On tree-decompositions, tangles and coarse graph theory

Dissertation

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

Universität Hamburg

vorgelegt
am Fachbereich Mathematik
von

Sandra Albrechtsen

Hamburg, 2025

Betreuer

Prof. Dr. Reinhard Diestel (Universität Hamburg)

Gutachter

Prof. Dr. Reinhard Diestel (Universität Hamburg)
Prof. Dr. Daniel Král' (Universität Leipzig)
Prof. Dr. Alexander Scott (University of Oxford)

Prüfungskommission

Vorsitz	:	Prof. Dr. Armin Iske	(Universität Hamburg)
Stellvertretender Vorsitz	:	Prof. Dr. Nathan Bowler	(Universität Hamburg)
Mitglied	:	Prof. Dr. Reinhard Diestel	(Universität Hamburg)
Mitglied	:	Dr. Matthias Hamann	(Universität Hamburg)
Mitglied und Schriftführer	:	PD Dr. Ralf Holtkamp	(Universität Hamburg)

Datum der Disputation: 17. September 2025

Contents

1 Introduction

Throughout the thesis, graphs will be simple and may be infinite, unless otherwise stated.

In a series of twenty papers spanning over 500 pages from 1983 to 2004, Robertson and Seymour [110] proved the Graph Minors Theorem: *in every infinite sequence of finite graphs there is one that is a minor of a later one.* As a consequence, every family of finite graphs that is closed under taking minors can be defined by a finite set of forbidden minors.

The Graph Minors Theorem has made a fundamental impact not only in graph theory, and its proof has introduced new methods and notions that are of quite independent interest. In this thesis we deal with two of these objects: *tangles* and *tree-decompositions*. Furthermore, we explore a branch of graph theory that only emerged quite recently, called *coarse graph theory*, and which is closely connected to *graph-decompositions*, a natural extension of tree-decompositions.

This thesis consists of three parts. The first two parts consider tangles and tree-decompositions – first in infinite and then in finite graphs. The third part deals with questions from coarse graph theory.

In what follows we give a brief overview of the parts and their chapters, using the terminology from [41]. We remark that, additionally, every chapter will feature its own independent and more comprehensive introduction.

1.1 Parts I and II: Tangles and tree-decompositions of finite and infinite graphs

1.1.1 Background

Tree-decompositions are a central object in structural graph theory. They were not only a crucial tool in Robertson and Seymour's [110] proof of the Graph Minor Theorem, but also attracted attention as several computationally hard problems can be solved efficiently on graphs that admit a tree-decomposition into small parts.

A *tree-decomposition* of a graph G is a pair $(T, \mathcal{V})$ of a tree T and a family $\mathcal{V} = (V_t)_{t\in V(T)}$ of subsets V_t of $V(G)$, its *bags*, that together cover G and which are arranged roughly along T. Since the bags of $(T, \mathcal{V})$ are arranged along T, the overall structure of G resembles that of the tree T. Because of this, tree-decompositions permit to transfer many properties of trees to more general graphs. However, this is only true really if the bags are not too large. Therefore, the *width* of a tree-decomposition is defined as the supremum of the sizes of its bags minus 1, and a graph has *tree-width* at most $k \in \mathbb{N}$ if it admits a tree-decomposition of width at most k.

With the notion of 'tree-width', the problem arose to characterise the graphs that have small tree-width. There are a number of substructures, e.g. large grid or clique minors, or k-blocks for large k, that are known to force a graph G to have large tree-width. While these objects differ in their concrete shape, they have one thing in common: for every low-order separations $\{A, B\}$ of G, they lie mostly on one of its sides A or B. Robertson and Seymour [114] distilled from this the notion of a tangle: formally, a k-*tangle* τ, for some $k \in \mathbb{N}$, in a graph G is an orientation of the separations $\{A, B\}$ of G of *order* $|A \cap B| < k$, as (A, B) or (B, A), such that there do not exist three separations $(A_i, B_i) \in \tau$ such that the union of the sides $G[A_i]$ not chosen by the tangle covers G.

Since their first introduction, the notion of a tangle and its framework of graph separations have been generalized to 'abstract separation systems' [42] and have found a wide range of applications outside of graph theory (for a summary, see [44]). In this thesis we will focus on graph tangles only.

1.1.2 Chapters 3 and 4: Linked and lean tree-decompositions of infinite graphs

A rooted[1] tree-decomposition $(T, \mathcal{V})$ of a graph G is *linked* if for every two comparable nodes $s < t$ of T there are $\min\{|V_e| \colon e \in E(sTt)\}$ pairwise disjoint V_s–V_t paths in G. It is *lean* if for every two (not necessarily distinct) nodes $s, t \in T$ and vertex sets $Z_s \subseteq V_s$ and $Z_t \subseteq V_t$ with $|Z_s| = |Z_t| =: \ell \in \mathbb{N}$, either G contains ℓ pairwise disjoint Z_s–Z_t paths or there exists an edge $e \in E(sTt)$ with $|V_e| < \ell$. Note that 'lean' is a stronger property than 'linked'.

Kříž and Thomas [95, 120] proved the following theorem about lean tree-decompositions, which forms a cornerstone both in Robertson and Seymour's work [113] on well-quasi-ordering finite graphs, and in Thomas's result [119] that the class of infinite graphs of tree-width $< k$ is well-quasi-ordered under the minor relation for all $k \in \mathbb{N}$.

Theorem 1.1.1. *Every graph of tree-width $< k \in \mathbb{N}$ admits a lean tree-decomposition of width $< k$.*

In Chapter 4 we discuss a number of counterexamples demonstrating the limits of possible generalisations of Theorem 1.1.1 to arbitrary infinite tree-width. In particular, we construct a locally finite, planar, connected graph that has no lean tree-decomposition (Example 8).

On the positive side, we obtain in Chapter 3 what might be the deepest and most comprehensive result in this thesis: a version of Theorem 1.1.1 for graphs of (possibly unbounded) finite tree-width:[2]

Theorem 1. *Every graph of finite tree-width admits a rooted tree-decomposition into finite parts that is linked, tight, and componental.*

[1]A tree-decomposition $(T, \mathcal{V})$ is *rooted* if its decomposition tree T has a *root*.
[2]A graph has *finite tree-width* if it admits a tree-decomposition into finite parts.

Note that the properties 'tight' and 'componental' ensure that the adhesion sets of the tree-decomposition are no larger than necessary (see Section 3.2.1 for the definition). We show in Chapter 4 that Theorem 1 is essentially best possible; in particular, 'linked' cannot be strengthened to its 'unrooted' version that requires the linked property between every pair of nodes and not just comparable ones (see Example 9).

While achieving just a subset of the properties 'linked', 'tight' and 'componental' may be significantly easier, having all three properties simultaneously is much more challenging to achieve. We substantiate the usefulness of Theorem 1 by presenting several applications demonstrating the surprisingly powerful interplay of the properties of being linked, tight, and componental (see Section 3.1). For example, we show that every graph without half-grid minor has a lean tree-decomposition into finite parts (Theorem 3).

1.1.3 Chapter 5: Tangle-tree duality in infinite graphs

Theorem 1 has another surprising application: It allows us to generalise the 'Tangle-tree duality theorem' to infinite graphs. This theorem is one of the two major results in Robertson and Seymour's [114] original work on tangles (rephrased here in the terminology of [41]):

Tangle-tree duality theorem. *For every finite graph G and $k \in \mathbb{N}$, exactly one of the following assertions holds:*

(i) *There exists a k-tangle in G.*

(ii) *There exists an $S_k(G)$-tree over $\mathcal{T}^*$.*

Note that $S_k(G)$-trees over $\mathcal{T}^*$ can be seen as a certain kind of tree-decomposition (see Section 2.3 for details). The Tangle-tree duality theorem implies an approximate duality for tangles and tree-width: every graph G with a k-tangle has tree-width at least $k-1$, while a tree as in (ii) induces a tree-decomposition of G of width at most $3k-4$.

While there are several papers that extend results about tangles in finite graphs to infinite ones [31, 64, 65, 90], no generalisation of the Tangle-tree duality theorem had been found so far. As the perhaps most gratifying result about tangles in this thesis, we address this shortcoming and extend Robertson and Seymour's Tangle-tree duality theorem to infinite graphs:

Theorem 13. *For every graph G and $k \in \mathbb{N}$, exactly one of the following assertions holds:*

(i) *There exists a principal k-tangle in G that is not induced by an end of combined degree $< k$.*

(ii) *There exists a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$.*

We discuss in Section 5.1 why Theorem 13 is the 'correct' generalisation of the Tangle-tree duality theorem to infinite graphs, even though it is not verbatim. In particular, Theorem 13 contains the finite Tangle-tree duality theorem as a special case. Moreover, Theorem 13 yields the same duality between tangles and tree-width as the Tangle-tree duality theorem.

The proof of Theorem 13 builds on the tree-decomposition from Theorem 1 and the 'refining techniques' developed in [4, 66]. Very broadly, the proof proceeds as follows: One can show that if a graph G contains no k-tangle as in (i), then G has finite tree-width. Hence, we can apply Theorem 1 to G, which yields a tree-decomposition $(T, \mathcal{V})$ of G. We then *refine* $(T, \mathcal{V})$, i.e. we decompose its parts by further tree-decompositions, so that the combined overall tree-decomposition of G induces a tree as in (ii).

1.1.4 Chapters 6 and 7: Refining trees of tangles

While the Tangle-tree duality theorem yields a tree-decomposition of a graph G if G has no k-tangles for some $k \in \mathbb{N}$, the other major tangle theorem from Robertson and Seymour [114] yields a tree-decomposition of G if G has some k-tangles.

Tree-of-tangles theorem. *For every $k \in \mathbb{N}$, every finite graph G has a tree-decomposition that distinguishes all the k-tangles in G efficiently.*

In such a tree-decomposition $(T, \mathcal{V})$ every k-tangle 'lives' in a different part. Such a tree-decomposition provides information about the overall structure of the graph and the location of the tangles inside it. However, a tangle-distinguishing tree-decompositions given by Tree-of-tangles theorem will in general not decompose G in an optimal way, in that its bags might be much larger than necessary. In Chapter 6 we show that, for every $k \in \mathbb{N}$, every tree-decomposition of a finite graph G which efficiently distinguishes all its k-tangles can be refined, i.e. further decomposed, to a tree-decomposition which is optimal in the sense that its bags are either too small to be home to a k-tangle, or as small as possible while being home to a k-tangle (see Theorem 17).

Strengthening the Tree-of-tangles theorem, Carmesin and Gollin [30] proved that, for every $k \in \mathbb{N}$, every finite graph has a canonical tree-decomposition $(T, \mathcal{V})$ of adhesion less than k that efficiently distinguishes every two distinct k-profiles,[3] and which has the further property that every separable k-block is equal to the unique bag of $(T, \mathcal{V})$ in which it is contained. Here, a *k-block* in a graph G, for some $k \in \mathbb{N}$, is a maximal set of at least k vertices no two of which can be separated in G by removing fewer than k other vertices. In Chapter 7 we give a new short proof of this result by showing that such a tree-decomposition can in fact be obtained from any canonical tight tree-decomposition of adhesion less than k by refining (see Theorem 18). As an application, we also obtain a generalisation of Carmesin and Gollin's result to locally finite graphs.

[3] *Profiles* are a generalisation of tangles (see Section 2.2 for the definition).

1.1.5 Chapter 8: On vertex sets inducing tangles

The notion of a tangle is designed so that every highly connected substructure of a graph G induces a tangle by *majority vote*: the collection of all low-order separations (A, B) of G whose side B contains most of the substructure will be a tangle. A natural question that was asked by Diestel, Hundertmark and Lemanczyk [47] is whether there is some sort of a converse: Is every tangle induced by majority vote of some set of vertices? More formally, does there exists for every k-tangle τ in a finite graph G some set $U \subseteq V(G)$ such that for every separation $(A, B) \in \tau$ we have $|U \cap A| < |U \cap B|$?

We reduce this question to graphs whose size is bounded by a function in k (see Corollary 22). For this, we show that any given k-tangle in a finite graph G is the lift of a k-tangle in some topological minor of G whose size is bounded in k (see Theorem 24). As an application, we strengthen a result by Elbracht, Kneip and Teegen [62]. They showed that for every $k \in \mathbb{N}$ every k-tangle τ in a finite graph G is induced by a weight function $f : V(G) \to \mathbb{N}$, in that $f(A) < f(B)$ for all $(A, B) \in \tau$. We now obtain that there is such a function f whose total weight is bounded in k.

1.2 Part III: Coarse graph theory

1.2.1 Background

Roughly speaking, tree-decompositions make it possible to describe the global structure of a graph: The smaller the tree-width of a graph, the more 'tree-like' its global structure. However, this only works well if the global structure is actually tree-like. Indeed, there are a number of graphs whose global structure is not tree-like, and tree-decompositions fail to capture the global structure of such graphs.

This motivated the definition of *graph-decompositions* [48], a natural extension of tree-decompositions which allow the bags V_h of decompositions $(H, \mathcal{V})$ to be arranged along general decomposition graphs H instead of just trees (see Section 9.3 for the definition). Recent applications of graph-decompositions include a local-global decomposition theorem [48] as well as the study of local separations [32] and of locally chordal graphs [1].

To measure how 'well' a graph-decomposition approximates a graph, we were looking for a notion of 'width', similar as to tree-decompositions. But as it turned out, defining the 'width' of a graph-decomposition analogously to tree-width, that is, as the minimal cardinality of a bag of the decomposition (minus 1), does not yield a meaningful width measure for arbitrary graph-decompositions.[4]

[4]See Section 11.1.2 for details. In particular, every graph admits a decomposition $(H, \mathcal{V})$ of width 2 (minus 1)

This inspired us to consider a metric perspective instead: To define the 'width' of a graph-decomposition $(H, \mathcal{V})$, we measure the size of its bags not in terms of their cardinality, but by the radius of its *parts* $G[V_h]$. More precisely, the *radial width* of a decomposition is the largest radius among its parts, and the *radial $\mathcal{H}$-width* of a graph G for a given class $\mathcal{H}$ of graphs is the smallest radial width among all decompositions $(H, \mathcal{V})$ of G with $H \in \mathcal{H}$.

With a suitable width measure at hand, our next goal was to identify obstructions to small radial $\mathcal{H}$-width and to characterise which graphs have small radial $\mathcal{H}$-width for given classes $\mathcal{H}$ of graphs. For this, we considered metric versions of minors as candidates for suitable obstructions, and conjectured that they appear in every graph of large radial width.

