$ we will recursively construct almost-standard colourings $c_n : E(G) \to [s]$, finite subsets $X_n \subset V(G)$, (n mod s)-components D_n of c_n and finite paths $D'_n \subseteq D_n$ such that for every $n \in \mathbb{N}$

1. $X_{n-1} \cup \{v_n\} \subseteq X_n$,

2. $V(D'_{n-1}) \subseteq X_n$,

3. $X_n \subseteq V(D'_n)$,

4. D'_n properly extends the path D'_{n-s} (the 'previous' path of colour n mod s) in both endpoints of D'_{n-s}, and

5. c_n agrees with c_{n-1} on $E(G[X_n])$.

Suppose inductively for some $n \in \mathbb{N}$ that c_n, X_n, D_n and D'_n have already been defined. Choose some $X_{n+1} \supseteq X_n \cup \{v_n\}$ large enough such that (1) and (2) are satisfied. Applying Lemma 12.3.1 with input c_n and X_{n+1} provides us with a colouring c_{n+1} such that (5) is satisfied and some ($n+1$ mod s)-component D_{n+1} covers X_{n+1}. Since c_{n+1} is almost-standard, D_{n+1} is a double-ray. Furthermore, since c_{n+1} agrees with c_n on $E(G[X_{n+1}])$, by the inductive hypothesis it agrees with c_k on $E(G[X_{k+1}])$ for each $k \leqslant n$.

Therefore, since $D'_{n+1-s} \subset X_{n-s+2}$ is a path of colour ($n+1$ mod s) in c_{n+1-s}, it follows that $D'_{n+1-s} \subset D_{n+1}$ and so we can extend D'_{n+1-s} to a sufficiently long finite path $D'_{n+1} \subset D_{n+1}$ such that (3) and (4) are satisfied at stage $n+1$.

Once the construction is complete, we define $T_1, \ldots, T_s \subset G$ by

$$T_i = \bigcup_{n \equiv i \mod s} D'_n$$

and claim that they form a decomposition of G into edge-disjoint Hamiltonian double-rays. Indeed, by (4), each T_i is a double-ray. That they are edge-disjoint can be seen as follows: Suppose for a contradiction that $e \in E(T_i) \cap E(T_j)$. Choose $n(i)$ and $n(j)$ minimal such that $e \in E(D'_{n(i)}) \subset E(T_i)$ and $e \in E(D'_{n(j)}) \subset E(T_j)$. We may assume that $n(i) < n(j)$, and so $e \in E(G[X_{n(i)+1}])$ by (2). Furthermore, by (5) it follows that $c_{n(j)}$ agrees with $c_{n(i)}$ on $E(G[X_{n(i)+1}])$. However by construction $c_{n(j)}(e) = j \neq i = c_{n(i)}(e)$ contradicting the previous line.

Finally, to see that each T_i is spanning, consider some $v_n \in V(G)$. By (1), $v_n \in X_n$. Pick $n' \geqslant n$ with $n' \equiv i \mod s$. Then by (3), $D'_{n'} \subset T_i$ covers $X_{n'}$ which in turn contains v_n, as $v_n \in X_n \subseteq X_{n'}$ by (1). □

12.4 Proof of the Covering Lemma

12.4.1 Blanket assumption.

Throughout this section, let us now fix

- a one-ended infinite abelian group Γ with finite generating set $S = \{g_1, \ldots, g_s\}$ such that every element of S has infinite order,

- an almost-standard colouring c of the Cayley graph $G = G(\Gamma, S)$,
- a finite subset $X \subseteq \Gamma$ such that c is standard on $V(G) \setminus X$,
- a colour i, say $i = 1$, and corresponding generator $g_1 \in S$, for which we want to show Lemma 12.3.1, and finally
- a second generator in S, say g_2, such that $\Delta := \langle g_1, g_2 \rangle \cong (\mathbb{Z}^2, +)$, see Corollary 12.2.5.

12.4.2 Overview of proof

We want to show Lemma 12.3.1 for the Cayley graph G, colouring c, generator g_1 and finite set X. The cosets of $\langle g_1, g_2 \rangle$ in Γ cover $V(G)$, and in the standard colouring the edges of colour 1 and 2 form a grid on $\langle g_1, g_2 \rangle$. So, since c is almost-standard, on each of these cosets the edges of colour 1 and 2 will look like a grid, apart from some finite set.

Our aim is to use the structure in these grids to change the colouring c to one satisfying the conclusions of Lemma 12.3.1. It will be more convenient to work with large finite grids, which we require, for technical reasons, to have an even number of rows. This is the reason for the slight asymmetry in the definition below.

Definition 12.4.1. *Let $g_i, g_j \in \Gamma$. For $N, M \in \mathbb{N}$ we write*

$$\langle g_i, g_j \rangle_{N,M} := \{ng_i + mg_j : n, m \in \mathbb{Z},\ -N \leqslant n \leqslant N,\ -M < m \leqslant M\} \subseteq \langle g_i, g_j \rangle \subseteq \Gamma.$$

The structure of our proof can be summarised as follows. First, in Section 12.4.3, we will show that there is some N_0 and some 'nice' finite set P of representatives of cosets of $\langle g_1, g_2 \rangle$ such that $P + \langle g_1, g_2 \rangle_{N_0,N_0}$ covers X. We will then, in Section 12.4.4 pick sufficiently large numbers $N_0 < N_1 < N_2 < N_3$ and consider the grids $P + \langle g_1, g_2 \rangle_{N_3,N_1}$. Using the structure of the grids we will make local changes to the colouring inside $P + (\langle g_1, g_2 \rangle_{N_3,N_1} \setminus \langle g_1, g_2 \rangle_{N_0,N_0})$ to construct our new colouring $\hat{c}$. This new colouring $\hat{c}$ will then agree with c on the subgraph induced by $P + \langle g_1, g_2 \rangle_{N_0,N_0} \supseteq X$, and be standard on $V(G) \setminus \big(P + \langle g_1, g_2 \rangle_{N_3,N_1}\big)$, and hence, as long as we ensure all the colour components are double-rays, almost-standard.

These local changes will happen in three steps. First, in Step 1, we will make local changes inside $x_\ell + (\langle g_1, g_2 \rangle_{N_3,N_1} \setminus \langle g_1, g_2 \rangle_{N_2,N_1})$ for each $x_\ell \in P$, in order to make every i-component meeting $P + \langle g_1, g_2 \rangle_{N_2,N_1}$ a finite cycle.

Next, in Step 2, we will make local changes inside $x_\ell + (\langle g_1, g_2 \rangle_{N_2,N_1} \setminus \langle g_1, g_2 \rangle_{N_1,N_1})$ for each $x_\ell \in P$, in order to combine the cycles meeting this translate of the grid into a single cycle.

Finally, in Step 3, we will make local changes inside $P + (\langle g_1, g_2 \rangle_{N_1,N_1} \setminus \langle g_1, g_2 \rangle_{N_0,N_0})$, in order to join the cycles for different x_ℓ into a single cycle covering $P + \langle g_1, g_2 \rangle_{N_0,N_0}$. We then make one final local change to turn this finite cycle into a double-ray.

12.4.3 Identifying the relevant cosets

Lemma 12.4.2. *There exist $N_0 \in \mathbb{N}$ and a finite set $P = \{x_0, \ldots, x_t\} \subset \Gamma$ such that*

- *$P^\Delta = \{x_0 + \Delta, \ldots, x_t + \Delta\}$ is a path in $G(\Gamma/\Delta, (S \setminus \{g_1, g_2\})^\Delta)$, and*
- *$X \subseteq P + \langle g_1, g_2 \rangle_{N_0,N_0}$.*

Proof. Since X is finite, there is a finite set $Y = \{y_1, \ldots, y_k\} \subset \Gamma$ such that the cosets in $Y^\Delta = \{y_1 + \Delta, \ldots, y_k + \Delta\}$ are all distinct and cover X. Moreover, since every $(y_\ell + \Delta) \cap X$ is finite, there exists $N_0 \in \mathbb{N}$ such that

$$(y_\ell + \langle g_1, g_2 \rangle) \cap X = (y_\ell + \langle g_1, g_2 \rangle_{N_0,N_0}) \cap X$$

for all $1 \leqslant \ell \leqslant k$. Then $X \subseteq Y + \langle g_1, g_2 \rangle_{N_0,N_0}$.

Next, by a result of Nash-Williams [100], every Cayley graph of a countably infinite abelian group has a Hamilton double-ray, and it is a folklore result (see [129]) that every Cayley graph of a finite abelian group has a Hamilton cycle. So in particular, the Cayley graph of $(\Gamma/\Delta, (S \setminus \{g_1, g_2\})^{\Delta})$, has a Hamilton cycle or double-ray, say H. Let $P \supseteq Y$ be a finite set of representatives of the cosets of Δ such that P^{Δ} is the set of vertices of a subpath of H. It is clear that P is as required. $\square$

- For the rest of this section let us fix $N_0 \in \mathbb{N}$ and $P = \{x_0, \ldots, x_t\} \subset \Gamma$ to be as given by Lemma 12.4.2.