At the same time, but independently of our considerations on graph decompositions, and instead inspired by Gromov's [80] coarse geometry viewpoint, Georgakopoulos and Papasoglu [75] had been developing a new perspective on the global structure of graphs. At its core, this perspective revolves around the concept of *quasi-isometry*, a generalisation of bi-Lipschitz maps which allows for an additive error. Roughly speaking, two metric spaces are quasi-isometric whenever their large scale geometry coincides (see Section 9.2 for the definition). Georgakopoulos and Papasoglu [75] presented results and questions regarding the interplay of geometry and graphs. At the heart of their paper, they proposed a conjecture [75, Conjecture 1.1], here stated in amended form:

Fat minor conjecture. *Let $\mathcal{X}$ be a finite set of finite graphs. Then there exists a function $f : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that, for all $K \in \mathbb{N}$, every graph with no K-fat X minor for any $X \in \mathcal{X}$ is $f(K)$-quasi-isometric to a graph with no X minor for any $X \in \mathcal{X}$.*

A graph X is a *K-fat minor* of a graph G for some $K \in \mathbb{N}$ if there is a model of X in G whose branch sets and branch paths are pairwise at least K far apart, except for incident vertex/edge pairs in X.

To our surprise, Georgakopoulos and Papasoglu had been considering the same kind of metric minors that we looked at, and their Fat minor conjecture turned out to be equivalent to our indented conjecture on graph-decompositions and small radial width that we mentioned earlier. Indeed, in Section 11.2 we show that two graphs are quasi-isometric if and only if each has a decomposition modelled on the other, subject to bounds on width parameters that correspond to the constants of the quasi-isometry.

The Fat minor conjecture is known to be true in some small cases (see [37, 70, 75] and Chapters 11 and 12), but has been disproved by Davies, Hickingbotham, Illingworth and McCarty for general sets $\mathcal{X}$ [38, Theorem 1].

---

along a planar graph H.

1.2.2 Chapter 10: A Menger-type theorem for two induced paths

In order to prove the Fat minor conjecture for any class $\mathcal{X}$ of graphs, it is inevitable to construct fat minors of the graphs in $\mathcal{X}$. As the branch paths of fat minors have to be pairwise far apart, we aimed to find a characterisation of when a graph contains many paths that are pairwise far apart. For this, we introduced the *Coarse Menger Conjecture* (Conjecture 28), of which we managed to prove the case $k = 2$: There exists a function $f(d) \in O(d)$, such that for every graph G and $X, Y \subseteq V(G)$, either there exist two X–Y paths P_1 and P_2 such that $d_G(P_1, P_2) \geqslant d$, or there exists a vertex $v \in V(G)$ such that the ball of radius $f(d)$ around v intersects every X–Y path (Theorem 27).[5]

Quite surprisingly, what we had developed as a tool to prove the K_4-case of the Fat minor conjecture, immediately attracted attention [72, 76, 88, 103]. In particular, the Coarse Menger Conjecture has since been disproved in its strongest form [103]. However, the case $k = 2$, which we proved, was all that we really needed to prove the case $\mathcal{X} = \{K_4\}$ of the Fat minor conjecture (see next section).

1.2.3 Chapters 11 and 12: Special cases of the Fat minor conjecture

In Chapter 12 we prove what, up to today, remains the case of the Fat minor conjecture with the most challenging proof amongst those that are known: the case $\mathcal{X} = \{K_4\}$ (see Theorem 33). Its twenty pages spanning proof not only relies on the Coarse Menger Theorem for two paths (Theorem 27) but also required a deep insight into the new theory of graph-decompositions. Our proof technique also yields a new short proof of the case $\mathcal{X} = \{K_4^-\}$, which was first established by Fujiwara and Papasoglu [70].

Additionally, we show in Chapter 11 that the Fat minor conjecture is true for finite graphs for $\mathcal{X} = \{K_3, K_{1,3}\}$ (quasi-isometric to a disjoint union of paths), $\mathcal{X} = \{K_{1,3}\}$ (cycles and paths) and $\mathcal{X} = \{K_3, W\}$ (subdivided stars), where W is the graph obtained from the disjoint union of two paths of length two by adding an edge between their inner vertices (see Figure 11.1). In fact, we prove a stronger result that finds a graph in $\mathcal{X}$ even as a 'quasi-geodesic topological minor' in G. These are a metric version of topological minors, whose model in G is 'quasi-geodesically' embedded (see Section 11.1.4).

1.2.4 Chapter 13: Asymptotic half-grid and full-grid minors

In Chapter 13 we apply the paradigm of coarse graph theory to quasi-transitive, infinite graphs. Consider any connected, locally finite, quasi-transitive graph whose cycle space is generated by

[5]We remark that an alternative proof was independently obtained by Georgakopoulos and Papasoglu [75], who also state the Coarse Menger Conjecture.

cycles of bounded lengths. Note that these include all locally finite Cayley graphs of finitely presented groups – a rich and well-studied class of groups. We prove that such graphs are quasi-isometric to trees if and only if they do not contain the full-grid[6] as an asymptotic minor.[7] This solves problems of Georgakopoulos and Papasoglu [75] and of Georgakopoulos and Hamann [74] for such graphs.

Additionally, we present a sufficient condition for when a (not necessarily quasi-transitive) graph whose cycle space is generated by cycles of bounded length contains the half-grid[8] as an asymptotic minor.

[6]The *full-grid* is the grid on $\mathbb{Z} \times \mathbb{Z}$.

[7]A graph contains a graph X as an asymptotic minor if it contains X as a K-fat minor for every $K \in \mathbb{N}$.

[8]The *half-grid* is the grid on $\mathbb{N} \times \mathbb{N}$.

2 Terminology

In this chapter we gather the definitions needed for Parts I and II. The definitions needed for Part III can be found in Chapter 9, which is located at the beginning of Part III. Additionally, some chapters have their own preliminary section, where we introduce definitions that are only needed for the respective chapter.

The notation and definitions in this thesis mostly follow Diestel's textbook [41]; in particular we may speak of a vertex $v \in G$ (rather than $v \in V(G)$) and an edge $e \in G$. We denote the set $\{1, \dots, n\}$ with $n \in \mathbb{N}$ by $[n]$.

Given a graph G and two subgraphs H_1, H_2 of G, the subgraph $H_1 \cup H_2$ is the subgraph with vertex set $V(H_1) \cup V(H_2)$ and edge set $E(H_1) \cup E(H_2)$.

If $P = p_0 \dots p_n$ is a path in a graph G, then we denote by $p_i P p_j$ for $i, j \in [n]$ the subpath $p_i p_{i+1} \dots p_j$ of P. Further, we denote by $\bar{P}$ the path $p_n \dots p_0$, and by $\mathring{P}$ the subpath $p_1 P p_{n-1}$ of P. A path is *trivial* if it has length 0.

For two sets U_1, U_2 of vertices of G, an *U_1–U_2 path* meets U_1 precisely in its first vertex and U_2 precisely in its last vertex. For a subgraph Y of G, a *Y-path* is a non-trivial path which meets Y precisely in its endvertices. Analogously, for a set U of vertices of G, a *U-path* is any Y-path where Y is the subgraph of G with vertex set U and no edges. A path *through* a set U of vertices of G is a path in G whose internal vertices are contained in U.

The *neighbourhood* $N_G(U)$ of a set U of vertices in a graph G is the set of all vertices in $V(G) \setminus U$ that are neighbours of some $u \in U$. We denote by $\mathcal{C}(G)$ the *set of components* of the graph G. Given sets $U' \subseteq U$ of vertices of G, a component C of $G - U$ *attaches* to U' if C has a neighbour in U'. The *boundary* $\partial_G Y$ of a subgraph Y of G is the set $N_G(V(G-Y))$ of vertices of Y that send in G an edge outside of Y. For example, the boundary $\partial_G C$ of a component of $G - U$ is $N_G(U) \cap V(C)$.

In the following Sections 2.1–2.5 we recall definitions about separation systems and tangles (Sections 2.1–2.3), tree-decompositions (Section 2.4), and infinite graphs (Section 2.5). In Sections 2.1–2.4 our presentation follows [3, 5], and in Section 2.5 we follow [12].

2.1 Separation systems of graphs

Let G be a graph. A *separation* of G is a set $\{A, B\}$ of subsets of $V(G)$ such that $A \cup B = V(G)$ and there is no edge in G between $A \setminus B$ and $B \setminus A$. The sets A and B are the *sides* of the separation $\{A,B\}$, and $A \setminus B$ and $B \setminus A$ are its *strict sides*. A separation $\{A,B\}$ of G is *proper* if none of its sides equals $V(G)$. The *order* $|\{A,B\}|$, or $|A,B|$ for short, of a separation $\{A,B\}$ is the size $|A \cap B|$ of its *separator* $A \cap B$. We write $U(G)$ for the set of all separations of G and let $S_k(G) := \{\{A,B\} \in U(G) : |A \cap B| < k\}$ for every $k \in \mathbb{N} \cup \{\aleph_0\}$. Note that $U(G) = S_{\aleph_0}(G)$ for every finite graph G.

The *orientations* of a separation $\{A,B\}$ are the *oriented separations* (A,B) and (B,A). We refer to A as the *small* side of the oriented separation (A,B) and to B as its *big* side. Given a set S of separations of G, we write $\vec{S} := \{(A,B) : \{A,B\} \in S\}$ for the set of all their orientations.

We will use terms defined for unoriented separations also for oriented ones and vice versa if that is possible without causing ambiguities. Moreover, if the context is clear, we will simply refer to both oriented and unoriented separations as 'separations'. If we do not need to know about the sides of a separation, we sometimes denote separations with s, and their orientations with $\vec{s}, \overleftarrow{s}$. Note that there are no default orientations: once we denoted one orientation by $\vec{s}$ the other one will be $\overleftarrow{s}$, and vice versa.

A separation (A,B) is *trivial (in $\vec{S}$)*, for some $S \subseteq \vec{S}_{\aleph_0}(G)$, if $(A,B) = (X, V(G))$ for some $X \subseteq V(G)$, and there exists $\{C,D\} \in S \setminus \{\{A,B\}\}$ with $X \subseteq A \cap B$. A separation (A,B) of G is *left-tight* (*right-tight*) if the neighbourhood in G of some component of $G[A \setminus B]$ ($G[B \setminus A]$) equals $A \cap B$. Moreover, $\{A,B\}$ is *tight* if (A,B) is left- and right-tight.

The oriented separations of a graph G are partially ordered by

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

Two separations $\{A,B\}$ and $\{C,D\}$ of G are *nested* if they have orientations which can be compared; otherwise they *cross*. A set of (unoriented) separations is *nested* if all its elements are pairwise nested. We use the terms 'nested' and 'cross' analogously for oriented separations: two separations, oriented or not, are *nested* if their underlying unoriented separations are nested; otherwise they *cross*.

For any pair of separations (A,B) and (C,D) also their *infimum* $(A,B) \wedge (C,D) := (A \cap C, B \cup D)$ and their *supremum* $(A,B) \vee (C,D) := (A \cup C, B \cap D)$ are separations of G; we call $\{A \cap C, B \cup D\}$, $\{A \cup C, B \cap D\}$, $\{B \cap C, A \cup D\}$ and $\{B \cup C, A \cap D\}$ the *corner separations* of $\{A,B\}$ and $\{C,D\}$.

Lemma 2.1.1 ([41, Lemma 12.5.5])**.** *Let r, s be two crossing separations of a graph G. Every separation of G that is nested with both r and s is also nested with all four corner separations of r and s.*

Moreover, it is easy to check that the orders of the infimum and supremum of any two separations $(A,B),(C,D)$ of a graph sum up to

$$|(A\cap C)\cap(B\cup D)|+|(A\cup C)\cap(B\cap D)|=|A\cap B|+|C\cap D|.$$

The inequality '$\leqslant$' implied by the equality is the important part here and is called *submodularity*.

A set $\sigma\subseteq\vec{S}_{\aleph_0}(G)\setminus\{(V(G),V(G))\}$ of separations is called a *star* if for any $(A,B),(C,D)\in\sigma$ it holds that $(A,B)\leqslant(D,C)$. A star σ is *proper (in $S_k(G)$)* if for every distinct $\vec{r},\vec{s}\in\sigma$ the relation $\vec{r}\leqslant\overleftarrow{s}$ is the only one, i.e. $\vec{r}\nleqslant\vec{s}$, $\vec{r}\ngeqslant\vec{s}$ and $\vec{r}\ngeqslant\overleftarrow{s}$, and $\sigma\neq\{\overleftarrow{s}\}$ for a trivial separation $\vec{s}\in\vec{S}_k(G)$.

The *interior* of a star $\sigma\subseteq\vec{S}_{\aleph_0}(G)$ is the intersection $\operatorname{int}(\sigma):=\bigcap_{(A,B)\in\sigma}B$, and the *torso of* σ, denoted by $\operatorname{torso}(\sigma)$, is the graph that is obtained from $G[\operatorname{int}(\sigma)]$ by adding the edge vu whenever $v\neq u\in\operatorname{int}(\sigma)$ lie together in some separator of a separation in σ. We call these added edges that lie in $\operatorname{torso}(\sigma)$ but not in G *torso edges*. We refer to $\operatorname{torso}(\sigma)$ as the *torso at* σ.

The partial order on $\vec{S}_{\aleph_0}(G)$ induces a partial order on the set of proper stars in $\vec{S}_{\aleph_0}(G)$: if $\sigma,\tau\subseteq\vec{S}_{\aleph_0}(G)$ are stars of proper separations, then $\sigma\leqslant\tau$ if and only if for every $(A,B)\in\sigma$ there exists some $(C,D)\in\tau$ such that $(A,B)\leqslant(C,D)$ [43]. Note that the restriction to proper stars is important in order to make the relation antisymmetric. A proper star in some set $\vec{S}\subseteq\vec{S}_{\aleph_0}(G)$ is *maximal* if it is a maximal element in the set of all proper stars in $\vec{S}$ with the partial order defined above.

2.2 Profiles and tangles in graphs

Let G be a graph. An *orientation* of a set $S\subseteq S_{\aleph_0}(G)$ is a set $O\subseteq\vec{S}$ which contains, for every $\{A,B\}\in S$, exactly one of its orientations (A,B) and (B,A). It is *consistent* if it does not contain both (B,A) and (C,D) whenever $(A,B)<(C,D)$ for distinct $\{A,B\},\{C,D\}\in S$. An orientation is *regular* if it does not contain $(V(G),A)$ for any subset $A\subseteq V(G)$.

An orientation O of S *lives* in a star $\sigma\subseteq\vec{S}$ (or equivalently σ *is home* to O) if $\sigma\subseteq O$. If $\mathcal{O}$ is a set of consistent orientations of S, we call a star $\sigma\subseteq\vec{S}$ *essential (for $\mathcal{O}$)* if some $O\in\mathcal{O}$ lives in σ. Otherwise σ is called *inessential (for $\mathcal{O}$)*. The star $\sigma\subseteq\vec{S}$ is *exclusive (for $\mathcal{O}$)* if it is contained in exactly one orientation in $\mathcal{O}$. If $O\in\mathcal{O}$ is that orientation, we say that σ is *O-exclusive (for $\mathcal{O}$)*.

A separation $\{A,B\}\in S$ *distinguishes* two orientations of S if they orient $\{A,B\}$ differently. $\{A,B\}$ distinguishes them *efficiently* if they are not distinguished by any separation of smaller order. A set of separations $N\subseteq S$ *(efficiently) distinguishes* a set $\mathcal{O}$ of consistent orientations of S if any two distinct orientations in $\mathcal{O}$ are (efficiently) distinguished by some separation in N.

Let $S\subseteq S_{\aleph_0}(G)$, and let $\mathcal{F}$ be a set of subsets of $\vec{S}_{\aleph_0}(G)$. We call an orientation O of S an *$\mathcal{F}$-tangle of S* if O is consistent and *avoids* $\mathcal{F}$, i.e. if O does not contain any element of $\mathcal{F}$ as a subset.

For some $k \in \mathbb{N} \cup \{\aleph_0\}$, a *k-tangle* in G (or *tangle of order* k) is a $\mathcal{T}_k$-tangle of $S_k(G)$ where

$$\mathcal{T}_k := \{\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} \subseteq \vec{S}_k(G) : G[A_1] \cup G[A_2] \cup G[A_3] = G\}.$$

We sometimes call the $\aleph_0$-tangles in G *infinite tangles*. We denote with $\mathcal{T}_k^*$ the set of all stars in $\mathcal{T}_k$, and we abbreviate $\mathcal{T}_{\aleph_0}, \mathcal{T}_{\aleph_0}^*$ with $\mathcal{T}, \mathcal{T}^*$, respectively. The elements of $\mathcal{T}$ are called *forbidden triples*. A forbidden triple in $\mathcal{T}$ such that two of the (A_i, B_i) are equal is a *forbidden tuple*. Note that k-tangles are consistent and regular. Moreover, it follows from the definition of $\mathcal{T}_k$ that if (A, B) is a $\leqslant$-maximal separation in a tangle in G, then $G[B \setminus A]$ is connected.

For some $k \in \mathbb{N}$, we call a consistent orientation P of $S_k(G)$ a *k-profile* in G if it satisfies that

$$\text{for all } (A, B), (C, D) \in P \text{ the separation } (B \cap D, A \cup C) \text{ does not lie in } P. \qquad (*)$$

Tangles and profiles of unspecified order are referred to as *tangles / profiles in G*. Clearly, every k-tangle is a k-profile, and hence has the *profile property*, i.e. it satisfies $(*)$. We say that a set $\mathcal{F}$ of subsets of $\vec{S}_k(G)$ is *profile-respecting (for $\vec{S}_k(G)$)* if every $\mathcal{F}$-tangle of $S_k(G)$ is a k-profile in G.

We can describe the k-profiles in G as $\mathcal{F}$-tangles of $S_k(G)$ as follows. Set

$$\mathcal{P}_k := \{\{(A, B), (B \cap C, A \cup D), (B \cap D, A \cup C)\} \subseteq \vec{S}_k(G) : (A, B), (C, D) \in \vec{S}_k(G)\} \qquad (2.1)$$

Then the set of all $\mathcal{P}_k$-tangles of $S_k(G)$ is exactly the set of all k-profiles in G (cf. [45, Lemma 11]).

Lemma 2.2.1 ([64, Lemma 6.1])**.** *Let P, P' be two distinct regular profiles in an arbitrary graph G. If $\{A, B\}$ is a separation of finite order that efficiently distinguishes P and P', then $\{A, B\}$ is tight.*

2.3 Tangle-tree duality in finite graphs

Let G be a graph and $S \subseteq S_{\aleph_0}(G)$. An *S-tree* is a pair (T, α) of a tree T and a map $\alpha : \vec{E}(T) \to \vec{S}$ from the oriented edges $\vec{E}(T)$ of T to $\vec{S}$ such that $\alpha(\overleftarrow{e}) = (B, A)$ whenever $\alpha(\vec{e}) = (A, B)$. If the tree T is finite, then we call (T, α) a *finite S-tree*. If $x \in V(T)$ is a leaf of T and $t \in V(T)$ its unique neighbour, then we call $\alpha(x, t) \in \vec{S}$ a *leaf separation (of T)*.

An S-tree (T, α) is *over* a set $\mathcal{F}$ of subsets of $\vec{S}$ if $\{\alpha(t', t) : (t', t) \in \vec{E}(T)\} \in \mathcal{F}$ for every node $t \in V(T)$. If (T, α) is over a set of stars in $\vec{S}$ and $t \in V(T)$, then we call the star $\sigma_t := \{\alpha(t', t) : (t', t) \in \vec{E}(T)\} \subseteq \vec{S}$ the *star associated with t (in T)*.

Let $S \subseteq S_{\aleph_0}(G)$, $\vec{r} \in \vec{S}_{\aleph_0}(G)$, and set $S_{\geqslant \vec{r}} := \{x \in S : \vec{x} \geqslant \vec{r} \text{ or } \overleftarrow{x} \geqslant \vec{r}\}$. Further, let $\vec{s} \in \vec{S}$. We say that $\vec{s}$ *emulates* $\vec{r}$ *in* $\vec{S}$ if $\vec{s} \geqslant \vec{r}$ and for every $\vec{x} \in \vec{S}$ with $\vec{x} \geqslant \vec{r}$ it holds that $\vec{s} \vee \vec{x} \in \vec{S}$. Given a set $\mathcal{F}$ of stars in $\vec{S}_{\aleph_0}(G)$, 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

$\sigma \in \mathcal{F}$ with $\sigma \subseteq \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$ that contains an element $\vec{t} \geqslant \vec{r}$ it holds that $\{\vec{s} \vee \vec{t}\} \cup \{\overleftarrow{s} \wedge \vec{t'} : \vec{t'} \in \sigma \setminus \{\vec{t}\}\}$ is again a star in $\mathcal{F}$.

A set $\mathcal{F}$ of stars in $\vec{S}$ is *closed under shifting* if whenever $\vec{s} \in \vec{S}$ emulates some $\vec{r} \in \vec{S}$, then it also emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$.

Given a graph G and a set $\mathcal{F}$ of stars in $\vec{S}_{\aleph_0}(G)$, a set $S \subseteq S_{\aleph_0}(G)$ is *$\mathcal{F}$-separable* if for every two non-trivial $\vec{r}, \vec{r'} \in \vec{S} \setminus \{(V(G), V(G))\}$ with $\vec{r} \leqslant \vec{r'}$ and $\{\overleftarrow{r}\}, \{\overleftarrow{r'}\} \notin \mathcal{F}$ there exists an $\vec{s} \in \vec{S}$ such that $\vec{s}$ and $\overleftarrow{s}$ emulate $\vec{r}$ and $\overleftarrow{r'}$, respectively, in $\vec{S}$ for $\mathcal{F}$. Further, $\mathcal{F}$ is *standard* for $\vec{S}$ if $\{\overleftarrow{r}\} \in \mathcal{F}$ for all $\vec{r} \in \vec{S}$ that are trivial in $\vec{S}$.

Theorem 2.3.1 ([55, Theorem 4.3])**.** *Let G be any graph, $S \subseteq S_{\aleph_0}(G)$ finite, and let $\mathcal{F}$ be a set of stars in $\vec{S}_{\aleph_0}(G)$, standard for $\vec{S}$. If $\vec{S}$ is $\mathcal{F}$-separable, exactly one of the following assertions holds:*

(i) *There exists an $\mathcal{F}$-tangle of S.*

(ii) *There exists an S-tree over $\mathcal{F}$.*

2.4 Tree-decompositions

A *tree-decomposition* of a graph G is a pair $(T, \mathcal{V})$ that consists of a tree T and a family $\mathcal{V} = (V_t)_{t \in T}$ of subsets of $V(G)$ indexed by the nodes of T and satisfies the following two conditions:

(T1) $G = \bigcup_{t \in T} G[V_t]$,

(T2) for every vertex v of G, the subgraph of T induced by $\{t \in T : v \in V_t\}$ is connected.

Whenever a tree-decomposition is introduced as $(T, \mathcal{V})$ in this thesis, we tacitly assume that $\mathcal{V} = (V_t)_{t \in T}$.

The sets V_t are the *bags* of the tree-decomposition, the induced subgraphs $G[V_t]$ on the bags are its *parts*, and T is its *decomposition tree*. The sets $V_e := V_{t_0} \cap V_{t_1}$ for edges $e = t_0 t_1$ of T are the *adhesion sets* of $(T, \mathcal{V})$ and its *adhesion* is the maximal size of an adhesion set. If all its adhesion sets are finite, then $(T, \mathcal{V})$ has *finite adhesion*.

A tree-decomposition $(T, \mathcal{V})$ has *width* less than $k \in \mathbb{N}$ if all its bags V_t have size at most k. A graph G has *tree-width* at most $k \in \mathbb{N}$ if it admits a tree-decomposition of width at most k. If there exists such a minimal $k \in \mathbb{N}$ we denote it by $\mathrm{tw}(G)$, and say that G has *finitely bounded tree-width*. We say that a graph has *finite tree-width* if it admits a tree-decomposition into finite parts.

In a tree-decomposition $(T, \mathcal{V})$ of G every (oriented) edge $\vec{e} = (t_0, t_1)$ of the decomposition tree T induces a separation of G as follows: Write T_0 for the component of $T - e$ containing t_0 and T_1 for the one containing t_1. Then (U_0, U_1) is a separation of G where $U_i := \bigcup_{t \in T_i} V_t$ for $i = 0, 1$ [41, Lemma 12.3.1]. We say that $\vec{s}_{\vec{e}} = \vec{s}_{(t_0, t_1)} := (U_0, U_1)$ and $s_e := \{U_0, U_1\}$ are *induced* by $\vec{e}$ and e, respectively, or more generally by $(T, \mathcal{V})$. Note that the separator of $\{U_0, U_1\}$ is V_e.

If $(T, \mathcal{V})$ and $(\tilde{T}, \tilde{\mathcal{V}})$ are both tree-decompositions of G, then $(T, \mathcal{V})$ *refines* $(\tilde{T}, \tilde{\mathcal{V}})$ if the set of separations induced by the edges of T is a superset of the set of separations induced by the edges of $\tilde{T}$.

For every node t of T, we write σ_t for the set of separations of G induced by the (inwards) oriented edges $\vec{e} = (s, t)$ for $s \in N_T(t)$, i.e. $\sigma_t = \{\vec{s}_{(u,t)} : u \in N_T(t)\}$. It is easy to see that such σ_t are stars of separations and that their interior is precisely V_t. We call σ_t the *star associated with* t, and we refer to torso(σ_t) as the *torso* of $(T, \mathcal{V})$ *at* t.

The *leaf separations* of a tree-decomposition $(T, \mathcal{V})$ are those separations of G that are induced by the oriented edges (s, t) of T where t is a leaf of T. We refer to this separation also as leaf separation *at* t.

A tree-decomposition $(T, \mathcal{V})$ of G is *canonical* if the construction Ψ of $(T, \mathcal{V})$ commutes with all isomorphisms $\varphi : G \to G'$: if φ maps the bags V_t of $(T, \mathcal{V})$ to bags $V'_{t'}$ of $(T', \mathcal{V}') := \Psi(G')$ such that $t \mapsto t'$ is an isomorphism $T \to T'$.

Given a tree-decomposition $(T, \mathcal{V})$ of a graph G, we say that a consistent orientation O of $S_k(G)$, for some $k \in \mathbb{N} \cup \{\aleph_0\}$, *lives* at a node t of T, or in the bag V_t, if $\sigma_t \subseteq O$. Further, given a set $\mathcal{O}$ of consistent orientations of $S_k(G)$, we call a node t of T and its bag V_t *essential (for $\mathcal{O}$)* if there is an orientation in $\mathcal{O}$ that lives at t and otherwise *inessential (for $\mathcal{O}$)*. A bag V_t of $(T, \mathcal{V})$ is *exclusive (for $\mathcal{O}$)* if precisely one orientation in $\mathcal{O}$ lives in V_t.

A tree-decomposition $(T, \mathcal{V})$ *(efficiently) distinguishes* two profiles if there is an edge $t_0t_1 \in E(T)$ such that $\{U_{t_0}, U_{t_1}\}$ (efficiently) distinguishes them.

2.5 Infinite graphs

A graph is *locally finite* if all its vertices have finite degree.

A *ray* is a one-way infinite path, and a *double ray* is a two-way infinite path. A graph is *rayless* if it contains no ray. A *tail* of a (double) ray R is any subray $S \subseteq R$. If $R = r_0r_1 \ldots$ is a ray, then we denote by r_iRr_j for $i, j \in \mathbb{N}$ the subpath $r_i \ldots r_j$ of R, and by r_iR or $R_{\geqslant i}$ the tail $r_ir_{i+1} \ldots$ of R. Further, we denote by Rr_i or $R_{\leqslant i}$ the subpath $r_0 \ldots r_i$ of R. We use these notions analogously for double rays; in particular, if $R = \ldots r_{-1}r_0r_1 \ldots$ is a double ray, then Rr_i and $R_{\leqslant i}$ denote the tail $r_ir_{i-1} \ldots$ of R.

A *comb* is a union of a ray R with infinitely many pairwise disjoint finite paths which have precisely their first vertex on R; we call the last vertices of these paths the *teeth* of the comb and refer to R as its *spine*. The following observation about combs in infinite graphs is well-known and follows immediately from the *Star-Comb Lemma*; see e.g. [41, Lemma 8.2.2] for a proof.

Lemma 2.5.1. *Let U be an infinite set of vertices in a locally finite, connected graph G. Then G contains a comb with all teeth in U.*

Moreover, every infinite, connected graph has a vertex of infinite degree or contains a ray.

2.5.1 Ends

An *end* of a graph G is an equivalence class of rays in G where two rays are equivalent if they are joined by infinitely many disjoint paths in G or, equivalently, if for every finite set $U \subseteq V(G)$ they have a tail in the same component of $G - U$. The set of all ends of a graph G is denoted by $\Omega(G)$. A *(double) ε-ray* is a (double) ray whose tails are all contained in ε.