12.4.4 Picking sufficiently large grids

In order to choose our grids large enough to be able to make all the necessary changes to our colouring, we will first need the following lemma, which guarantees that we can find, for each $k \neq 1, 2$ and $x \in \Gamma$, many distinct standard k-double-rays which go between the cosets $x + \Delta$ and $(x + g_k) + \Delta$.

Lemma 12.4.3. *For any $g_k \in S \setminus \{g_1, g_2\}$ and any pair of distinct cosets $x + \Delta$ and $(g_k + x) + \Delta$, there are infinitely many distinct standard k-double-rays R for the colouring c with $E(R) \cap E(x + \Delta, (g_k + x) + \Delta) \neq \emptyset$.*

Proof. It clearly suffices to prove the assertion for $c = c_{\text{std}}$. We claim that either

$$\mathcal{R}_1 = \{\leftrightsquigarrow(x + mg_1, g_k) \colon m \in \mathbb{Z}\} \quad \text{or} \quad \mathcal{R}_2 = \{\leftrightsquigarrow(x + mg_2, g_k) \colon m \in \mathbb{Z}\}$$

is such a collection of disjoint standard k-double-rays.

Suppose that $\mathcal{R}_1$ is not a collection of disjoint double-rays. Then there are $m \neq m' \in \mathbb{Z}$ and $n, n' \in \mathbb{Z}$ such that

$$mg_1 + ng_k = m'g_1 + n'g_k.$$

Since g_1 has infinite order, it follows that $n \neq n'$, too, and so we can conclude that there are $\ell, \ell' \in \mathbb{Z} \setminus \{0\}$ such that $\ell g_1 = \ell' g_k$. Similarly, if $\mathcal{R}_2$ is not a collection of disjoint double-rays, then we can find $q, q' \in \mathbb{Z} \setminus \{0\}$ such that $qg_2 = q'g_k$. However, it now follows that

$$q'\ell g_1 = q'(\ell' g_k) = \ell'(q'g_k) = \ell' q g_2,$$

contradicting the fact that $\langle g_1, g_2 \rangle \cong (\mathbb{Z}^2, +)$. This establishes the claim.

Finally, observe that if say $\mathcal{R}_1$ is a disjoint collection, then for every $R_m = \leftrightsquigarrow(x + mg_1, g_k) \in \mathcal{R}_1$ we have $(x + mg_1, x + mg_1 + g_k) \in E(R_m) \cap E(x + \Delta, (g_k + x) + \Delta)$ as desired. $\square$

We are now ready to define our numbers $N_0 < N_1 < N_2 < N_3$. Recall that N_0 and $P = \{x_0, \ldots, x_t\}$ are given by Lemma 12.4.2. For each $\ell \in [t]$, let $g_{n(\ell)}$ be some generator in $S \setminus \{g_1, g_2\}$ that induces the edge between $x_{\ell-1} + \Delta$ and $x_\ell + \Delta$ on the path P^{Δ}. Note that $n(\ell) \in [s] \setminus \{1, 2\}$ for all ℓ.

By Lemma 12.4.3, we may find t^2 many disjoint standard double-rays

$$\mathcal{R} = \left\{ R_\ell^k \colon 1 \leqslant k, \ell \leqslant t \right\}$$

such that for every ℓ, the double-rays in $\left\{ R_\ell^k = \leftrightsquigarrow\left(y_\ell^k, g_{n(\ell)}\right) \colon k \in [t] \right\}$ are standard $n(\ell)$-double-rays containing an edge

$$e_\ell^k = (y_\ell^k, y_\ell^k + g_{n(\ell)}) \in E(R_\ell^k) \cap E(x_{\ell-1} + \Delta, x_\ell + \Delta)$$

so that all $T_\ell^k = \blacksquare\big(y_\ell^k, g_1, g_{n(\ell)}\big)$ are $(1, n(\ell))$-standard squares for c which have empty intersection with $\{x_{\ell-1}, x_\ell\} + \langle g_1, g_2\rangle_{N_0,N_0}$. Furthermore we may assume that these standard squares are all edge-disjoint. Then

- let $N_1 > N_0$ be sufficiently large such that the subgraph induced by $P + \langle g_1, g_2\rangle_{N_1-3,N_1-3}$ contains all standard squares T_ℓ^k mentioned above.
- Let N_2 be arbitrary with $N_2 \geqslant 5N_1$.
- Let N_3 be arbitrary with $N_3 \geqslant N_2 + 2N_1$.

12.4.5 The cap-off step

Our main tool for locally modifying our colouring is the following notion of 'colour switchings', which is also used in [95]. Informally, given a four cycle on which the edge colouring alternates between two colours, to perform a colour switching on this square we exchange the colours of the edges.

Definition 12.4.4 (Colour switching of standard squares)**.** *Given an edge colouring $c\colon E(G(\Gamma, S)) \to [s]$ and an (i,j)-standard square $\blacksquare(x, g_i, g_j)$, a* colour switching *on $\blacksquare(x, g_i, g_j)$ changes the colouring c to the colouring c' such that*

- *$c' = c$ on $E \setminus \blacksquare(x, g_i, g_j)$,*
- *$c'\big((x, x+g_i)\big) = c'\big((x+g_j, x+g_i+g_j)\big) = j$,*
- *$c'\big((x, x+g_j)\big) = c'\big((x+g_i, x+g_i+g_j)\big) = i$.*

It would be convenient if colour switchings maintained the property that a colouring is almost-standard. Indeed, if c is standard on $E(G) \setminus F$ then c' is standard on $E(G) \setminus (F \cup \blacksquare(x, g_i, g_j))$. Also, it is a simple check that if the i and j-subgraphs of G for c are 2-regular and spanning, then the same is true for c'. However, some i or j-components may change from double-rays to finite cycles, and vice versa.

Step 1 (Cap-off step)**.** *There is a colouring c' obtained from c by colour switchings of finitely many $(1,2)$-standard squares such that*

- *$c' = c$ on $E(G[X])$;*
- *every 1-component in c' meeting $P + \langle g_1, g_2\rangle_{N_2,N_1}$ is a finite cycle intersecting both $P + (\langle g_1, g_2\rangle_{N_3,N_1} \setminus \langle g_1, g_2\rangle_{N_2,N_1})$ and $P + \langle g_1, g_2\rangle_{N_1,N_1}$;*
- *every other 1-component, and all other components of all other colour classes of c' are double-rays;*
- *c' is standard outside of $P + \langle g_1, g_2\rangle_{N_3,N_1}$ and inside of $P + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$;*
- *for each $x_\ell \in P$, the set of vertices*

$$\{x_l + ng_1 + mg_2 \colon N_1 \leqslant |n| \leqslant N_2, m \in \{N_1, N_1 - 1\}\}$$

is contained in a single 1-component of c'.

Proof. For $\ell \in [t]$ and $q \in [N_1]$ let $R_q^\ell = \blacksquare\big(v_q^\ell, g_1, g_2\big)$ and $L_q^\ell = \blacksquare\big(w_q^\ell, g_1, g_2\big)$ be the $(1,2)$-squares with base point $v_q^\ell = x_\ell + (N_3 + 1 - 2q) \cdot g_1 + (N_1 + 1 - 2q) \cdot g_2$ and $w_q^\ell = x_\ell - (N_3 + 2 - 2q) \cdot g_1 + (N_1 + 1 - 2q) \cdot g_2$ respectively. The square L_q^ℓ is the mirror image of R_q^ℓ with respect to the y-axis of the grid $x_\ell + \langle g_1, g_2 \rangle$, however the base points are not mirror images, accounting for the slight asymmetry in the definitions.

Since $N_3 \geqslant N_2 + 2N_1$, it follows that

$$R_q^\ell \cup L_q^\ell \subseteq E(x_\ell + (\langle g_1, g_2 \rangle_{N_3,N_1} \setminus \langle g_1, g_2 \rangle_{N_2,N_1}))$$

for all $q \in [N_1]$, and so by assumption on c, all R_q^ℓ and L_q^ℓ are indeed standard $(1,2)$-squares. We perform colour switchings on R_q^ℓ and L_q^ℓ for all $\ell \in [t]$ and $q \in [N_1]$, and call the resulting edge colouring c'. It is clear that $c' = c$ on $E(G[X])$ and that c' is standard outside of $P + \langle g_1, g_2 \rangle_{N_3,N_1}$ and inside of $P + (\langle g_1, g_2 \rangle_{N_2,N_1} \setminus \langle g_1, g_2 \rangle_{N_0,N_0})$. Let $C \subset G$ denote the region consisting of all

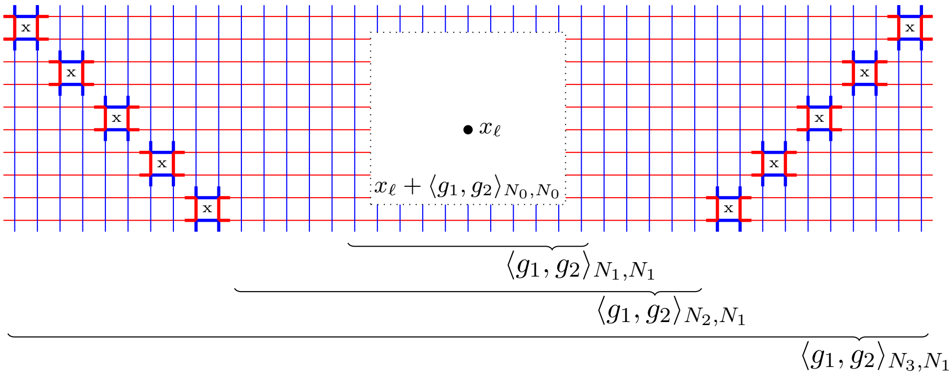

Figure 12.1: Performing colour switchings of standard squares at positions indicated by 'x' in a copy $x_\ell + \langle g_1, g_2 \rangle_{N_3,N_1}$ of a finite grid.