For every finite set $X \subseteq V(G)$ and every end ε of G, there is a unique component of $G - X$ which contains a tail of some, or equivalently every, ε-ray; we denote this component by $C_G(X, \varepsilon)$ and say that ε *lives* in $C_G(X, \varepsilon)$. We denote by $\Omega_G(X, \varepsilon)$ the set of all end of G which live in $C_G(C, \varepsilon)$. The collection of all the $\Omega_G(X, \varepsilon)$ with finite $X \subseteq V(G)$ and $\varepsilon \in \Omega(G)$ form a basis for a topology of $\Omega(G)$.

A vertex v of G *dominates* and end ε of G if $v \in C_G(X, \varepsilon)$ for every finite $X \subseteq V(G - v)$. We write $\mathrm{Dom}(\varepsilon)$ for the set of all vertices of G which dominate ε. Note that ends of locally finite graphs have no dominating vertices. For a connected subgraph C of G with finite $N_G(C)$, we write $\mathrm{Dom}(C)$ for all vertices which dominate some end of G that lives in C. Note that $\mathrm{Dom}(C)$ is a subset of $V(C) \cup N_G(C)$.

The *degree* of an end ε of G is defined as $\deg(\varepsilon) := \sup\{|\mathcal{R}| : \mathcal{R}$ is a family of disjoint rays in $\varepsilon\}$. Halin [81, Satz 1] showed that this supremum is always attained. If an end has infinite degree, then it is *thick*. Together with the number $\mathrm{dom}(\varepsilon) := |\mathrm{Dom}(\varepsilon)|$ of vertices dominating ε, the degree sums up to the *combined degree* $\Delta(\varepsilon) := \deg(\varepsilon) + \mathrm{dom}(\varepsilon)$ of ε.

2.5.2 Critical vertex sets

Given a set U of vertices of a graph G, a component C of $G - U$ is *tight at U in G* if $N_G(C) = U$. By slight abuse of notation, we will refer to such components C as *tight* components of $G - U$. We write $\mathcal{C}_U := \mathcal{C}(G - U)$ for the set of components of $G - U$ and $\breve{\mathcal{C}}_U := \breve{\mathcal{C}}(G - U) \subseteq \mathcal{C}_U$ for the set of all tight components C of $G - U$. A *critical vertex set* of G is any finite set $U \subseteq V(G)$ such that the set $\breve{\mathcal{C}}_U$ is infinite. We denote by $\mathrm{crit}(G)$ the set of all critical vertex sets of G.

A graph is *tough* if it has no critical vertex set, or equivalently, if deleting finitely many vertices never leaves infinitely many components. The following theorem was first proved by Polat [107, Theorems 3.3 & 3.8]; see also [12, Theorem 2.5] for a short proof.

Theorem 2.5.2. *Every tough, rayless graph is finite.*

Part I

Tangles and tree-decompositions of infinite graphs

3 Linked tree-decompositions into finite parts

We prove that every graph which admits a tree-decomposition into finite parts has a rooted tree-decomposition into finite parts that is linked, tight and componental.

As an application, we obtain that every graph without half-grid minor has a lean tree-decomposition into finite parts, strengthening the corresponding result by Kříž and Thomas for graphs of finitely bounded tree-width. In particular, it follows that every graph without half-grid minor has a tree-decomposition which efficiently distinguishes all ends and critical vertex sets, strengthening results by Carmesin and by Elm and Kurkofka for this graph class.

As a second application of our main result, it follows that every graph which admits a tree-decomposition into finite parts has a tree-decomposition into finite parts that displays all the ends of G and their combined degrees, resolving a question of Halin from 1977. This latter tree-decomposition yields short, unified proofs of the characterisations due to Robertson, Seymour and Thomas of graphs without half-grid minor, and of graphs without binary tree subdivision.

This chapter is based on [12] and joint work with Raphael W. Jacobs, Paul Knappe and Max Pitz.

3.1 Introduction

All graphs in this chapter may be infinite, unless otherwise stated.

3.1.1 The main result

Our point of departure is Kříž and Thomas's result on linked tree-decompositions, which forms a cornerstone both in Robertson and Seymour's work [113] on well-quasi-ordering finite graphs, and in Thomas's result [119] that the class of infinite graphs of tree-width $< k$ is well-quasi-ordered under the minor relation for all $k \in \mathbb{N}$.

Theorem 3.1.1 (Thomas 1990 [120], Kříž and Thomas 1991 [95])**.** *For every $k \in \mathbb{N}$, every (finite or infinite) graph of tree-width $< k$ has a linked rooted tree-decomposition*[1] *of width $< k$.*

[1]There is also an unrooted version of this theorem but this is not needed for the well-quasi-ordering applications.

To make this result precise, recall that a *(rooted) tree-decomposition* $(T, \mathcal{V})$ of a possibly infinite graph G is given by a (rooted) *decomposition tree* T whose nodes t are assigned *bags* $V_t \subseteq V(G)$ or *parts* $G[V_t]$ of the underlying graph G such that $\mathcal{V} = (V_t)_{t \in T}$ covers G in a way that reflects the separation properties of T: Similarly as the deletion of an edge $e = st$ from T separates it into components $T_s \ni s$ and $T_t \ni t$, the corresponding sets $A^e_s = \bigcup_{x \in T_s} V_x$ and $A^e_t = \bigcup_{x \in T_t} V_x$ in the underlying graph G are separated by the *adhesion set* $V_e = V_s \cap V_t$. A graph G has

- *tree-width* $< k$ if it admits a (rooted) tree-decomposition into parts of size $\leqslant k$,
- *finitely bounded tree-width* if it has tree-width $< k$ for some $k \in \mathbb{N}$, and
- *finite tree-width* if it admits a (rooted) tree-decomposition into finite parts (of possibly unbounded size).[2]

Given a tree T rooted at a node r, its *tree-order* is given by $x \leqslant y$ for $x, y \in V(T) \cup E(T)$ if x lies on the (unique) $\subseteq$-minimal path rTy from r to y. For an edge $e = st$ of T with $s < t$, the *part above* e is $G{\uparrow}e = G[A^e_t]$ and the *part strictly above* e is $G\mathring{\uparrow}e = G{\uparrow}e - V_e$. A rooted tree-decomposition $(T, \mathcal{V})$ of a graph G is

- *linked* if for every two comparable nodes $s < t$ of T there are $\min\{|V_e| \colon e \in E(sTt)\}$ pairwise disjoint V_s–V_t paths in G [119],
- *tight* if for every edge e of T some component C of $G\mathring{\uparrow}e$ satisfies $N(C) = V_e$, and
- *componental* if for every edge e of T the part $G\mathring{\uparrow}e$ strictly above e is connected.

Kříž and Thomas proved Theorem 3.1.1 for every integer $k \in \mathbb{N}$. Our main result is that Theorem 3.1.1 also holds for $k = \aleph_0$. Graphs of tree-width $< \aleph_0$ are those graphs which admit a tree-decomposition into parts of size less than $< \aleph_0$, i.e. precisely the graphs of finite tree-width.

Theorem 1. *Every graph of finite tree-width admits a rooted tree-decomposition into finite parts that is linked, tight, and componental.*

We remark that, in contrast to Theorem 3.1.1, the linked rooted tree-decomposition in Theorem 1 is additionally tight and componental.

Note that achieving just a subset of the properties in Theorem 1 may be significantly easier. Recall that a *normal spanning tree* of a graph G is a rooted spanning tree T such that the endvertices of every edge of G are comparable in its tree-order. It is well known that the connected graphs of finite tree-width are precisely the graphs with normal spanning trees (see Theorem 3.2.2 below). Given a normal spanning tree T with root r, by assigning to each of its nodes t the bag $V_t := V(rTt)$ we obtain a rooted tree-decomposition $(T, \mathcal{V}_{\mathrm{NT}})$ into finite parts that is componental and linked, the latter albeit for the trivial reason that $s < t \in T$ implies $V_s \subseteq V_t$. However, this

[2] Note that for a graph to have finite tree-width we do not require that the tree-decomposition into finite parts also satisfies that $\liminf_{e \in E(R)} V_e$ is finite for all rays R in T, as it is sometimes [41, 115] done.

tree-decomposition clearly fails to be tight. Following an observation by Diestel [39], one can restore tightness by taking as bags the subsets $V_t' \subseteq V_t$ that consist only of those vertices in V_t that send a G-edge to a vertex above t in T. The new tree-decomposition $(T, \mathcal{V}'_{\mathrm{NT}})$ is then tight and componental, but in general no longer linked.

Having all three properties simultaneously is more challenging to achieve. In what follows, we hope to convince the reader of the usefulness of Theorem 1 by demonstrating the surprisingly powerful interplay of the properties of being linked, tight, and componental. We do so by presenting several applications in the following sections.

In Chapter 4 we provide examples showing that Theorem 1 appears to lie at the frontier of what is still true. For example, *linked* cannot be strengthened to its 'unrooted' version that requires the *linked* property between every pair of nodes and not just comparable ones (see Example 9).

3.1.2 Displaying end structure and excluding infinite minors

As our first application, we show that every tree-decomposition as in Theorem 1 displays the combinatorial and topological end structure of the underlying graph as follows, resolving a question by Halin from 1977 [83, §6]. (Halin's original work yields proofs of Theorems 1 and 2 for locally finite connected graphs with at most two ends.)

Theorem 2. *Every graph G of finite tree-width admits a rooted tree-decomposition into finite parts that homeomorphically displays all the ends of G, their dominating vertices, and their combined degrees.*

Recall that an *end* ε of a graph G is an equivalence class of rays in G where two rays are equivalent if for every finite set X they have a tail in the same component of $G - X$. We refer to this component of $G - X$ as $C_G(X, \varepsilon)$. Write $\Omega(G)$ for the *set of all ends* of G. A vertex v of G *dominates* an end ε of G if it lies in $C_G(X, \varepsilon)$ for every finite set X of vertices other than v. We denote the set of all vertices of G which dominate an end ε by $\mathrm{Dom}(\varepsilon)$. The *degree* $\deg(\varepsilon)$ of an end ε of G is the supremum over all cardinals κ such that there exists a set of κ pairwise disjoint rays in ε, and its *combined degree* is $\Delta(\varepsilon) := \deg(\varepsilon) + |\,\mathrm{Dom}(\varepsilon)|$.

In a rooted tree-decomposition $(T, \mathcal{V})$ of a graph G into finite parts, every end ε of G *gives rise* to a single rooted ray R_ε in T that starts at the root and then always continues upwards along the unique edge $e \in T$ with $C_G(V_t, \varepsilon) \subseteq G\mathring{\uparrow}e$. This yields a map $\varphi\colon \Omega(G) \to \Omega(T)$, $\varepsilon \mapsto R_\varepsilon$. A rooted tree-decomposition $(T, \mathcal{V})$

- *displays* the ends of G if φ is a bijection [26],
- *displays the ends homeomorphically* if φ is a homeomorphism [93],

- *displays all dominating vertices* if $\liminf_{e \in E(R_\varepsilon)} V_e = \mathrm{Dom}(\varepsilon)^3$ for all $\varepsilon \in \Omega(G)$, and
- *displays all combined degrees* if $\liminf_{e \in E(R_\varepsilon)} |V_e| = \Delta(\varepsilon)$ for all $\varepsilon \in \Omega(G)$.

Roughly, every componental, rooted tree-decomposition homeomorphically displays all the ends [93, Lemma 3.1], every tight such tree-decomposition displays all dominating vertices (Lemma 3.3.2), and every such linked tight tree-decomposition displays all combined degrees (Lemma 3.3.3).

Diestel's tree-decomposition $(T, \mathcal{V}'_{\mathrm{NT}})$ mentioned above was not linked, but still displays all the ends and their dominating vertices. This observation allowed Diestel [39] to obtain a short proof of Robertson, Seymour and Thomas's characterisation that graphs without subdivided infinite cliques are precisely the graphs that admit a tree-decomposition $(T, \mathcal{V})$ into finite parts such that for every ray R of T, the set $\liminf_{e \in E(R)} V_e$ is finite. Extending this idea, we now use Theorem 2 to provide short, unified proofs for two other related results by Robertson, Seymour and Thomas.

Corollary 3.1.2 (Robertson, Seymour & Thomas 1995 [115, (2.6)])**.** *A graph contains no half-grid minor if and only if it admits a tree-decomposition $(T, \mathcal{V})$ into finite parts such that for every ray R of T we have $\liminf_{e \in E(R)} |V_e| < \infty$.*

Proof. We prove here only the hard implication.[4] So assume that G is a graph without half-grid minor. By a result of Halin [84, 104], every connected graph without a subdivided $K_{\aleph_0}$ has a normal spanning tree. Applying this to every component of G we see that G has finite tree-width. So we may apply Theorem 2.

By Halin's grid theorem [81, 97], every end of infinite degree contains a half-grid minor. And by a routine exercise, every end dominated by infinitely many vertices contains a subdivided $K_{\aleph_0}$. So every end of G has finite combined degree. Let $(T, \mathcal{V})$ be the tree-decomposition into finite parts from Theorem 2. Then $\liminf_{e \in R} |V_e|$ equals the combined degree of the end of G giving rise to R, and hence is finite, as desired. □

Corollary 3.1.3 (Seymour & Thomas 1993 [116, (1.5)])**.** *A graph contains no subdivision of the binary tree T_2 if and only if it admits a tree-decomposition $(T, \mathcal{V})$ into finite parts such that for every ray R in T we have $\liminf_{e \in E(R)} |V_e| < \infty$ and T contains no subdivision of T_2.*

Proof. Again, we only prove the hard implication.[4] Assume that G is a graph without a subdivision of T_2. In particular, G contains no half-grid minor. As above, G has finite tree-width and every end of G has finite combined degree. Let $(T, \mathcal{V})$ be the tree-decomposition into finite parts from Theorem 2. Then $\liminf_{e \in E(R)} |V_e|$ equals the combined degree of the end of G giving rise to R, and hence is finite.

[3] The set-theoretic $\liminf_{n \in \mathbb{N}} A_n$ consists of all points that are contained in all but finitely many A_n. For a ray $R = v_0 e_0 v_1 e_1 v_1 \dots$ in T, one gets $\liminf_{e \in E(R_\varepsilon)} V_e = \bigcup_{n \in \mathbb{N}} \bigcap_{i \geqslant n} V_{e_i}$.

[4] See [115, 116] for details about the 'easy' implications.

We complete the proof by showing that T contains no subdivided binary tree. Recall Jung's characterisation that a graph G contains no end-injective[5] subdivided T_2 if and only if the end space of G is scattered[6] [91, §3]. Hence, if G contains no subdivided binary tree, then its end space is scattered. Since $(T, \mathcal{V})$ displays the ends of G homeomorphically, we conclude that the end space of T is scattered, so T contains no end-injective subdivided T_2, by the converse of Jung's result. As all subtrees of trees are end-injective, this yields the desired result. □

Note that Theorem 1 shows that one can always require the witnessing tree-decompositions for Corollaries 3.1.2 and 3.1.3 to be linked, tight, and componental.

3.1.3 Lean tree-decompositions

We already mentioned that there exist even stronger versions of the Kříž-Thomas Theorem 3.1.1. See the articles by Bellenbaum and Diestel [16] and by Erde [67] for a modern treatment of the finite case.

Theorem 3.1.4 (Thomas 1990 [120], Kříž and Thomas 1991 [95])**.** *For every $k \in \mathbb{N}$, every (finite or infinite) graph of tree-width $< k$ has a lean tree-decomposition of width $< k$.*

Here, a tree-decomposition of a graph G is *lean* if for every two (not necessarily distinct) nodes $t_1, t_2 \in T$ and vertex sets $Z_1 \subseteq V_{t_1}$ and $Z_2 \subseteq V_{t_2}$ with $|Z_1| = |Z_2| =: \ell \in \mathbb{N}$, either G contains ℓ pairwise disjoint Z_1–Z_2 paths or there exists an edge $e \in t_1 T t_2$ with $|V_e| < \ell$. There are two ways in which 'lean' is stronger than 'linked': First, every lean tree-decomposition $(T, \mathcal{V})$ satisfies that for *every* two nodes $s \neq t \in T$ there are $\min\{|V_e| \colon e \in E(sTt)\}$ pairwise disjoint V_s–V_t paths in G. In a way, this is the unrooted version of the property 'linked' as introduced above which only required this property when s and t are comparable. The other difference between 'linked' and 'lean' is that a bag of a lean tree-decomposition is only large if it is highly connected, which follows by considering the case $t_1 = t_2$ in the definition of 'lean'.

Now it is natural to ask whether Theorem 3.1.4 extends to graphs of finite tree-width. Given that the tree-decomposition $(T, \mathcal{V}_{\mathrm{NT}})$ from above is linked but not necessarily lean, the following question appears to be non-trivial:

Does every graph of finite tree-width admit a lean tree-decomposition into finite parts?

Unfortunately, the answer to this question is in the negative. In Chapter 4 we construct a locally finite and planar graph which does not admit any lean tree-decomposition. In particular, not even graphs without $K_{\aleph_0}$ minor admit lean tree-decompositions.

[5]A rooted subtree T of a graph G is *end-injective* if distinct rooted rays in T belong to distinct ends of G.

[6]The precise definition of 'scattered' is not relevant here, it is only important that 'scattered' is a property of topological spaces and hence preserved under homeomorphisms.

Yet, we can extend Theorem 3.1.4 in the optimal way to graphs without half-grid minor by post-processing the tree-decomposition from Theorem 1 using the finite case of Theorem 3.1.4.

Theorem 3. *Every graph G without half-grid minor admits a lean tree-decomposition into finite parts. Moreover, if the tree-width of G is finitely bounded, then the lean tree-decomposition can be chosen to have width* $\operatorname{tw}(G)$.

The first part of Theorem 3 is a new result, while the 'moreover'-part reobtains the infinite case in Theorem 3.1.4. We recall that Robertson, Seymour and Thomas's strongest version [115, (12.11)] of Corollary 3.1.2 yields witnessing tree-decompositions which are additionally linked (even in its unrooted version). An immediate application of our detailed version of Theorem 3 (see Theorem 3' in Section 3.8) strengthens their result by showing that the witnessing tree-decomposition in Corollary 3.1.2 may even be chosen to be lean.

3.1.4 Displaying all infinities

A crucial tool for the proof of Theorem 1 are 'critical vertex sets', the second kind of infinities besides ends: A set X of vertices of G is *critical* if there are infinitely many *tight* components of $G - X$, that is components C of $G - X$ with $N_G(C) = X$. Polat [107] already noted that an infinite graph always contains either a ray/end or a critical vertex set.

As an intermediate step in the proof of Theorem 1 we obtain a similar tree-decomposition which 'displays all the infinities' of the underlying graph: A rooted tree-decomposition $(T, \mathcal{V})$

- *displays the critical vertex sets* if the map $t \mapsto V_t$ restricted to the infinite-degree nodes of T whose V_t is finite is a bijection to the critical vertex sets. In this context, we denote for every critical vertex set X the unique infinite-degree node of T whose bag is X by t_X,
- *displays the tight components of every critical vertex set cofinitely* if it displays the critical vertex sets such that for every critical vertex set X every $G\mathring{\uparrow}e$ with $e = t_X t \in T$ with $t_X <_T t$ is a tight component of $G - X$ and cofinitely many tight components of $G - X$ are some such $G\mathring{\uparrow}e$, and
- *displays the infinities* if it displays the ends homeomorphically, their combined degree, their dominating vertices, the critical vertex sets and their tight components cofinitely.

In general, a tree-decomposition that cofinitely displays the tight components of every critical vertex set can no longer be componental; but we still ensure that it is *cofinally componental*, i.e. along every rooted ray of T there are infinitely many edges e such that $G\mathring{\uparrow}e$ is connected.

Theorem 4. *Every graph of finite tree-width admits a linked, tight, cofinally componental, rooted tree-decomposition into finite parts which displays the infinities.*

Further, a tree-decomposition $(T, \mathcal{V})$ *efficiently distinguishes all the ends and critical vertex sets* if for any pair of ends and/or critical vertex sets one of the adhesion sets of $(T, \mathcal{V})$ is a size-wise minimal separator for them. Carmesin [26, Theorem 5.12] showed that for every graph there is a nested set of separations which efficiently distinguishes all its ends. Elm and Kurkofka [65, Theorem 1] extended this result by showing that there always exists a nested set of separations which efficiently distinguishes all ends and all critical vertex sets. If one aims not only for a nested set of separations but for a tree-decomposition, then the best current result is that every locally finite graph without half-grid minor admits even a tree-decomposition which distinguishes all its ends efficiently [90, Theorem 1]. In Chapter 4 (Construction 4.2.1) we present a planar graph witnessing that this result cannot be extended to graphs without $K_{\aleph_0}$ minor [11, Lemma 3.2], even if they are locally finite. Yet we extend it from locally finite graphs to arbitrary infinite graphs without half-grid minor:

Corollary 5. *Every graph G without half-grid minor admits a tree-decomposition of finite adhesion which distinguishes all its ends and critical vertex sets efficiently.*

We obtain Corollary 5 as an application of Theorem 3: Our proof of Theorem 3 yields that the obtained tree-decomposition also displays the infinities of G and hence in particular distinguishes them. Since the tree-decomposition is lean, one can conclude that it even does so efficiently, and thus is as desired for Corollary 5.

An equivalent way of stating Corollary 5 is that every graph G without half-grid minor admits a tree-decomposition of finite adhesion which efficiently distinguishes all its 'combinatorially distinguishable infinite tangles'.[7] This fits into a series of results [31, 64, 65, 90] extending the 'tree-of-tangles' theorem from Robertson and Seymour [114, (10.3)] to infinite graphs.

Additionally, in Chapter 5, we apply Theorem 4 to obtain a 'tangle-tree duality' theorem for infinite graphs,[8] which extends Robertson and Seymour's [114, (4.3) & (5.1)] other fundamental theorem about tangles to infinite graphs.

3.1.5 An open problem

Thomas famously conjectured that the class of countable graphs is well-quasi-ordered under the minor relation [119, (10.3)]. In light of the observation that all known counterexamples to well-quasi-orderings of infinite graphs [53, 94, 106, 118] do not have finite tree-width by the characterisation of finite tree-width graphs (i.e. graphs with normal spanning trees, see Theorem 3.2.2 below) in [105], it might be interesting to also consider the following, even stronger conjecture:

[7] See Section 3.9 for definitions.

[8] Every graph with no principal k-tangle not induced by an end of combined degree $< k$ admits a tree-decomposition witnessing this.

Conjecture 3.1.5. *The graphs of finite tree-width are well-quasi-ordered under the minor relation.*

We remark that already the simplest case of the conjecture – namely for the class of infinite graphs where all components are finite – is open.

3.1.6 How this chapter is organised

In Section 3.2 we present further definitions concerning tree-decompositions and their interplay with ends and critical vertex sets. We show in Section 3.3 that the tree-decomposition from Theorem 1 already satisfies Theorem 2. Section 3.4 consists of three parts. First, we give in Section 3.4.1 a brief overview over the proof of Theorem 1, which includes the statement of Theorems 6 and 7, our two main ingredients to the proof of Theorem 1. In Section 3.4.2, we state detailed versions of Theorems 1 and 4 which assert further properties of the respective tree-decompositions. In Section 3.4.3, we then reduce Theorems 1 and 4 to Theorems 6 and 7. We prove Theorem 6 in Section 3.5. In Section 3.6, we show some basic behaviour of paths and rays in torsos to prove Theorem 7 in Section 3.7. Finally, in Section 3.8, we prove Theorem 3 and Corollary 5.

3.2 Preliminaries

In this chapter, a tree T often contains a special node $\mathrm{root}(T)$, its *root*. A rooted tree T has a natural partial order $\leqslant_T$ on its nodes and edges, which depends on $r = \mathrm{root}(T)$: for two nodes $s, t \in V(T)$, we write sTt for the unique $\subseteq$-minimal path in T which contains x and y. We then set $x \leqslant_T t$ if $x \in rTt$ for every $t \in V(T)$ and $x \in V(T) \cup E(T)$. Further, for $e = st \in E(T)$ with $s \leqslant_T t$, we set $x \leqslant_T e$ if $x \in rTs$ or $x = e$. In a rooted tree, a *leaf* is any maximal node in the tree-order. The *down-closure* $\lceil x \rceil$ and *up-closure* $\lfloor x \rfloor$ of x in T are $\{y \in V(T) \mid y \leqslant x\}$ and $\{y \in V(T) \mid y \geqslant x\}$, respectively. We write $\lceil \mathring{x} \rceil$ and $\lfloor \mathring{x} \rfloor$ as shorthand for $\lceil x \rceil \smallsetminus \{x\}$ and $\lfloor x \rfloor \smallsetminus \{x\}$.

A rooted tree T in a graph G is *normal* if the endvertices of every T-path in G are $\leqslant_T$-comparable.

3.2.1 Tree-decompositions

A tree-decomposition $(T, \mathcal{V})$ is *rooted* if its decomposition tree T is rooted.

If $(T, \mathcal{V})$ is a tree-decomposition of a graph G, then a tree T' obtained from T by edge-contractions *induces* the tree-decomposition $(T', \mathcal{V}')$ of G whose bags are $V'_t = \bigcup_{s \in t} V_s$ for every node t of T', where we denote the vertex set of T' as the set of branch sets, that is the $\subseteq$-maximal subtrees of T consisting of contracted edges. Whenever $t \in T'$ is a subtree of T on a single vertex s, we may reference to t by s, as well. If T is a rooted tree, then $\mathrm{root}(T')$ is the node of T' containing $\mathrm{root}(T)$.

Let $(T, \mathcal{V})$ be a rooted tree-decomposition. Given an edge $e = t_0t_1$ of T with $t_0 <_T t_1$, we abbreviate the sides of its induced separation by $G{\downarrow}e := G\left[\bigcup_{t \in T_0} V_t\right]$ and $G{\uparrow}e := G\left[\bigcup_{t \in T_1} V_t\right]$. Further, we write $G\mathring{\downarrow}e := G{\downarrow}e - V_e$ and $G\mathring{\uparrow}e := G{\uparrow}e - V_e$. For a node $t \in T$, we set $G{\downarrow}t := G{\downarrow}e$, $G\mathring{\downarrow}t := G\mathring{\downarrow}e$, $G{\uparrow}t := G{\uparrow}e$ and $G\mathring{\uparrow}t := G\mathring{\uparrow}e$ where $e = st$ is the unique edge with $s <_T t$. It is easy to see that $G{\uparrow}x \supseteq G{\uparrow}y$, $G\mathring{\uparrow}x \supseteq G\mathring{\uparrow}y$, $G{\downarrow}x \subseteq G{\downarrow}y$ and $G\mathring{\downarrow}x \subseteq G\mathring{\downarrow}y$ for every two nodes or edges $x, y \in T$ with $x \leqslant_T y$, and also $\bigcap_{e \in R} G\mathring{\uparrow}e = \emptyset$ for every rooted ray R in T.

A rooted tree-decomposition $(T, \mathcal{V})$ is *componental* if $G\mathring{\uparrow}e$ is connected for every edge $e \in T$, and it is *cofinally componental* if $G\mathring{\uparrow}e$ is connected for cofinally many edges e of every $\subseteq$-maximal $\leqslant_T$-chain[9] in T. It is *tight* if, for every edge $e \in T$, there is a component C of $G\mathring{\uparrow}e$ with $N_G(C) = V_e$, and if additionally all components C of $G\mathring{\uparrow}e$ satisfy $N_G(C) = V_e$, then $(T, \mathcal{V})$ is *fully tight*. Every tight, componental, rooted tree-decomposition is fully tight. A rooted tree-decomposition $(T, \mathcal{V})$ is *linked* if for every two edges $e \leqslant_T e'$ of T, there is an edge $f \in T$ with $e \leqslant_T f \leqslant_T e'$ and a family $\{P_v \mid v \in V_f\}$ of pairwise disjoint V_e–$V_{e'}$ paths, or equivalently $(G{\downarrow}e)$ – $(G{\uparrow}e')$ paths, in G such that $v \in P_v$. In particular, the size of the family of pairwise disjoint paths equals the size of V_f. Given a set X of vertices of G, the rooted tree-decomposition $(T, \mathcal{V})$ of G is *X-linked* if $X \subseteq V_{\mathrm{root}(T)}$ and if for every edge $e \in T$ there exists an edge $f \leqslant_T e$ and a family $\{P_v \mid v \in V_f\}$ of pairwise disjoint X–V_e paths in G such that $v \in P_v$.