vertices that lie in $x_\ell + (\langle g_1, g_2 \rangle_{N_3,N_1}$ for some ℓ between a pair L_q^ℓ and R_q^ℓ for some q, i.e.

$$C = \bigcup_{\ell=1}^{t} \bigcup_{q=1}^{N_1} \bigcup_{m=1}^{2} \{x_\ell + ng_1 + (N_1 + m - 2q)g_2 \colon |n| \leqslant N_3 + 1 - 2q\}.$$

Then $P + \langle g_1, g_2 \rangle_{N_2,N_1} \subseteq C$. By construction, there are no edges of colour 1 in c' leaving C, that is, $E(C, V(G) \setminus C) \cap c'^{-1}(1) = \emptyset$. In particular, since the 1-subgraph of G under c' remains 2-regular and spanning, as remarked above, all 1-components under c' inside C are finite cycles, whose union covers C.

Also, since each 1-component of c is a double-ray, it must leave the finite set $P + \langle g_1, g_2 \rangle_{N_3,N_1}$ and hence meets some R_q^ℓ or L_q^ℓ. Therefore, by construction each 1-component of c' inside C meets some R_q^ℓ or L_q^ℓ and so, since c' is standard outside of $P + \langle g_1, g_2 \rangle_{N_0,N_0}$ except at the squares R_q^ℓ or L_q^ℓ, each such 1-component meets both $P + (\langle g_1, g_2 \rangle_{N_3,N_1} \setminus \langle g_1, g_2 \rangle_{N_2,N_1})$ and $P + \langle g_1, g_2 \rangle_{N_1,N_1}$.

Moreover, all other colour components remain double-rays. This is clear for all k-components of G if $k \neq 1, 2$ (as the colours switchings of $(1,2)$-standard squares did not affect these other colours). However, it is also clear for the 1-coloured double-rays outside of C and also for all 2-coloured components, as we chose our standard squares R_q^ℓ and L_q^ℓ 'staggered', so as not to create any finite monochromatic cycles, see Figure 12.1 (recall that every $x_\ell + \Delta$ is isomorphic to the grid).

Finally, since $N_1 > N_0$, the edge set

$$\begin{aligned}
&\{(x_\ell + ng_1 + N_1g_2, x_\ell + (n+1)g_1 + N_1g_2)\colon -N_3 \leqslant |n| < N_3 - 1\} \\
&\quad \cup \left\{(v_1^\ell, v_1^\ell + g_2), ((w_1^\ell + g_1, w_1^\ell + g_1 + g_2))\right\} \\
&\quad \cup \{(x_\ell + ng_1 + (N_1-1)g_2, x_\ell + (n+1)g_1 + (N_1-1)g_2)\} : -N_3 \leqslant n < -N_1 \\
&\quad \cup \{(x_\ell + ng_1 + (N_1-1)g_2, x_\ell + (n+1)g_1 + (N_1-1)g_2)\} : N_1 \leqslant n < N_3
\end{aligned}$$

meets only R_1^ℓ and L_1^ℓ and therefore is easily seen to be part of the same 1-component of c'. In Figure 12.1, these edges correspond to the red line at the top, and the two lines below it on either side of $x_\ell + \langle g_1, g_2\rangle_{N_1,N_1}$. □

12.4.6 Combining cycles inside each coset of Δ

In the previous step we chose the $(1,2)$-standard squares at which we performed colour switchings in a staggered manner in the grids $x_l + \langle g_1, g_2\rangle_{N_3,N_1}$, so that we could guarantee that all the 2-components were still double-rays afterwards. In later steps we will no longer be able to be as explicit about which standard squares we perform colour switchings at, and so we will require the following definitions to be able to say when it is 'safe' to perform a colour switching at a standard square.

Definition 12.4.5 (Crossing edges)**.** *Suppose $R = \{(v_i, v_{i+1})\colon i \in \mathbb{Z}\}$ is a double-ray and $e_1 = (v_{j_1}, v_{j_2})$ and $e_2 = (v_{k_1}, v_{k_2})$ are edges with $j_1 < j_2$ and $k_1 < k_2$. We say that e_1 and e_2* cross *on R if either $j_1 < k_1 < j_2 < k_2$ or $k_1 < j_1 < k_2 < j_2$.*

Lemma 12.4.6. *For an edge-colouring $c\colon E(G(\Gamma, S)) \to [s]$, suppose that $\blacksquare(x, g_i, g_k)$ is an (i,k)-standard square with $g_i \neq -g_k$, and further that the two k-coloured edges $(x, x+g_k)$ and $(x+g_i, x+g_i+g_k)$ of $\blacksquare(x, g_i, g_k)$ lie on the same* standard k-double-ray *$R = \leftrightsquigarrow(x, g_k)$. Then the two i-coloured edges of $\blacksquare(x, g_i, g_k)$ cross on R.*

Proof. Write $e_1 = (x, x+g_i)$ and $e_2 = (x+g_k, x+g_k+g_i)$ for the two i-coloured edges of $\blacksquare(x, g_i, g_k)$. The assumption that $(x, x+g_k)$ and $(x+g_i, x+g_i+g_k)$ both lie on $\leftrightsquigarrow(x, g_k)$ implies that $g_i = rg_k$ for some $r \in \mathbb{Z} \setminus \{-1, 0, 1\}$. If $r > 1$, we have $x < x+g_k < x+g_i < x+g_k+g_i$ (where $<$ denotes the natural linear order on the vertex set of the double-ray), and if $r < -1$, we have $x+g_i < x+g_k+g_i < x < x+g_k$, and so the edges e_1 and e_2 indeed cross on R. □

Definition 12.4.7 (Safe standard square)**.** *Given an edge colouring $c\colon E(G(\Gamma, S)) \to [s]$ we say an (i,k)-standard square $T = \blacksquare(x, g_i, g_k)$ is* safe *if $g_i \neq -g_k$ and either*

- *the k-components for c meeting T are distinct double-rays, or*
- *there is a unique k-component for c meeting T, which is a double-ray on which $(x, x+g_i)$ and $(x+g_k, x+g_i+g_k)$ cross.*

The following lemma tells us, amongst other things, that if we perform a colour switching at a safe $(1,k)$-standard square then the k-components in the resulting colouring meeting that square will still be double-rays.

Lemma 12.4.8. *Let $c\colon E(G(\Gamma, S)) \to [s]$ be an edge colouring, $T = \blacksquare(x, g_i, g_k)$ be an (i,k)-standard square with $g_i \neq -g_k$, and c' be the colouring obtained by performing a colour switching on T. Then the following statements are true:*

- *If there are two distinct i-components C_1 and C_2 for c meeting T which are both 2-regular, at least one of which is a finite cycle, then there is a single i-component for c' meeting T which is 2-regular and whose vertex set is $V(C_1) \cup V(C_2)$;*

- *If the k-components for c meeting T are distinct double-rays then the k-components for c' meeting T are distinct double-rays;*

- *If there is a unique k-component for c meeting T, which is a double-ray on which $(x, x+g_i)$ and $(x+g_k, x+g_i+g_k)$ cross, then there is unique k-component for c' meeting T, which is a double-ray.*

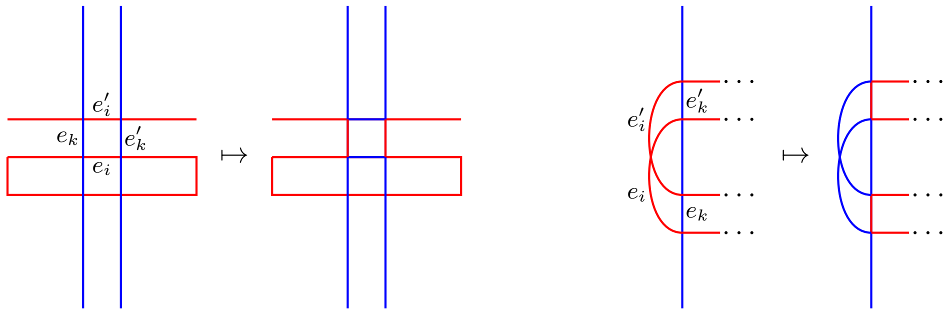

Figure 12.2: The two situations in Lemma 12.4.8 with i in red and k in blue.

Proof. Let us write $e_i = (x, x+g_i)$, $e_k = (x, x+g_k)$, $e_i' = (x+g_k, x+g_i+g_k)$ and $e_k' = (x+g_i, x+g_i+g_k)$, so that $\blacksquare(x, g_i, g_j) = \{e_i, e_k, e_i', e_k'\}$.