Lemma 3.2.1. *Let $(T, \mathcal{V})$ be a rooted tree-decomposition of a graph G and $X \subseteq V(G)$. Further let $e, f \in E(T)$, $t \in V(T)$ with $t \leqslant_T e$ and f is the unique edge in tTe incident with t. If $V_e = X \subseteq V_t$, $G\mathring{\uparrow}e$ is non-empty and $G\mathring{\uparrow}f$ is connected, then $V_s = V_t$ for every node $s \in tTe$ other than t.*

Proof. Suppose for a contradiction that there exists a node $s \in tTe$ other than t with $V_s \neq X$. Let s be a $\leqslant_T$-minimal such node. By (T2), $X \subseteq V_s$. Thus, there exists some $v \in V_s \setminus X$. Then $v \in G\mathring{\uparrow}f \setminus G\mathring{\uparrow}e$, since $V_e = X$. Moreover, $G\mathring{\uparrow}e \subseteq G\mathring{\uparrow}f$, and $G\mathring{\uparrow}f$ is connected by assumption. So there is a v–$G\mathring{\uparrow}e$ path in $G\mathring{\uparrow}f$, because $G\mathring{\uparrow}f$ is non-empty; in particular, this path avoids X, as it avoids $V_f \supseteq X$. So it also avoids $N_G(G\mathring{\uparrow}e) \subseteq X$, which is a contradiction as $v \notin G\mathring{\uparrow}e$. □

The following result shows that the finite tree-width graphs are essentially the graphs with normal spanning trees; the latter have been heavily investigated and are by now well-understood. In particular, we have Jung's *Normal Spanning Tree Criterion* [91, Satz 6'], that a connected graph G admits a normal spanning tree if its vertex set is a countable union of dispersed sets; here a set of vertices U in G is *dispersed* if for every end ε of G there is some finite set $X \subseteq V(G)$ such that U is disjoint from $C_G(X, \varepsilon)$.

Theorem 3.2.2. *A graph has finite tree-width if and only if each of its components admits a normal spanning tree.*

[9] These are the $\subseteq$-maximal rooted paths or rays.

Proof. Assume first that each component C of a given graph G admits a normal spanning tree T_C. This induces a tree-decomposition $(T_C, \mathcal{V}_C)$ of C into finite parts, given by assigning each vertex $t \in T_C$ its down-closure $\lceil t \rceil_{T_C}$ as bag. Consider the tree $T = \{r\} \sqcup \bigsqcup_C T_C$ where the new root r is adjacent to each $\operatorname{root}(T_C)$. If we assign to r the empty bag and keep all other bags, then we get a tree-decomposition of G into finite parts as desired.

Conversely, let $(T, \mathcal{V})$ be a tree-decomposition of a given graph G into finite parts. We may assume that G is connected. Since the parts of $(T, \mathcal{V})$ are finite, every end ε of G gives rise to a rooted ray $R^\varepsilon = v_0, e_0, v_1, e_1, v_2, \ldots$ in T. We claim that the union U_i of the bags assigned to nodes of the decomposition on a fixed level i is dispersed: Indeed, for every end ε of G, the finite adhesion set V_{e_i} separates U_i from $C_G(V_{e_i}, \varepsilon)$. As $V(G) = \bigcup_{i \in \mathbb{N}} U_i$, we are done by Jung's normal spanning tree criterion. □

3.2.2 Ends and tree-decompositions

Every end η of a rooted tree T contains precisely one *rooted ray* R in T, i.e. a ray in T that starts in $\operatorname{root}(T)$; we will frequently use this one-to-one correspondence between the ends of T and the rooted rays in T. Let $(T, \mathcal{V})$ be a rooted tree-decomposition $(T, \mathcal{V})$ of a graph G with finite adhesion. An end ε of G *lives in* an end η of T if some, or equivalently every, ray in ε has a tail in $G\mathring{\uparrow}e$ for every edge e of the unique rooted ray R in η; with the above relation between the rooted rays in T and the ends of T in mind, we also say that ε *gives rise to* R and R *arises from* ε. We remark that every end of G gives rise to at most one rooted ray in T, since $(T, \mathcal{V})$ has finite adhesion. If every end of G gives rise to some rooted ray of T, we encode this correspondence between the ends of G and T by a map $\varphi \colon \Omega(G) \to \Omega(T)$. In particular, this is the case if every torso of $(T, \mathcal{V})$ is rayless. If φ is a bijection between the ends of G and the ends of T, then we say that the rooted tree-decomposition $(T, \mathcal{V})$ *displays all ends* of G. If φ is a homeomorphism from $\Omega(G)$ to $\Omega(T)$, then $(T, \mathcal{V})$ *homeomorphically displays all ends* of G.

For a sequence $(V_i)_{i \in \mathbb{N}}$ of sets, we set $\liminf_{i \in \mathbb{N}} V_i := \bigcap_{i \in \mathbb{N}} \bigcup_{j \geqslant i} V_i$. If the sequence of sets V_{e_i} is indexed by the edges of a ray $R = v_0 e_0 v_1 e_1 \ldots$ in a graph G, then we also write $\liminf_{e \in R} V_e$ for $\liminf_{i \in \mathbb{N}} V_{e_i}$. If $(T, \mathcal{V})$ displays the ends of G and additionally has the property that the unique rooted ray R in T which arises from the end ε of G satisfies $\liminf_{e \in R} |V_e| = \Delta_G(\varepsilon)$, then the tree-decomposition $(T, \mathcal{V})$ of G *displays the combined degrees of every end*. If a rooted tree-decomposition $(T, \mathcal{V})$ of G displays all ends of G and additionally $\liminf_{e \in R} V_e = \operatorname{Dom}(\varepsilon)$ for every rooted ray R of T and its arising end ε of G, then $(T, \mathcal{V})$ *displays the dominating vertices of every end.*

Let X and Y be two sets of vertices in a graph G. Then X is *linked to* Y in G if there is a family $\{R_x \mid x \in X\}$ of pairwise disjoint X–Y paths in G such that $x \in R_x$. Let ε be an end of G. An *X–ε ray* in G is a ray which starts in X and is contained in ε. An *X–ε path* in G is an X–$\mathrm{Dom}(\varepsilon)$ path in G. A finite set $S \subseteq V(G)$ is an X–ε *separator* in G if $C_G(S,\varepsilon)$ does not meet X. The set X is *linked to* an end ε of G if there is a family $\{R_x \mid x \in X\}$ of pairwise disjoint X–ε paths and rays such that $x \in R_x$. Further, a rooted tree-decomposition $(T,\mathcal{V})$ is *end-linked* if for every edge e of T there exists some end ε of G which lives in $G{\uparrow}e$ and to which V_e is linked. The following lemma is clear.

Lemma 3.2.3. *Every end-linked, rooted tree-decomposition* $(T,\mathcal{V})$ *is tight.* □

3.2.3 Critical vertex sets and tree-decompositions

A tree-decomposition $(T,\mathcal{V})$ of a graph G *displays the critical vertex sets* if the map $t \mapsto V_t$ restricted to the infinite-degree nodes of T whose V_t is finite is a bijection to the critical vertex sets of G. In this case, for a critical vertex set X, we denote by t_X the unique infinite-degree node t with $V_t = X$. If a rooted tree-decomposition $(T,\mathcal{V})$ not only displays the critical vertex sets but also for every critical vertex set X cofinitely many tight components of $G - X$ are $G\mathring{\uparrow}e$ for some $e = t_X t \in T$ with $t_X <_T t$ and every such $G\mathring{\uparrow}e$ is a tight component of $G - X$, then $(T,\mathcal{V})$ *displays the tight components of every critical vertex set cofinitely.* We remark that if such a $(T,\mathcal{V})$ is tight, then for every finite proper subset Y of any (possibly infinite) bag V_t there are at most finitely many edges $e = ts \in T$ with $t <_T s$ such that $V_e = Y$. A rooted tree-decomposition $(T,\mathcal{V})$ of G into finite parts *displays the infinities* of G, if it displays the ends of G homeomorphically, their combined degrees, their dominating vertices, the critical vertex sets and their tight components cofinitely.

As two vertices in a critical vertex set can always be joined by a path which avoids any given finite set of other vertices, a greedy argument yields the following.

Lemma 3.2.4. *Assume that for a critical vertex set X of a graph G we have two path families $\mathcal{P},\mathcal{Q}$ of $k \in \mathbb{N}$ disjoint Y–X paths and of k disjoint X–Z paths, respectively, for some $Y, Z \subseteq V(G)$, such that the paths in $\mathcal{P} \cup \mathcal{Q}$ are disjoint outside of X. Then there exists a family of k disjoint Y–Z paths in G.* □

3.2.4 Critical vertex sets of torsos

The following lemma, which we will use in the proofs of Theorems 3 and 4, says that the critical vertex sets of a torso are closely related to the critical vertex sets of the underlying graph; in particular, torsos in tough graphs are tough.

Lemma 3.2.5. *Let σ be a star of finite-order separations of a graph G such that for cofinitely many separations $(A,B) \in \sigma$ the side A contains a tight component of $G - (A \cap B)$. Then* $\mathrm{crit}(\mathrm{torso}(\sigma)) \subseteq \mathrm{crit}(G)$. *Moreover,*

(i) *If* $X \in \mathrm{crit}(\mathrm{torso}(\sigma))$, *then there are infinitely many tight components of* $G - X$ *which meet* $\mathrm{torso}(\sigma)$.

(ii) *If* $X \subseteq \mathrm{int}(\sigma)$, *then the set* $V(C) \cap \mathrm{int}(\sigma)$ *induces a tight component of* $\mathrm{torso}(\sigma) - X$ *for cofinitely many* $C \in \check{\mathcal{C}}(G - X)$ *which meet* $\mathrm{int}(\sigma)$.

Proof. We remark that (i) immediately yields that $\mathrm{crit}(\mathrm{torso}(\sigma)) \subseteq \mathrm{crit}(G)$. Let U be the union of the finite separators of those finitely many $(A, B) \in \sigma$ whose side A does not contain a tight component of $G - (A \cap B)$; in particular, U is finite.

(i): Since U is finite, only finitely many components of $\mathrm{torso}(\sigma) - X$ meet U. For every component C' of $\mathrm{torso}(\sigma) - X$ which avoids U, the subgraph

$$C := C' \cup \bigcup \{G[A] \colon (A, B) \in \sigma,\ V(C') \cap A \neq \emptyset\}$$

is connected by the definition of U. Moreover, since the separators $A \cap B$ of separations $(A, B) \in \sigma$ are complete in $\mathrm{torso}(\sigma)$, the component C' contains the entire $(A \cap B) \setminus X$ as soon as it meets $A \cap B$. Hence, C is a component of $G - X$, and by definition it contains no other components of $\mathrm{torso}(\sigma) - X$ than C'. It thus suffices to show for all infinitely many components C' of $\mathrm{torso}(\sigma) - X$ that avoid U that the component C of $G - X$ which contains C' is tight. For this, it suffices to prove that whenever there is a torso edge from $u' \in C'$ to $v \in X$, then there also is some edge from C to v in G. By the definition of torso, there is a separation $(A, B) \in \sigma$ with $u', v \in A \cap B$. The side A of (A, B) contains a tight component K of $G - (A \cap B)$ as C' avoids U; in particular, $K \subseteq C$ by the definition of C, and K sends an edge to v in G.

(ii): Let $X \subseteq \mathrm{int}(\sigma)$. It suffices to show that $C' := C \cap \mathrm{int}(\sigma)$ is a tight component of $\mathrm{torso}(\sigma) - X$ for every tight component C of $G - X$ which meets $\mathrm{int}(\sigma)$ but avoids the finite set U. The definition of torso immediately yields that C' induces a connected subgraph of $\mathrm{torso}(\sigma) - X$ with $N_{\mathrm{torso}(\sigma)}(C') \supseteq X$ because $X \subseteq \mathrm{int}(\sigma)$. It remains to show that $N_{\mathrm{torso}(\sigma)}(C') \subseteq X$. For this it suffices to prove that whenever there is a torso edge from $u' \in C'$ to $v \notin C'$ the vertex v is already in X. Let $(A, B) \in \sigma$ such that $u', v \in A \cap B$. As C avoids U, the component C, or equivalently C', only meets separators of separations $(A, B) \in \sigma$ whose side A contains a tight component of $G - (A \cap B)$. Thus, we have $(A \cap B) \setminus X \subseteq A \setminus X \subseteq C$ as soon as C meets $A \cap B$. Since $v \in (A \cap B)$ but not in C, we thus have $v \in X$, as desired. □

3.3 Displaying ends

In this short section we show that Theorem 2 follows from Theorem 1, that is, we show that a linked, tight, componental, rooted tree-decomposition into finite parts homeomorphically displays all ends, their dominating vertices and their combined degrees.

The proof is divided into three lemmas. The first follows immediately from [93, Lemma 3.1] applied to the tree-decomposition induced by contracting all edges which violate componental.

Lemma 3.3.1. *Let $(T, \mathcal{V})$ be a cofinally componental*[10]*, rooted tree-decomposition of a graph G which has finite adhesion. Then every rooted ray R of T arises from precisely one end of G. Moreover, if all torsos of $(T, \mathcal{V})$ are rayless, then $(T, \mathcal{V})$ displays all ends of G homeomorphically.* □

Lemma 3.3.2. *Let $(T, \mathcal{V})$ be a tight, componental, rooted tree-decomposition of a graph G which has finite adhesion. Then $\liminf_{e \in R} V_e = \mathrm{Dom}(\varepsilon)$ for every rooted R of T and the unique end ε of G which gives rise to R.*

Proof. Since $(T, \mathcal{V})$ has finite adhesion, $\liminf_{e \in R} V_e$ contains $\mathrm{Dom}(\varepsilon)$. Conversely, let v be a vertex in $\liminf_{e \in R} V_e$. Let $Q \in \varepsilon$ be arbitrary. We aim to construct an infinite v–Q fan in G, which then shows that $v \in \mathrm{Dom}(\varepsilon)$. Note that we may recursively find infinitely many pairwise internally disjoint v–Q paths in G if for each finite set $X \subseteq V(G) \setminus \{v\}$ there is a v–Q path in G avoiding X. As $(T, \mathcal{V})$ is a tree-decomposition, there is an edge $e \in R$ for which $G \mathring{\uparrow} e$ avoids any given finite set $X \subseteq V(G)$. Since ε gives rise to R, the ray Q has a tail in $G \mathring{\uparrow} e$. Thus, we find the desired v–Q path in the connected subgraph $G \uparrow e$, as $(T, \mathcal{V})$ is tight and componental. □

Lemma 3.3.3. *Let $(T, \mathcal{V})$ be a linked, rooted tree-decomposition of a graph G which has finite adhesion. Suppose that an end ε of G gives rise to a ray R in T which arises from no other end of G and that $\liminf_{e \in R} V_e = \mathrm{Dom}(\varepsilon)$. Then $\liminf_{e \in R} |V_e| = \Delta(\varepsilon)$.*

In particular, if $(T, \mathcal{V})$ displays all ends of G and their dominating vertices, then $(T, \mathcal{V})$ also displays the combined degree of each end of G.

Proof. Let $R = v_0 e_0 v_1 e_1 \dots$ be the unique rooted ray in T which arises from the end ε. If $\mathrm{dom}(\varepsilon)$ is infinite, then $(T, \mathcal{V})$ displays the combined degree of ε by assumption. Thus, we may assume $\mathrm{dom}(\varepsilon)$ to be finite. Moving to a tail of R, we may assume that $\mathrm{Dom}(\varepsilon) \subseteq V_e$ for every $e \in R$.

From the sequence $(V_{e_i})_{i \in \mathbb{N}}$ of adhesion sets along R, we extract a subsequence by letting $i_0 \in \mathbb{N}$ with $|V_{e_{i_0}}| = \min_{e \in R} |V_e|$ and recursively choosing $i_{n+1} \in \mathbb{N}$ with $i_{n+1} > i_n$ and $|V_{e_{i_{n+1}}}| = \min_{j > i_n} |V_{e_j}|$ for $n \in \mathbb{N}$. Write $S_n := V_{e_{i_n}}$. Then $(|S_n|)_{n \in \mathbb{N}}$ is non-decreasing, satisfies $\liminf_{e \in R} |V_e| = \liminf_{n \in \mathbb{N}} |S_n|$ and $\liminf_{n \in \mathbb{N}} S_n = \liminf_{e \in R} V_e = \mathrm{Dom}(\varepsilon)$ by definition.

Since $\liminf_{e \in R} S_n = \mathrm{Dom}(\varepsilon)$, we may assume that $S_n \cap S_m \subseteq \mathrm{Dom}(\varepsilon)$ for all $m \geqslant n$ by passing to a further subsequence. We also have $C_G(S_m, \varepsilon) \subseteq C_G(S_n, \varepsilon)$ since ε gives rise to R. Moreover, $\bigcap_{n \in \mathbb{N}} C_G(S_n, \varepsilon) = \emptyset$, since ε gives rise to R and $(T, \mathcal{V})$ is a tree-decomposition. This implies by [77, Corollary 5.7] that $\liminf_{e \in R} |V_e| = \liminf_{n \in \mathbb{N}} |S_n| \geqslant \Delta_G(\varepsilon)$.

For $\liminf_{e \in R} |V_e| = \liminf_{n \in \mathbb{N}} |S_n| \leqslant \Delta_G(\varepsilon)$, we use that $(T, \mathcal{V})$ is linked: By the construction of the sequence $(S_n)_{n \in \mathbb{N}}$, the linkedness yields a family $\mathcal{P}_n$ of $|S_n|$ disjoint S_n–S_{n+1} paths in G for

[10] Our definition of *componental* agrees with the definition of *upwards connected* from [93].

every $n \in \mathbb{N}$. Then $\mathcal{R} = \bigcup_{n\in\mathbb{N}} \mathcal{P}_n$ is a family of pairwise disjoint paths and rays; the rays in $\mathcal{R}$ are necessarily ε-rays since ε is the unique end of G that gives rise to R, and the finite paths in $\mathcal{R}$ necessarily have a vertex of $\mathrm{Dom}(\varepsilon)$ as their last vertex. Hence, $|\mathcal{R}| \leqslant \Delta(\varepsilon)$. Moreover, we clearly have $|\mathcal{R}| = \liminf_{n\in\mathbb{N}} |S_n|$, which implies that $\liminf_{n\in\mathbb{N}} |S_n| = |\mathcal{R}| \leqslant \Delta_G(\varepsilon)$, as desired. □

Proof of Theorem 2 given Theorem 1. By Lemmas 3.3.1 to 3.3.3 the tree-decomposition from Theorem 1 is as desired. □

We also have the following lemma, which follows immediately from Lemmas 3.3.1 to 3.3.3, and which we will need in the next chapter.

Lemma 3.3.4. *Let $(T, \mathcal{V})$ be a rooted tree-decomposition of a graph G which has finite adhesion and which is linked, tight and componental. Then $(T, \mathcal{V})$ displays all ends of G, their dominating vertices and their combined degrees.* □

3.4 A high-level proof of Theorem 1

3.4.1 A two-step approach to Theorem 1

By Theorem 2.5.2, every tough, rayless graph is finite. Hence, to arrive at a tree-decomposition into finite parts, we first construct a tree-decomposition $(T, \mathcal{V})$ whose torsos are tough. Next, for each torso of $(T, \mathcal{V})$ corresponding to a node $t \in T$ we construct another tree-decomposition $(T^t, \mathcal{V}^t)$ with rayless torsos. By refining $(T, \mathcal{V})$ with all the $(T^t, \mathcal{V}^t)$, we obtain a tree-decomposition $(T', \mathcal{V}')$ into tough and rayless parts, which then must be finite by Theorem 2.5.2.

But how to guarantee that the resulting tree-decomposition is linked? Lemma 3.2.4 ensures that if we begin in the above two-step approach with a tree-decomposition $(T, \mathcal{V})$ whose adhesion sets are all critical vertex sets, and then refine by tree-decompositions $(T^t, \mathcal{V}^t)$ that are each linked when considered individually, then the arising combined tree-decomposition is again linked.[11]

A little more formally, the first step in our two-step approach is the following theorem:

Theorem 6. *Every graph of finite tree-width admits a tight, cofinally componental, rooted tree-decomposition whose adhesion sets are critical vertex sets, whose torsos are tough and which displays the critical vertex sets and their tight components cofinitely.*

The second step in our two-step approach is formally given by the following theorem:

[11]This is not completely true: In order to ensure that the arising tree-decomposition is linked we need a 'refinement version' of Theorem 7 (see Theorem 7') that ensures that the adhesion sets of $(T, \mathcal{V})$ corresponding to edges incident with a given node $t \in T$ appear as adhesion sets in $(T^t, \mathcal{V}^t)$. But this is only a technical issue which does not add much complexity to the proof of Theorem 7.

Theorem 7. *Every graph of finite tree-width admits a linked, tight, componental, rooted tree-decomposition of finite adhesion whose torsos are rayless.*

Our aim in the remainder of this section is to demonstrate formally how Theorem 6 and Theorem 7 can be combined to yield a proof of Theorem 1. The subsequent Sections 3.5 and 3.7 are then concerned with proving Theorems 6 and 7 respectively.

3.4.2 Detailed versions of the main theorems

In fact, we will prove a version of Theorem 1 not as stated in the introduction, but with three additional properties that are important for technical reasons, but which we believe are also of interest in their own right.

Theorem 1' (Detailed version of Theorem 1)**.** *Every graph G of finite tree-width admits a rooted tree-decomposition $(T, \mathcal{V})$ into finite parts that is linked, tight, and componental. In particular, $(T, \mathcal{V})$ displays all the ends of G homeomorphically, their combined degrees and their dominating vertices. Moreover, we may assume that*

(L1) *for every $e \in E(T)$, the adhesion set V_e is either linked to a critical vertex set of G that is included in $G{\uparrow}e$ or linked to an end of G that lives in $G{\uparrow}e$,*

(L2) *for every $e <_T e' \in E(T)$ with $|V_e| \leqslant |V_{e'}|$, each vertex of $V_e \cap V_{e'}$ either dominates some end of G that lives in $G{\uparrow}e'$, or is contained in a critical vertex set of G that is included in $G{\uparrow}e'$,*

(L3) *the bags of $(T, \mathcal{V})$ are pairwise distinct.*

Before we continue, let us quickly comment on these additional properties: Property (L1) is a minimality condition: Since every end of G lives in an end of $(T, \mathcal{V})$, there will be infinitely many edges of T whose adhesion set is linked to that end. Moreover, since all parts of $(T, \mathcal{V})$ are finite, one can easily see that every critical vertex set will appear as some adhesion set of $(T, \mathcal{V})$ (cf. Lemma 3.5.1). So (L1) says that we did not decompose G 'more than necessary'. In particular, every bag at a leaf of $(T, \mathcal{V})$ will be of the form $X \cup V(C)$ for a critical vertex set X of G and a finite tight component C of $G - X$.

Let us now turn to property (L2). Recall that Halin [82, Theorem 2] showed that every locally finite connected graph has a linked ray-decomposition into finite parts with disjoint adhesion sets. In light of this, (L2) describes how close we can come to having 'disjoint adhesion sets' in the general case; Example 9 in Chapter 4 shows that even for locally finite graphs, it may be impossible to get a tree-decomposition with disjoint adhesion sets, so the condition 'with $|V_e| \leqslant |V_{e'}|$' in (L2) is indeed necessary.

As already indicated in the introduction, we first show Theorem 4 and then derive Theorem 1 from it. In fact, we derive the detailed version Theorem 1' from the following detailed version of Theorem 4.

Theorem 4' (Detailed version of Theorem 4)**.** *Every graph G of finite tree-width admits a fully tight, cofinally componental, linked, rooted tree-decomposition into finite parts which displays the infinities of G and which satisfies* (L1) *and* (L2) *from Theorem 1'. Moreover,*

(I1) *if $G\mathring{\uparrow}e$ is disconnected for $e = st \in E(T)$ with $s <_T t$, then* $V_s \supseteq V_t \in \mathrm{crit}(G)$ *and* $\deg(t) = \infty$,

(I2) *if $V_t = V_e$ for some node $t \in T$ and the unique edge $e = st \in T$ with $s <_T t$, then* $\deg(t) = \infty$ *and* $V_t \in \mathrm{crit}(G)$.

We remark that whenever a tree-decomposition $(T, \mathcal{V})$ displays the critical vertex sets and satisfies (I1), then it is automatically cofinally componental, because if two successive comparable edges $rs, st \in T$ violate componental, then (I1) yields that $\deg(t), \deg(s) = \infty$ and $V_s \supseteq V_t$, which implies that $V_s \supsetneq V_t$, since $(T, \mathcal{V})$ displays the critical vertex sets and hence $V_s \neq V_t$. Moreover, every tree-decomposition as in Theorem 4' already 'nearly' satisfies (L3), that is, has 'almost' distinct bags.

Lemma 3.4.1. *Let $(T, \mathcal{V})$ be a tight, rooted tree-decomposition of a graph G which displays the critical vertex sets of G and satisfies* (I1) *and* (I2)*. Then the bags of $(T, \mathcal{V})$ are pairwise distinct, unless they are critical vertex sets, which may appear as at most two bags associated with adjacent nodes of T.*

Proof. Suppose there are distinct nodes $t, s \in T$ such that $V_t = V_s$. Then tTs contains from at least one of t, s, say from t, its unique down-edge e. Then (T2) ensures that $V_e = V_t$, so by (I2) we have $V_t \in \mathrm{crit}(G)$ and $\deg(t) = \infty$. In particular, since $(T, \mathcal{V})$ displays the critical vertex sets of G, we have $\deg(s) \neq \infty$ and hence $s <_T t$, again by (I2). Now if $V_{s'} = V_s$ for the unique neighbour of s in sTt, then $\deg(s') = \infty$ by (I2), so $s' = t$ because $(T, \mathcal{V})$ displays the critical vertex sets of G. Hence, we may assume that $V_{s'} \neq V_s$, so $V_{s'} \supsetneq V_s$ by (T2). Then (I1) implies that $G\mathring{\uparrow}f$ is connected. But then Lemma 3.2.1 implies that $V_{s'} = V_t = V_s$, a contradiction. □

Proof of Theorem 1' given Theorem 4'. Let $(T', \mathcal{V}')$ be the tree-decomposition obtained from Theorem 4'. Let $(T, \mathcal{V})$ be the tree-decomposition induced by contracting every edge e of T' whose $G\mathring{\uparrow}e$ is disconnected. It is immediate from the construction that $(T, \mathcal{V})$ is componental. Since $(T', \mathcal{V}')$ is fully tight and satisfies (L1) and (L2) from Theorem 1', so does every tree-decomposition induced by edge-contractions from $(T', \mathcal{V}')$. We note that whenever an edge $e = st \in E(T')$ with $s <_{T'} t$ has been contracted then no other edge f incident with t has been contracted and for $V_f = V_t = V_e$ for all such edges $f \in T'$ incident with t, because $(T', \mathcal{V}')$ displays the critical vertex sets and their tight components cofinitely and it also satisfies (I1). Hence, such as $(T', \mathcal{V}')$ does, the induced tree-decomposition $(T, \mathcal{V})$ still is linked, displays the ends of G homeomorphically, their combined degrees and their dominating vertices, and also its parts are finite.

Moreover, $(T, \mathcal{V})$ satisfies (L3). Indeed, by Lemma 3.4.1, we only need to check that every critical vertex set appears as at most one bag of $(T, \mathcal{V})$. By Lemma 3.4.1, every $X \in \operatorname{crit}(G)$ can appear as at most two bags of $(T', \mathcal{V}')$, which then need to be adjacent. So assume $V_s = V_t = X$ with s being the unique down-neighbour of t. Then by (I2), $\deg(t) = \infty$. Since $(T', \mathcal{V}')$ displays the critical vertex sets and their tight components cofinitely, it follows that $G\mathring{\uparrow}(st)$ is disconnected. Thus, we have contracted the edge st in the construction of $(T, \mathcal{V})$. Hence, also every critical vertex set of G appears as at most one bag of $(T, \mathcal{V})$, so its bags are paiwise distinct. □

In order to prove Theorem 4', we still follow the promised two-step approach, but need the following detailed versions of Theorems 6 and 7.

Theorem 6' (Detailed version of Theorem 6). *Let G be a graph of finite tree-width. Then G admits a fully tight, cofinally componental, rooted tree-decomposition $(T, \mathcal{V})$ whose adhesion sets are critical vertex sets, whose torsos are tough and which displays the critical vertex sets and their tight components cofinitely.*

Moreover, it satisfies (I1) *and* (I2) *from Theorem 4'.*

In order to state the detailed version of Theorem 7, we need one more definition. Recall that a separation (A, B) of a graph G is *left-tight* if some components C of $G[A \setminus B]$ satisfies $N_G(C) = A \cap B$. Moreover, a separation (A, B) of a graph G is *left-fully-tight* if all components C of $G[A \setminus B]$ satisfy $N_G(C) = A \cap B$.