For the first item, let the i-components for c be $e_i \in C_1$ and $e_i' \in C_2$, where without loss of generality C_2 is a finite cycle. Then $C_2 - e_i'$ is a finite path, and $C_1 - e_i$ has at most 2 components, one containing x and one containing $x+g_i$. Hence, the i-component for c' meeting T, $(C_1 \cup C_2) - \{e_i, e_i'\} + \{e_k, e_k'\}$, is connected and has vertex set $V(C_1) \cup V(C_2)$.

For the second item, let the k-components for c be $e_k \in D_1$ and $e_k' \in D_2$. Then $D_1 - e_k$ has two components, a ray starting at x and a ray starting at $x+g_k$. Similarly, $D_2 - e_k'$ has two components, a ray starting at $x+g_i$ and a ray starting at $x+g_i+g_k$. Hence, the k-components for c' meeting T, which are the components of $(D_1 \cup D_2) - \{e_k, e_k'\} + \{e_i, e_i'\}$, are distinct double-rays.

Finally, if there is a single k-component D for c meeting T such that D is a double-ray, then $D - \{e_k, e_k'\}$ consist of three components. Since e_i and e_i' cross on D there are two cases as to what these components are. Either the components consist of two rays, starting at x and $x+g_i+g_k$ and a finite path from $x+g_k$ to $x+g_i$, or the components consist of two rays, starting at $x+g_i$ and $x+g_k$, and a finite path from $x+g_i+g_k$ to x. In either case, the k-component for c' meeting T, namely $D - \{e_k, e_k'\} + \{e_i, e_i'\}$, is a double-ray. $\square$

Lemma 12.4.8 is also useful as the first item allows us to use $(1,k)$ colour switchings to combine two 1-components into a single 1-component which covers the same vertex set.

Step 2 (Combining cycles step)**.** *We can change c' from Step 1 via colour switchings of finitely many $(1,2)$-standard squares to a colouring c'' satisfying*

- *$c'' = c' = c$ on $E(G[X])$;*

- *every 1-component in c'' meeting $P + \langle g_1, g_2\rangle_{N_2,N_1}$ is a finite cycle intersecting both $P + (\langle g_1, g_2\rangle_{N_3,N_1} \setminus \langle g_1, g_2\rangle_{N_2,N_1})$ and $P + \langle g_1, g_2\rangle_{N_1,N_1}$;*

- *every other* 1*-component, and all other components of all other colour classes of* c'' *are double-rays;*
- *every* 1*-component in* c'' *meeting some* $x_k + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ *covers* $x_k + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$*;*
- c'' *is standard outside of* $P + \langle g_1, g_2\rangle_{N_3,N_1}$ *and inside of* $P + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$*.*

Proof. Our plan will be to go through the 'grids' $x_k + \langle g_1, g_2\rangle_{N_2,N_1}$ in order, from $k = 0$ to t, and use colour switchings to combine all the 1-components which meet $x_k + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ into a single 1-component. We note that, since c' is not standard on X, it may be the case that these 1-components also meet $x_{k'} + \langle g_1, g_2\rangle_{N_2,N_1}$ for $k' \neq k$.

We claim inductively that there exists a sequence of colourings $c' = c_0, c_1, \ldots, c_t = c''$ such that for each $0 \leqslant \ell \leqslant t$:

- $c_\ell = c' = c$ on $E(G[X])$;
- every 1-component in c_ℓ meeting $P + \langle g_1, g_2\rangle_{N_2,N_1}$ is a finite cycle intersecting both $P + (\langle g_1, g_2\rangle_{N_3,N_1} \setminus \langle g_1, g_2\rangle_{N_2,N_1})$ and $P + \langle g_1, g_2\rangle_{N_1,N_1}$;
- for every $k \leqslant \ell$, every 1-component in c_ℓ meeting $x_k + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ covers $x_k + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$;
- for every $k > \ell$, $c_\ell = c'$ on $x_k + \langle g_1, g_2\rangle_{N_2,N_1}$
- every other 1-component, and all other components of all other colour classes of c_ℓ are double-rays;
- c_ℓ is standard outside of $P + \langle g_1, g_2\rangle_{N_3,N_1}$ and inside of $P + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$.

In Step 1 we constructed $c_0 = c'$ such that this holds. Suppose that $0 < \ell \leqslant t$, and that we have already constructed c_k for $k < \ell$.

For $q \in [4N_1 - 2]$ we define $T_q = \blacksquare(v_q, g_1, g_2)$ to be the $(1, 2)$-square with base point

$$v_q = \begin{cases} x_\ell + (N_2 + 2 - 2q)g_1 + (N_1 - q)g_2 & \text{if } q \leqslant 2N_1 - 1, \text{ and} \\ x_\ell - (N_2 + 3 - 2q')g_1 + (N_1 - q')g_2 & \text{if } q' = q - (2N_1 - 1) \geqslant 1. \end{cases}$$

With these definitions, T_{2N_1-1+q} is the mirror image of T_q for all $q \in [2N_1 - 1]$ along the y-axis. Moreover, since $N_2 \geqslant 5N_1$, each T_q is contained within $x_k + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_1,N_1})$, see Figure 12.3.

We will combine the 1-components in $c_{\ell-1}$ which meet $x_\ell + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ into a single component by performing colour switchings at some of the $(1, 2)$-squares T_q. Let us show first that most of the induction hypotheses are maintained regardless of the subset of the T_q we make switchings at.

We note that, since $c_{\ell-1}$ is standard inside of $x_\ell + (\langle g_1, g_1\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ and outside of $P + \langle g_1, g_2\rangle_{N_3,N_1}$, and $g_1 \neq -g_2$, each T_q is a safe $(1, 2)$-standard square for $c_{\ell-1}$. Furthermore, by construction, even if we perform colour switchings at any subset of the T_q, the remaining squares remain standard and safe.

Hence, by Lemma 12.4.8 and the induction assumption, after performing colour switchings at any subset of the standard squares T_q all 2-components of the resulting colouring will be double-rays. Secondly, these colour switchings will not change the colouring outside of $P + \langle g_1, g_2\rangle_{N_2,N_1}$ and inside of $P + \langle g_1, g_2\rangle_{N_1,N_1}$, or in any $x_k + \langle g_1, g_2\rangle_{N_2,N_1}$ with $k \neq \ell$. In particular, every 1-component not meeting $P + \langle g_1, g_2\rangle_{N_2,N_1}$ will still be a double-ray. Finally, again by Lemma

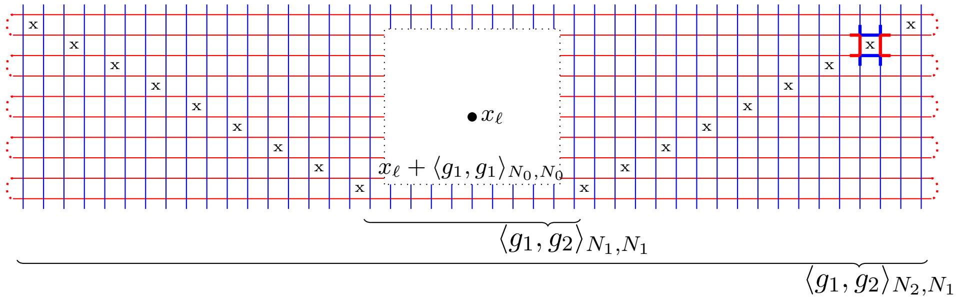

Figure 12.3: The standard squares T_q, with a colour switching performed at T_2.

12.4.8, every 1-component of the resulting colouring meeting $P + \langle g_1, g_2\rangle_{N_2,N_1}$ will be a finite cycle which covers the vertex set of some union of 1-components in $c_{\ell-1}$, and hence will intersect both $P + (\langle g_1, g_2\rangle_{N_3,N_1} \setminus \langle g_1, g_2\rangle_{N_2,N_1})$ and $P + \langle g_1, g_2\rangle_{N_1,N_1}$.

Let us write $e_q = (v_q, v_q + g_1)$ for each $q \in [4N_1 - 2]$. Since $c_{\ell-1} = c'$ on $x_\ell + \langle g_1, g_2\rangle_{N_2,N_1}$, and by Step 1 c' is standard on $x_\ell + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$, each 1-component of $c_{\ell-1}$ that meets $x_\ell + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ contains at least one e_q. Also, e_1 and e_{2N_1} belong to the same 1-component by the last claim in Step 1. Let us write $\mathcal{C}$ for the collection of such cycles, and consider the map

$$\alpha\colon \mathcal{C} \to \{1, \dots, 4N_1 - 1\},\ C \mapsto \min\{q\colon e_q \in E(C)\},$$

which maps each cycle to the first e_q that it contains. Since $\mathcal{C}$ is a disjoint collection of cycles, the map α is injective. Now let c_ℓ be the colouring obtained from $c_{\ell-1}$ by switching all standard squares in

$$\mathcal{T} = \{T_q\colon q \in \operatorname{im}(\alpha)\} \setminus \{T_1\}.$$

We claim that c_ℓ satisfies our induction hypothesis for ℓ. By the previous comments it will be sufficient to show

Claim 12.4.9. *Every 1-component in c_ℓ meeting $x_\ell + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ covers $x_\ell + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$.*

To see this, we index $\mathcal{C} = \{C_1, \dots, C_r\}$ such that $u < v$ implies $\alpha(C_u) < \alpha(C_v)$, and consider the sequence of colourings $\{c^z\colon z \in [r]\}$ where $c^1 = c_{\ell-1}$ and each c^z is obtained from c^{z-1} by switching the standard square $T_{\alpha(C_z)}$.

Let us show by induction that for every $z \in [r]$ there is a 1-component of c^z which covers $\bigcup_{y \leqslant z} C_y$. For $z = 1$ the claim is clearly true. So, suppose $z > 1$. Since $\alpha(C_z)$ is minimal in $\{\alpha(C_y)\colon y \geqslant z\}$ it follows that $e_q \in \bigcup_{y<z} C_y$ for every $q < \alpha(C_z)$. Note that, since $c_{\ell-1} = c'$ on $x_\ell + \langle g_1, g_2\rangle_{N_2,N_1}$, it follows from the final claim in the Cap-off step that C_1 contains both e_1 and e_{2N_1}, and so $\alpha(C_z) \neq 2N_1$.

Consider the standard square $T_{\alpha(C_z)}$. Since $c_{\ell-1} = c'$ on $x_\ell + \langle g_1, g_2\rangle_{N_2,N_1}$, by construction the edge 'opposite' to $e_{\alpha(C_z)}$ in $T_{\alpha(C_z)}$, that is, $e_{\alpha(C_z)} + g_j$, is in the same 1-component in $c_{\ell-1}$ as $e_{\alpha(C_z)-1}$, and hence is contained in $\bigcup_{y<z} C_y$.

Therefore, by Lemma 12.4.8, after performing an $(1, 2)$-colour switching at $T_{\alpha(C_z)}$, the 1-component in c^z contains $\bigcup_{y \leqslant z} C_y$.

Hence, there is a 1-component of $c_\ell = c^r$ which covers $\bigcup_{y \leqslant r} C_y$, and so there is a unique 1-component of c_ℓ meeting $x_\ell + (\langle g_1, g_2\rangle_{N_2,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ which covers it, establishing the claim. $\square$

12.4.7 Combining cycles across different cosets of Δ

In the third and final step we join the finite cycles covering each $x_\ell + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ into a single finite cycle, and then make one final switch to absorb this cycle into a double-ray. The resulting colouring will then satisfy the conditions of Lemma 12.3.1.

Step 3 (Combining cosets step)**.** *We can change c'' from the previous lemma to an almost-standard colouring $\hat{c}$ such that*

- $\hat{c} = c'' = c' = c$ *on* $E(G[X])$*;*
- *Some component in colour* 1 *covers* $P + \langle g_1, g_2\rangle_{N_1,N_1}$.

Proof. Recall that $P = \{x_0, \dots, x_t\}$ is such that $P^\Delta = \{x_0 + \Delta, \dots, x_t + \Delta\}$ is a finite, graph-theoretic path in the Cayley graph of the quotient Γ/Δ with generating set $(S \setminus \{g_1, g_2\})^\Delta$. Moreover, recall from Section 12.4.4 that $N_1 > N_0$ was chosen so that for the initial colouring c there were t^2 many disjoint standard double-rays

$$\mathcal{R} = \left\{R_\ell^k \colon 1 \leqslant k, \ell \leqslant t\right\}$$

such that for every ℓ, the double-rays in $\left\{R_\ell^k = \leftrightsquigarrow\!\left(y_\ell^k, g_{n(\ell)}\right) \colon k \in [t]\right\}$ are standard $n(\ell)$-double-rays containing an edge

$$e_\ell^k = (y_\ell^k, y_\ell^k + g_{n(\ell)}) \in E(R_\ell^k) \cap E(x_{\ell-1} + \Delta, x_\ell + \Delta)$$

so that all $T_\ell^k = \blacksquare\!\left(y_\ell^k, g_1, g_{n(\ell)}\right)$ are edge-disjoint $(1, n(\ell))$-standard squares for the colouring c contained in the subgraph induced by $P + \langle g_1, g_2\rangle_{N_1-3,N_1-3}$ which have empty intersection with $\{x_{\ell-1}, x_\ell\} + \langle g_1, g_2\rangle_{N_0,N_0}$. However, since we only altered the $(1,2)$-subgraphs of G in Step 1 and 2, it is clear that all these standard double-rays and standard squares for c remain standard also for the colourings c' and in particular c''.

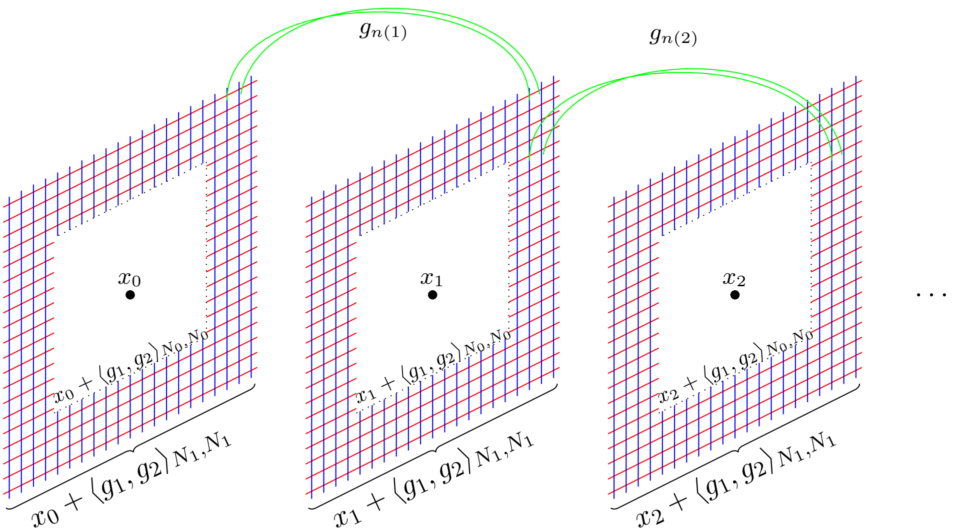

Figure 12.4: Using $(1, n(\ell))$-standard squares to join up different cosets. For this picture, we assume wlog that $x_{\ell+1} = x_\ell + g_{n(\ell+1)}$.

We claim that there exists a function $k \colon [t] \to [t] \cup \{\perp\}$ such that iteratively switching $T_\ell^{k(\ell)}$ (or not doing anything at all if $k(\ell) = \perp$) results in a sequence of colourings $c'' = c_0, c_1, \dots, c_t$ such that for each $0 \leqslant \ell \leqslant t$,

1. a single finite 1-component in c_ℓ covers $\{x_0, \ldots, x_\ell\} + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$,
2. for every k, every 1-component in c_ℓ meeting $x_k + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$ is a finite cycle covering $x_k + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$, and
3. every other 1-component, and all other components of all other colour classes in c_ℓ are double-rays.

In Step 2 we constructed a colouring $c_0 = c''$ for which properties (1)–(3) are satisfied. Now suppose that $\ell \geqslant 1$, and that the colouring $c_{\ell-1}$ obtained by switching the standard squares $\left\{T_{\ell'}^{k(\ell')} : \ell' \in [\ell-1], k(\ell') \neq \perp\right\}$ satisfies (1)–(3). By construction, each such standard square $T_{\ell'}^{k(\ell')}$ is has an edge in common with the ray $R_{\ell'}^{k(\ell')}$ and potentially one further $n(\ell')$-component. But since we had reserved more that $\ell - 1$ different rays $R_\ell^1, \ldots, R_\ell^t$, it follows that some ray R_ℓ^K remains a standard $n(\ell)$-coloured component for $c_{\ell-1}$.

Both edges $(y_\ell^K, y_\ell^K + g_i)$ and $(y_\ell^K + g_{n(\ell)}, y_\ell^K + g_{n(\ell)} + g_i)$ of T_ℓ^K are contained in $\{x_{\ell-1}, x_\ell\} + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$, and hence are, by assumption (2), covered by finite 1-cycles in $c_{\ell-1}$. If both edges lie in the same finite 1-cycle, there is nothing to do and we set $k(\ell) := \perp$, so that $c_\ell = c_{\ell-1}$. However, if they lie on different finite cycles, set $k(\ell) := K$. Then, in our procedure we perform a colour switching on the standard square $T_\ell^{k(\ell)}$ and claim that the resulting c_ℓ is as required. By Lemma 12.4.8, the two finite 1-components merge into a single finite cycle, and so (1) and (2) are certainly satisfied for c_ℓ.

To see (3), we need to verify that $T_\ell^{k(\ell)}$ is, when we perform the switching, safe. However, $T_\ell^{k(\ell)}$ was chosen so that the edge $(y_\ell^{k(\ell)}, y_\ell^{k(\ell)} + g_{n(\ell)}) \in T_\ell^{k(\ell)}$ lies on a standard double-ray $R = R_\ell^{k(\ell)}$ of $c_{\ell-1}$. Also, by the inductive assumption (3), the second $n(\ell)$-coloured edge $(y_\ell^{k(\ell)} + g_i, y_\ell^{k(\ell)} + g_i + g_{n(\ell)}) \in T_\ell^{k(\ell)}$ lies on an $n(l)$-coloured double-ray R' in $c_{\ell-1}$. If R and R' are distinct, then $T_\ell^{k(\ell)}$ is safe, and if $R = R'$ then, since R is a standard $n(\ell)$-double-ray, Lemma 12.4.6 implies that $T_\ell^{k(\ell)}$ is safe. Hence c_ℓ satisfies (3). This completes the induction step.

Thus, by (1) and (3), we obtain an edge-colouring c_t for G such that a single finite 1-component covers $P + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$, and all other 1-components and all other components of other colour classes in c_t are double-rays. Furthermore, since every 1-component which meets $P + \langle g_1, g_2\rangle_{N_0,N_0}$ must meet $P + (\langle g_1, g_2\rangle_{N_1,N_1} \setminus \langle g_1, g_2\rangle_{N_0,N_0})$, it follows that the 1-component in fact covers $P + \langle g_1, g_2\rangle_{N_0,N_0}$. Moreover, since $T_\ell^{k(\ell)} \subset P + \langle g_1, g_2\rangle_{N_1-3,N_1-3}$ for all $\ell \in [t]$, it follows that c_t is standard on $x_0 + (\langle g_1, g_2\rangle_{N_1,\infty} \setminus \langle g_1, g_2\rangle_{N_1-3,N_1-3})$, and that it is standard outside of $P + \langle g_1, g_2\rangle_{N_3,N_1}$. Hence, the square $\blacksquare(x, g_1, g_2)$ with base point $x = x_0 + (N_1 - 2)g_1 + N_1 g_2$ is a standard $(1,2)$-square such that

- the edge $(x, x + g_1)$ lies on the finite 1-cycle of c_t,
- the edge $(x+g_2, x+g_2+g_1)$ lies on a standard 1-double-ray $\leftrightsquigarrow(x + g_2, g_1)$ (lying completely outside of $P + \langle g_1, g_2\rangle_{N_3,N_1}$) of c_t, and
- the edges $(x, x+g_2)$ and $(x+g_1, x+g_2+g_1)$ lie on distinct standard 2-double-rays $\leftrightsquigarrow(x, g_2)$ and $\leftrightsquigarrow(x + g_1, g_2) \subseteq x_0 + (\langle g_1, g_2\rangle_{N_1,\infty} \setminus \langle g_1, g_2\rangle_{N_1-3,N_1-3})$.

Therefore, we may perform a colour switching on $\blacksquare(x, g_1, g_2)$, which results, by Lemma 12.4.8, in an almost-standard colouring of G such that a single 1-component covers $P + \langle g_1, g_2\rangle_{N_1,N_1}$, and hence X. □

12.5 Hamiltonian decompositions of products

The techniques from the previous section can also be applied to give us the following general result about Hamiltonian decompositions of products of graphs.

Theorem 12.1.4. *If G and H are countable multi-graphs which both have Hamilton decompositions, then so does their product $G\Box H$.*

Proof. Suppose that $\{R_i : i \in I\}$ and $\{S_j : j \in J\}$ form decompositions of G and H into edge-disjoint Hamiltonian double-rays, where I, J may be finite or countably infinite. Note that, for each $i \in I, j \in J$, $R_i\Box S_j$ is a spanning subgraph of $G\Box H$, and is isomorphic to the Cayley graph of $(\mathbb{Z}^2, +)$ with the standard generating set.

Let $\pi_G : G\Box H \to G$ and $\pi_H : G\Box H \to H$ the projection maps from $G\Box H$ onto the respective coordinates. As our *standard colouring* for $G\Box H$ we take the map

$$c\colon E(G\Box H) \to I\dot{\cup}J,\ e \mapsto \begin{cases} i & \text{if } e \in \pi_G^{-1}(E(R_i)), \\ j & \text{if } e \in \pi_H^{-1}(E(S_j)). \end{cases}$$

Then each $R_i\Box S_j$ is 2-coloured (with colours i and j), and this colouring agrees with the standard colouring of $C_{\mathbb{Z}^2} = G((\mathbb{Z}^2, +), \{(1, 0), (0, 1)\})$ from Section 12.3. We also define an *almost-standard colouring* of $G\Box H$ as in Definition 12.2.2.

We may suppose that $V(G) = \mathbb{N} = V(H)$. Fix a surjection $f : \mathbb{N} \to I \cup J$ such that every colour appears infinitely often.

By starting with $c_0 = c$ and applying Lemma 12.3.1 recursively inside the spanning subgraphs $R_{f(k)}\Box S_1$, if $f(k) \in I$, or inside $R_1\Box S_{f(k)}$, for $f(k) \in J$, we find a sequence of almost-standard edge-colourings $c_k : G\Box H \to I \cup J$ and natural numbers $M_k \leqslant N_k < M_{k+1}$ such that

- c_{k+1} agrees with c_k on the subgraph of $G\Box H$ induced by $[0, M_{k+1}]^2$,
- there is a finite path D_k of colour $f(k)$ in c_k covering $[0, N_k]^2$, and
- M_{k+1} is large enough such that $D_k \subset [0, M_{k+1}]^2$.

To be precise, suppose we already have a finite path D_k of colour $f(k)$ in c_k covering $[0, N_k]^2$, and at stage $k + 1$ we have say $f(k + 1) \in I$, and so we are considering $R_{f(k+1)}\Box S_1 \cong C_{\mathbb{Z}^2}$. We choose

- $M_{k+1} > N_k$ large enough such that $D_k \subset [0, M_{k+1}]^2 \subset G\Box H$, and
- $N_{k+1} > M_{k+1}$ large enough such that $Q_1 = [0, N_{k+1}]^2 \subset G\Box H$ contains all edges where c_k differs from the standard colouring c.

Next, consider an isomorphism $h : R_{f(k+1)}\Box S_1 \cong C_{\mathbb{Z}^2}$. Pick a 'square' $Q_2 \subset R_{f(k+1)}\Box S_1$ with $Q_1 \subset Q_2$, i.e. a set Q_2 such that h restricted to Q_2 is an isomorphism to the subgraph of $C_{\mathbb{Z}^2}$ induced by $[-\tilde{N}_{k+1}, \tilde{N}_{k+1}]^2 \subseteq \mathbb{Z}^2$ for some $\tilde{N}_{k+1} \in \mathbb{N}$, and then apply Lemma 12.3.1 to $R_{f(k+1)}\Box S_1$ and Q_2 to obtain a finite path D_{k+1} of colour $f(k + 1)$ in c_{k+1} covering Q_2.

It follows that the double-rays $\{T_i : i \in I\} \cup \{T_j : j \in J\}$ with $T_\ell = \bigcup_{k\in f^{-1}(\ell)} D_k$ give the desired decomposition of $G\Box H$. $\Box$

12.6 Open Problems

As mentioned in Section 12.2, the finitely generated abelian groups can be classified as the groups $\mathbb{Z}^n \oplus \bigoplus_{i=1}^r \mathbb{Z}_{q_i}$, where $n, r, q_1, \ldots, q_r \in \mathbb{Z}$. Theorem 12.1.2 shows that Alspach's conjecture holds for every such group with $n \geqslant 2$, as long as each generator has infinite order. The question however remains as to what can be said about Cayley graphs $G(\Gamma, S)$ when S contains elements of finite order.

Problem 9. *Let* Γ *be an infinite, finitely-generated, one-ended abelian group and* S *be a generating set for* Γ *which contains elements of finite order. Show that* $G(\Gamma, S)$ *has a Hamilton decomposition.*

Alspach's conjecture has also been shown to hold when $n = 1$, $r = 0$, and the generating set S has size 2, by Bryant, Herke, Maenhaut and Webb [30]. In a paper in preparation [60], the first two authors consider the general case when $n = 1$ and the underlying Cayley graph is 4-regular. Since the Cayley graph is 2-ended, it can happen for parity reasons that no Hamilton decomposition exists. However, this is the only obstruction, and in all other cases the Cayley graphs have a Hamilton decomposition. Together with the result of Bermond, Favaron and Maheo [19] for finite abelian groups, and the case $\Gamma \cong (\mathbb{Z}^2, +)$ of Theorem 12.1.2, this fully characterises the 4-regular connected Cayley graphs of finite abelian groups which have Hamilton decompositions. A natural next step would be to consider the case of 6-regular Cayley graphs.

Problem 10. *Let* Γ *be a finitely generated abelian group and let* S *be a generating set of* Γ *such that* $C(\Gamma, S)$ *is* 6*-regular. Characterise the pairs* (Γ, S) *such that* $G(\Gamma, S)$ *has a decomposition into spanning double-rays.*

Bibliography

[1] B. Alspach. Research problems. *Discrete Mathematics*, 50:115, 1984.

[2] B. Alspach. Unsolved problem 4.5. *Annals of Discrete Mathematics*, 27:464, 1985.

[3] B. Alspach. The wonderful walecki construction. *Bull. Inst. Combin. Appl*, 52:7–20, 2008.

[4] B. Alspach, J. C. Bermond, and D. Sotteau. Decomposition into cycles I: Hamilton decompositions. In *Cycles and rays*, pages 9–18. Springer, 1990.

[5] B. Alspach and Christopher D. Godsil. *Cycles in graphs*, volume 27. Elsevier, 1985.

[6] O. Amini, F. Mazoit, N. Nisse, and S. Thomassé. Submodular partition functions. *Discrete Mathematics*, 309(20):6000–6008, 2009.

[7] T. Andreae. Bemerkung zu einem Problem aus der Theorie der unendlichen Graphen. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 46, page 91. Springer, 1977.

[8] T. Andreae. Über eine Eigenschaft lokalfiniter, unendlicher Bäume. *Journal of Combinatorial Theory, Series B*, 27(2):202–215, 1979.

[9] T. Andreae. On a problem of R. Halin concerning infinite graphs II. *Discrete Mathematics*, 29(3):219–227, 1980.

[10] T. Andreae. On the Reconstruction of Locally Finite Trees. *Journal of Graph Theory*, 5(2):123–135, 1981.

[11] T. Andreae. Simpler counterexamples to the edge-reconstruction conjecture for infinite graphs. *Journal of Combinatorial Theory, Series B*, 32(3):258 – 263, 1982.

[12] T. Andreae. On reconstructing rooted trees. *Journal of Combinatorial Theory, Series B*, 62(2):183 – 198, 1994.

[13] T. Andreae. On disjoint configurations in infinite graphs. *Journal of Graph Theory*, 39(4):222–229, 2002.

[14] T. Andreae. Classes of locally finite ubiquitous graphs. *Journal of Combinatorial Theory, Series B*, 103(2):274–290, 2013.

[15] J. Azzato. Linked tree-decompositions of represented infinite matroids. *Journal of Combinatorial Theory, Series B*, 101(3):123–140, 2011.

[16] J. Barát. Directed path-width and monotonicity in digraph searching. *Graphs and Combinatorics*, 22(2):161–172, 2006.

[17] P. Bellenbaum and R. Diestel. Two short proofs concerning tree-decompositions. *Combinatorics, Probability & Computing*, 11(06):541–547, 2002.

[18] J. C. Bermond. Hamiltonian decompositions of graphs, directed graphs and hypergraphs. *Annals of Discrete Mathematics*, 3:21–28, 1978.

[19] J. C. Bermond, O. Favaron, and M. Maheo. Hamiltonian decomposition of cayley graphs of degree 4. *Journal of Combinatorial Theory, Series B*, 46(2):142–153, 1989.

[20] D. Berwanger, A. Dawar, P. Hunter, S. Kreutzer, and J. Obdržálek. The DAG-width of directed graphs. *Journal of Combinatorial Theory, Series B*, 102(4):900–923, 2012.

[21] D. Bienstock, N. Robertson, P. Seymour, and R. Thomas. Quickly excluding a forest. *Journal of Combinatorial Theory, Series B*, 52(2):274–283, 1991.

[22] J. A. Bondy and R. Hemminger. Reconstructing infinite graphs. *Pacific Journal of Mathematics*, 52(2):331–340, 1974.

[23] J. A. Bondy and R. L. Hemminger. Graph reconstruction – a survey. *Journal of Graph Theory*, 1(3):227–268, 1977.

[24] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs I : Topological ubiquity of trees. in preparation.

[25] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs II : Ubiquity of graphs with non-linear end structure. in preparation.

[26] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs III : Ubiquity of graphs with extensive tree decompositions. in preparation.

[27] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs IV : Ubiquity of graphs of bounded tree-width. in preparation.

[28] N. Bowler, J. Erde, P. Heinig, F. Lehner, and M. Pitz. A counterexample to the reconstruction conjecture for locally finite trees. *Bull. Lond. Math. Soc.*, 49(4):630–648, 2017.

[29] N. Bowler, J. Erde, P. Heinig, F. Lehner, and M. Pitz. Non-reconstructible locally finite graphs. *Journal of Combinatorial Theory, Series B*, 133:122–152, 2018.

[30] D. Bryant, S. Herke, B. Maenhaut, and B. Webb. On hamilton decompositions of infinite circulant graphs. *arXiv preprint arXiv:1701.08506*, 2017.

[31] D. Bryant and V. Moulton. Neighbor-net: an agglomerative method for the construction of phylogenetic networks. *Mol. Biol. Evol.*, 21(2):255–265, 2004.

[32] G. Călinescu, A. Dumitrescu, and J. Pach. Reconfigurations in graphs and grids. *SIAM Journal on Discrete Mathematics*, 22(1):124–138, 2008.

[33] J. Carmesin, R. Diestel, and F. Hamann, M.and Hundertmark. Canonical tree-decompositions of finite graphs II. Essential parts. *Journal of Combinatorial Theory, Series B*, 118:268–283, 2016.

[34] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. k-Blocks: a connectivity invariant for graphs. *SIAM Journal on Discrete Mathematics*, 28(4):1876–1891, 2014.

[35] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs I. Existence and algorithms. *Journal of Combinatorial Theory, Series B*, 116:1–24, 2016.

[36] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34(1):11–46, 2014.

[37] J. Carmesin and J. P. Gollin. Canonical tree-decompositions of a graph that display its k-blocks. *Journal of Combinatorial Theory, Series B*, 122:1–20, 2017.

[38] M. Chudnovsky and P. Seymour. A well-quasi-order for tournaments. *Journal of Combinatorial Theory, Series B*, 101(1):47–53, 2011.

[39] R. Diestel. Tangles in the empirical sciences: indirect clustering by separation. *In preparation.*

[40] R. Diestel. Tangles in the social sciences: a new mathematical model to identify types and predict behaviour. *In preparation.*

[41] R. Diestel. Graph minors I: A short proof of the path-width theorem. *Combinatorics, Probability and Computing*, 4(1):27–30, 1995.

[42] R. Diestel. Locally finite graphs with ends: A topological approach, i. basic theory. *Discrete Mathematics*, 311(15):1423–1447, 2011.

[43] R. Diestel. *Graph Theory (5th edition).* Springer-Verlag, 2017.

[44] R. Diestel. Abstract separation systems. *Order*, 35(1):157–170, 2018.

[45] R. Diestel. Tree sets. *Order*, 35(1):171–192, 2018.

[46] R. Diestel, P. Eberenz, and J. Erde. Duality theorems for blocks and tangles in graphs. *SIAM Journal on Discrete Mathematics*, 31(3):1514–1528, 2017.

[47] R. Diestel, J. Erde, and D. Weißauer. Structural submodularity and tangles in abstract separation systems. *Journal of Combinatorial Theory, Series A*, 167:155–180, 2019.

[48] R. Diestel, F. Hundertmark, and S. Lemanczyk. Profiles of separations: in graphs, matroids, and beyond. *Combinatorica*, pages 1–39, 2016.

[49] R. Diestel and D. Kühn. Graph-theoretical versus topological ends of graphs. *Journal of Combinatorial Theory, Series B*, 87:197–206, 2003.

[50] R. Diestel and S. Oum. Unifying duality theorems for width parameters in graphs and matroids. II. General duality. *arXiv preprint arXiv:1406.3798*, 2014.

[51] R. Diestel and S. Oum. Tangle-tree duality in abstract separation systems. *arXiv preprint arXiv:1701.02509*, 2017.

[52] R. Diestel and S. Oum. Tangle-tree duality: in graphs, matroids and beyond. *arXiv preprint arXiv:1701.02651*, 2017.

[53] R. Diestel and R. Thomas. Excluding a countable clique. *Journal of Combinatorial Theory, Series B*, 76(1):41–67, 1999.

[54] R. Diestel and G. Whittle. Tangles and the Mona Lisa. *arXiv preprint arXiv:1603.06652*, 2016.

[55] G. A. Dirac. On rigid circuit graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 25(1-2):71–76, 1961.

[56] Z. Dvořák. A stronger structure theorem for excluded topological minors. *arXiv preprint arXiv:1209.0129*, 2012.

[57] C. Elbracht, J. Kneip, and M. Teegen. Trees of tangles in abstract separation systems. *In preparation.*

[58] J. Erde. Refining a tree-decomposition which distinguishes tangles. *SIAM Journal on Discrete Mathematics*, 31(3):1529–1551, 2017.

[59] J. Erde. A unified treatment of linked and lean tree-decompositions. *Journal of Combinatorial Theory, Series B*, 130:114–143, 2018.

[60] J. Erde and F. Lehner. Hamilton decompositions of infinite 4-regular Cayley graphs. In preparation.

[61] J. Erde, F. Lehner, and M. Pitz. Hamilton decompositions of one-ended Cayley graphs. *Journal of Combinatorial Theory, Series B*, 2019. to appear.

[62] J. Fisher. A counterexample to the countable version of a conjecture of Ulam. *Journal of Combinatorial Theory*, 7(4):364–365, 1969.

[63] J. Fisher, R. L. Graham, and F. Harary. A simpler counterexample to the reconstruction conjecture for denumerable graphs. *Journal of Combinatorial Theory, Series B*, 12(2):203–204, 1972.

[64] L. Fuchs. *Abelian groups.* Springer, 2015.

[65] F. Gavril. The intersection graphs of subtrees in trees are exactly the chordal graphs. *Journal of Combinatorial Theory. Series B*, 16:47–56, 1974.

[66] J. Geelen and B. Joeris. A generalization of the grid theorem. *arXiv preprint arXiv:1609.09098*, 2016.

[67] J.F. Geelen, A.M.H Gerards, N. Roberston, and G. Whittle. Obstructions to branch-decomposition of matroids. *Journal of Combinatorial Theory, Series B*, 96:560–570, 2006.

[68] J.F. Geelen, A.M.H Gerards, and G. Whittle. Branch-width and well-quasi-ordering in matroids and graphs. *Journal of Combinatorial Theory, Series B*, 84(2):270–290, 2002.

[69] P. Gollin and K. Heuer. k-connected sets in infinite graphs. in preparation.

[70] M. Grohe and D. Marx. Structure theorem and isomorphism test for graphs with excluded topological subgraphs. *SIAM Journal on Computing*, 44(1):114–159, 2015.

[71] R. Halin. Über die Maximalzahl fremder unendlicher Wege in Graphen. *Mathematische Nachrichten*, 30(1-2):63–85, 1965.

[72] R. Halin. Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen. *Mathematische Nachrichten*, 44(1-6):119–127, 1970.

[73] R Halin. A problem in infinite graph-theory. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 43, pages 79–84. Springer, 1975.

[74] R. Halin. S-functions for graphs. *Journal of geometry*, 8(1-2):171–186, 1976.

[75] F. Harary. On the reconstruction of a graph from a collection of subgraphs. In *Theory of Graphs and its Applications (Proc. Sympos. Smolenice, 1963)*, pages 47–52, 1964.

[76] F. Harary, A. J. Schwenk, and R. L. Scott. On the reconstruction of countable forests. *Publ. Inst. Math.*, 13(27):39–42, 1972.

[77] D. Harvey and D. Wood. Parameters tied to treewidth. *Journal of Graph Theory*, 84(4):364–385, 2017.

[78] P. Hliněný and G. Whittle. Matroid tree-width. *European Journal of Combinatorics*, 27(7):1117–1128, 2006.

[79] P. Hliněný and G. Whittle. Addendum to matroid tree-width. *European Journal of Combinatorics*, 30(4):1036–1044, 2009.

[80] F. Hundertmark. Profiles. An algebraic approach to combinatorial connectivity. arXiv:1110.6207, 2011.

[81] P. Hunter and S. Kreutzer. Digraph measures: Kelly decompositions, games, and orderings. *Theoretical Computer Science*, 399(3):206–219, 2008.

[82] T. Johnson, N. Robertson, P. Seymour, and R. Thomas. Addendum to "Directed treewidth". *unpublished manuscript*, 2001.

[83] T. Johnson, N. Robertson, P. Seymour, and R. Thomas. Directed tree-width. *Journal of Combinatorial Theory, Series B*, 82(1):138–154, 2001.

[84] P.J. Kelly. A congruence theorem for trees. *Pacific J. Math.*, 7(1):961–968, 1957.

[85] I. Kim and P. Seymour. Tournament minors. *Journal of Combinatorial Theory, Series B*, 112:138–153, 2015.

[86] S. Kintali. Directed minors III. directed linked decompositions. *arXiv preprint arXiv:1404.5976*, 2014.

[87] D. Kornhauser, G. Miller, and P. Spirakis. Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. In *Proceedings of the 25th Annual Symposium on Foundations of Computer Science, 1984*, SFCS '84, pages 241–250. IEEE Computer Society, 1984.

[88] A. Kotzig. Every cartesian product of two circuits is decomposable into two hamiltonian circuits. *Centre de Recherche Math ematiques, Montreal*, 1973.

[89] I. Kříž and R. Thomas. The menger-like property of the tree-width of infinite graphs. *Journal of Combinatorial Theory, Series B*, 52(1):86–91, 1991.

[90] J. Lagergren. Upper bounds on the size of obstructions and intertwines. *Journal of Combinatorial Theory, Series B*, 73(1):7–40, 1998.

[91] J. Lake. A problem concerning infinite graphs. *Discrete Mathematics*, 14(4):343–345, 1976.

[92] R. Laver. On Fraïssé's order type conjecture. *Annals of Mathematics*, pages 89–111, 1971.

[93] R. Laver. Better-quasi-orderings and a class of trees. *Studies in foundations and combinatorics*, 1:31–48, 1978.

[94] C. H. Liu and R. Thomas. Excluding subdivisions of bounded degree graphs. *Journal of Combinatorial Theory, Series B*, 134:1–35, 2019.

[95] J. Liu. Hamiltonian decompositions of cayley graphs on abelian groups. *Discrete Mathematics*, 131(1-3):163–171, 1994.

[96] J. Liu. Hamiltonian decompositions of cayley graphs on abelian groups of odd order. *Journal of Combinatorial Theory, Series B*, 66(1):75–86, 1996.

[97] J. Liu. Hamiltonian decompositions of cayley graphs on abelian groups of even order. *Journal of Combinatorial Theory, Series B*, 88(2):305–321, 2003.

[98] W. Mader. Kreuzungsfreie a, b–Wege in endlichen Graphen. *Abh. Math. Sem. Univ. Hamburg*, 42:187–204, 1974.

[99] W. Mader. Über n-fach zusammenhängende Eckenmengen in Graphen. *Journal of Combinatorial Theory, Series B*, 25(1):74–93, 1978.

[100] C. St. J. A. Nash-Williams. Abelian groups, graphs and generalized knights. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 55, pages 232–238. Cambridge Univ Press, 1959.

[101] C. St. J. A. Nash-Williams. On well-quasi-ordering infinite trees. *Mathematical Proceedings of the Cambridge Philosophical Society*, 61(3):697–720, 1965.

[102] C. St. J. A. Nash-Williams. Reconstruction of locally finite connected graphs with at least three infinite wings. *Journal of Graph Theory*, 11(4):497–505, 1987.

[103] C. St. J. A. Nash-Williams. Reconstruction of infinite graphs. *Discrete Mathematics*, 95(1):221–229, 1991.

[104] C. St. J. A. Nash-Williams. Reconstruction of locally finite connected graphs with two infinite wings. *Discrete Mathematics*, 92(1):227–249, 1991.

[105] J. Obdržálek. DAG-width: connectivity measure for directed graphs. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 814–821. Society for Industrial and Applied Mathematics, 2006.

[106] S. Oum. Rank-width and well-quasi-ordering. *SIAM Journal on Discrete Mathematics*, 22(2):666–682, 2008.

[107] B. Reed. Introducing directed tree width. *Electronic Notes in Discrete Mathematics*, 3:222–229, 1999.

[108] N. Robertson and P. Seymour. Graph minors. IV. tree-width and well-quasi-ordering. *Journal of Combinatorial Theory, Series B*, 48(2):227–254, 1990.

[109] N. Robertson and P. D. Seymour. Graph minors. III. Planar tree-width. *Journal of Combinatorial Theory, Series B*, 36(1):49–64, 1984.

[110] N. Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52(2):153–190, 1991.

[111] N. Robertson and P. D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.

[112] N. Robertson and P. D. Seymour. Graph minors XXIII. Nash-Williams' immersion conjecture. *J. Comb. Theory, Ser. B*, 100(2):181–205, 2010.

[113] N. Robertson and P.D. Seymour. Graph minors. V. excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.

[114] N. Robertson and P.D. Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63:65–110, 1995.

[115] P. D. Robertson, N.and Seymour. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004.

[116] M. Safari. D-width: A more natural measure for directed tree width. In *International Symposium on Mathematical Foundations of Computer Science*, pages 745–756. Springer, 2005.

[117] P. Scott and T. Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36, pages 137–203, 1979.

[118] P. D. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B*, 58(1):22–33, 1993.

[119] S. Shelah. A compactness theorem for singular cardinals, free algebras, whitehead problem and tranversals. *Israel Journal of Mathematics*, 21(4):319–349, 1975.

[120] R. Stong. Hamilton decompositions of cartesian products of graphs. *Discrete Mathematics*, 90(2):169–190, 1991.

[121] R.E. Tarjan. Decomposition by clique separators. *Discrete mathematics*, 55(2):221–232, 1985.

[122] R. Thomas. A counter-example to 'Wagner's conjecture' for infinite graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 103, pages 55–57. Cambridge University Press, 1988.

[123] R. Thomas. Well-quasi-ordering infinite graphs with forbidden finite planar minor. *Transactions of the American Mathematical Society*, 312(1):279–313, 1989.

[124] R. Thomas. A menger-like property of tree-width: the finite case. *Journal of Combinatorial Theory, Series B*, 48(1):67–76, 1990.

[125] C. Thomassen. Reconstructing 1-coherent locally finite trees. *Commentarii Mathematici Helvetici*, 53(1):608–612, 1978.

[126] W. T. Tutte. *Graph Theory.* Addison-Wesley, 1984.

[127] D. Weißauer. On the block number of graphs. *arXiv preprint arXiv:1702.04245*, 2017.

[128] E. Westlund. Hamilton decompositions of certain 6-regular cayley graphs on abelian groups with a cyclic subgroup of index two. *Discrete Mathematics*, 312(22):3228–3235, 2012.

[129] D. Witte and J. Gallian. A survey: Hamiltonian cycles in cayley graphs. *Discrete Mathematics*, 51(3):293–304, 1984.

[130] D. R. Woodall. A note on a problem of Halin's. *Journal of Combinatorial Theory, Series B*, 21(2):132–134, 1976.