Theorem 7' (Detailed version of Theorem 7). *Let G be a graph, and let σ be a star of left-well-linked left-fully-tight finite-order separations of G such that* $\operatorname{torso}(\sigma)$ *has finite tree-width. Further, let $X \subseteq \operatorname{int}(\sigma)$ be a finite set of vertices of G. Then G admits a linked, X-linked, fully tight, rooted tree-decomposition $(T, \mathcal{V})$ of finite adhesion such that*

- (R1) *its torsos at non-leaves are rayless and its leaf separations are precisely $\{(B, A) \mid (A, B) \in \sigma\}$,*
- (R2) *for all edges e of T, the adhesion set V_e is either linked to an end living in $G\uparrow e$ or linked to a set $A \cap B \subseteq G\uparrow e$ with $(A, B) \in \sigma$,*
- (R3) *for every $e <_T e' \in E(T)$ with $|V_e| \leqslant |V_{e'}|$, either each vertex of $V_e \cap V_{e'}$ dominates some end of G that lives in $G\uparrow e'$, or $V_e \cap V_{e'}$ is contained in $A \cap B \subseteq G\uparrow e'$ for some $(A, B) \in \sigma$,*
- (R4) *$V_s \supsetneq V_e \subsetneq V_t$ for all edges $e = st \in T$ with $s <_T t$ and $s \neq r :=$ $\operatorname{root}(T)$. Moreover, if $X \subsetneq \operatorname{int}(\sigma)$, $G - X$ is connected and $N_G(G - X) = X$, then $X \subsetneq V_r$ and also $V_r \supsetneq V_e \subsetneq V_t$ for all edges $e = rt \in T$.*

Note that (R4) implies that the bags of $(T, \mathcal{V})$ are pairwise distinct.

3.4.3 Proof of the main result

According to our two-step approach, we prove Theorem 4' by applying Theorem 7' to the torsos of the tree-decomposition given by Theorem 6'. For this, we need to ensure that all torsos again have finite tree-width:

Lemma 3.4.2. *Let G be a graph of finite tree-width, and let G' be obtained from G by adding an edge between $u, v \in V(G)$ whenever there are infinitely many internally disjoint u–v paths in G. Then every tree-decomposition of G of finite adhesion is also a tree-decomposition of G'.*

In particular, if G has finite tree-width, then the torso at a star of separations of G whose separators are critical vertex sets in G has finite tree-width.

Proof. Assume that $(T, \mathcal{V})$ is a tree-decomposition of G of finite adhesion, and consider any two vertices $u, v \in V(G)$ with $uv \in E(G') \setminus E(G)$. If there exists a bag V_t containing both u and v, then the edge uv in G' cannot violate that $(T, \mathcal{V})$ is a tree-decomposition of G'. To find such a bag V_t, recall that there are infinitely many internally disjoint u–v paths in G. In particular, no finite set of vertices other than u and v separates u and v in G. Since $(T, \mathcal{V})$ has finite adhesion, u and v must be contained in some bag V_t of $(T, \mathcal{V})$, as desired.

For the 'in particular'-part assume that G has finite tree-width and that σ is a star of separations of G whose separators are critical in G. Let $(T, \mathcal{V})$ be a tree-decomposition of G into finite parts. Then by the first part and because critical vertex sets are infinitely connected, $(T, \mathcal{V}')$ with $V'_t := V_t \cap \operatorname{int}(\sigma)$ is a tree-decomposition of the torso at σ in G, as desired. □

Proof of Theorem 4' given Theorem 6' and Theorem 7'. Let $(T^1, \mathcal{V}^1)$ be the rooted tree-decomposition from Theorem 6' whose adhesion sets are critical vertex sets. In particular, $(T^1, \mathcal{V}^1)$ displays the critical vertex sets of G and their tight components cofinitely. Moreover, its torsos are tough and it satisfies (I1) from Theorem 4'. Let t be a node of T^1. We describe how we refine the torso at t in T^1 using Theorem 7':

If t is the root of T^1, then we set $G^t := G$, $X_t := \emptyset$ and $\sigma'_t = \sigma_t$. Else let $s \in T^1$ be the (unique) predecessor of t and let G^t be the graph obtained from $G{\uparrow}st$ by adding all edges between vertices of V^1_{st}. Set $X_t := V^1_{st}$ and $\sigma'_t := \{(A, B \setminus V(G\mathring{\downarrow}st)) \mid (A_{st}, B_{st}) \neq (A, B) \in \sigma_t\}$, where (A_{st}, B_{st}) is the separation induced by (s, t).

First, we assume that t is a node of T^1 with $V^1_t \in \operatorname{crit}(G)$; in particular, all infinite-degree nodes whose corresponding bag is finite are such t, since $(T^1, \mathcal{V}^1)$ displays the critical vertex sets. Then we set $(T^t, \mathcal{V}^t)$ to be the tree-decomposition of G^t whose decomposition tree is a star rooted in its centre t and with bag $V^t_t := V^1_t$ while its leaf separations are precisely $((B, A) \mid (A, B) \in \sigma'_t)$. Note that all adhesion sets of this tree-decomposition $(T^t, \mathcal{V}^t)$ are $V^1_t \in \operatorname{crit}(G)$. Thus, $(T^t, \mathcal{V}^t)$ is a rooted tree-decomposition as in the conclusion of Theorem 7'.

Secondly, we now assume that t is a node of T^1 with $V^1_t \notin \text{crit}(G)$. By construction, the torso of σ'_t in G^t is equal to the torso of σ_t in G; so by Lemma 3.4.2, this torso has finite tree-width. We claim that G^t, X_t and σ'_t are as required to apply Theorem 7'. For this it suffices that all separations in σ'_t are left-well-linked. Let $(A,B) \in \sigma'_t$. Then $X =: A \cap B$ is some critical vertex set of G, as it is an adhesion set of $(T^1, \mathcal{V}^1)$. Since $(T^1, \mathcal{V}^1)$ displays the critical vertex sets, there is a unique infinite-degree node $t_X \in T^1$ with $V^1_{t_X} = X$. If infinitely many tight component of $G - X$ are contained in A, then (A,B) is left-well-linked. Thus, it now suffices to show that $t <_{T^1} t_X$, as $(T^1, \mathcal{V}^1)$ cofinitely displays also the tight components of the critical vertex sets. Because $V_t \notin \text{crit}(G)$, $t \neq t_X$; hence, we suppose towards a contradiction that $t >_{T^1} t_X$. Now $G[A \setminus B]$ is in particular non-empty, as $(T^1, \mathcal{V}^1)$ is fully tight. Also $G\mathring{\uparrow} f$ is a (tight) component of $G - X$ where f is the unique edge on $t_X T^1 t$ incident with t, since $(T^1, \mathcal{V}^1)$ cofinitely displays the tight components of the the critical vertex sets. But then Lemma 3.2.1 yields $V_t = X \in \text{crit}(G)$ which contradicts the assumptions on t. All in all, we may now apply Theorem 7' to obtain the rooted tree-decomposition $(T^t, \mathcal{V}^t)$ of G^t.

We remark that all these rooted tree-decompositions $(T^t, \mathcal{V}^t)$ in particular contain X_t in its root part and have precisely $((B,A) \mid (A,B) \in \sigma'_t)$ as its leaf separations. To build the desired tree-decomposition of G, we first stick all these tree-decompositions $(T^t, \mathcal{V}^t)$ together along $(T^1, \mathcal{V}^1)$: Formally, the tree T^2 arises from the disjoint union of the trees T^t by identifying a leaf u of T^t with the root of T^s if $G\mathring{\uparrow} u$ (with respect to T^t) is equal to $G^s - X_s$; we say that the edge of T^t (and hence of T^2) incident with the leaf u *belongs to* T^1 and that it *corresponds to ts*. All edges of T^t that do not belong to T^1 are said to *belong to* T^t. We remark that every edge belongs to precisely one of the T^t or T^1. We set the root of T^2 to be $\text{root}(T^{\text{root}(T^1)})$. For each node $s \in T^2$, we then set V^2_s to be V^t_s, where t is the (unique) node of T^1 such that s is either a non-leaf of T^t or the unique node of T^t. We say that s *belongs to* T^t. Now we claim that $(T^2, \mathcal{V}^2)$ has all the desired properties.

Let us first note that the construction of T^2 immediately ensures that

- if $e \in T^2$ belongs to T^t for some $t \in T^1$, then $G\mathring{\uparrow} e = G^t \mathring{\uparrow} e$ and $V^2_e = V^t_e$, and
- if $e \in T^2$ belongs to T^1, then $G\mathring{\uparrow} e = G\mathring{\uparrow} f$ and $V^2_e = V^1_f$ for the edge $f \in T^1$ which e corresponds to.

Thus, $(T^2, \mathcal{V}^2)$ is fully tight, since $(T^1, \mathcal{V}^1)$ and all the $(T^t, \mathcal{V}^t)$ are fully tight. The tree-decomposition $(T^2, \mathcal{V}^2)$ satisfies (I1) from Theorem 4', since $(T^1, \mathcal{V}^1)$ and all the $(T^t, \mathcal{V}^t)$ satisfy (I1). It also satisfies (I2). Indeed, by Theorem 6', $(T^1, \mathcal{V}^1)$ satisfies (I2). Moreover, by (R4) from Theorem 7', the $(T^t, \mathcal{V}^t)$ have the property that $V^t_x \supsetneq V^t_e \subsetneq V^t_y$ for all edges $e = xy \in T^t$ with $x <_{T^t} y$ and also $V^1_e \subsetneq V^s_{\text{root}(T^s)}$ where $e = st \in T^1$ with $s <_{T^1} t$. We remark that we used here that every node t and its unique edge $e = st \in T^1$ with $s <_{T^1} t$ to which we applied Theorem 7' satisfies $X_t = V^1_{st} \subsetneq V^1_t = \text{int}(\sigma) = \text{int}(\sigma'_t)$ by (I2) from Theorem 6'' and also $G^t - X_t$ is connected with

$N_G^t(G^t - X) = X$ by (I1) and since $(T^1, \mathcal{V}^1)$ is fully tight. It follows that $V_x^2 \supsetneq V_e^2 \subsetneq V_y^2$ for all edges $e = xy$ with $x <_{T^2} y$ except those where $V_x^2 = V_x^1 = V_y^1 = V_y^2 \in \text{crit}(G)$. In particular, $(T^2, \mathcal{V}^2)$ satisfies (I2).

Let us now show that all parts of $(T^2, \mathcal{V}^2)$ are finite. By Theorem 2.5.2, it suffices to show that the torso at every $s \in T^2$ is rayless and tough. Let t be the node of T^1 to whose T^t the node s belongs. It is immediate from the construction of $(T^2, \mathcal{V}^2)$ that the torso G_s^2 at $s \in T^2$ in $(T^2, \mathcal{V}^2)$ is equal to the torso G_s^t at s in $(T^t, \mathcal{V}^t)$; in particular, these torsos are rayless by (R1). Suppose for a contradiction that the torso G_s^2 at $s \in T^2$ in $(T^2, \mathcal{V}^2)$ contains a critical vertex set X. By Lemma 3.2.5 (i), infinitely many tight components of $G - X$ meet the torso G_s^2; in particular, they meet the torso G_t^1 at $t \in T^1$ in $(T^1, \mathcal{V}^1)$. Now by Lemma 3.2.5 (ii), cofinitely many of these tight components of $G - X$ restrict to tight components of the torso G_t^1 at t in $(T^1, \mathcal{V}^1)$. Thus, X is a critical vertex set of the tough torso G_t^1 which is a contradiction.

Since the adhesion sets of $(T^1, \mathcal{V}^1)$ are critical vertex sets of G, (L1) and (L2) follow immediately from (R2) and (R3), respectively.

It remains to show that $(T^2, \mathcal{V}^2)$ is linked. So let $e \leqslant_{T^2} e'$ be given. Suppose first that there is a node t of T^1 such that all edges in eT^2e' belong either to T^t or correspond to edges of T^1 incident with t. The tree-decomposition $(T^t, \mathcal{V}^t)$ obtained from Theorem 7' is linked and X_t-linked, where $X_t = V_{st}^1$ for the predecessor s of t in T^1. Thus, there exists an edge $f \in eT^te'$ with $e \leqslant_{T^t} f \leqslant_{T^t} e'$ and a family $\mathcal{P}$ of $k := |V_f^t|$ disjoint V_e^t–$V_{e'}^t$ paths in the auxiliary graph G^t such that $v \in P_v$. As $G^t \subseteq G$, these paths are also paths in G. Since the adhesion sets of T^t and the tree-order in T^t directly transfer to T^2 by construction, this completes the first case.

To conclude the proof that $(T^2, \mathcal{V}^2)$ is linked, let k be the minimum size of an adhesion set V_g among all edges $g \in eT^2e'$. Further, let $f_1, \dots, f_\ell$ be the edges on the path eT^2e' that belong to T^1 ordered by $\leqslant_{T^2}$. To find k disjoint V_e^2–$V_{e'}^2$ paths in G, we apply the above argument to each subpath $f_iT^2f_{i+1}$ with $i \in \{1, \dots, n-1\}$. By the choice of k, we get a family $\mathcal{P}_i$ of k disjoint $V_{f_i}^2$–$V_{f_{i+1}}^2$ paths in G for every $i \in \{1, \dots, n-1\}$. As $(T^2, \mathcal{V}^2)$ is a tree-decomposition, we have that for every $g_1 \leqslant_{T^2} g_2 \leqslant_{T^2} g_3$, V_{g_2} separates $V_{g_1}^2$ and $V_{g_3}^2$. Therefore, $P_i \in \mathcal{P}_i$ and $P_j \in \mathcal{P}_j$ are internally disjoint for $i \neq j$. We remark that $V_{f_{i+1}}^2$ is an adhesion set of $(T^1, \mathcal{V}^1)$ and thus critical in G. Hence, Lemma 3.2.4 yields the desired path family.

In particular, $(T^2, \mathcal{V}^2)$ displays all the ends of G homeomorphically, their dominating vertices and their combined degrees by Lemmas 3.3.1 to 3.3.3. It is immediate from the construction that $(T^2, \mathcal{V}^2)$ still displays the critical vertices and their tight components cofinitely. □

3.5 Tree-decomposition along critical vertex sets

In this section we prove Theorem 6', which we restate here for convenience.

Theorem 6'' (Detailed version of Theorem 6)**.** *Let G be a graph of finite tree-width. Then G admits a fully tight, cofinally componental, rooted tree-decomposition $(T, \mathcal{V})$ whose adhesion sets are critical vertex sets, whose torsos are tough and which displays the critical vertex sets and their tight components cofinitely.*

Moreover,

(I1) *if $G\mathring{\uparrow}e$ is disconnected for $e = st \in E(T)$ with $s <_T t$, then $V_s \supseteq V_t \in \mathrm{crit}(G)$ and $\deg(t) = \infty$,*

(I2) *if $V_t = V_e$ for some node $t \in T$ and the unique edge $e = st \in T$ with $s <_T t$, then $\deg(t) = \infty$ and $V_t \in \mathrm{crit}(G)$.*

Recall that these (I1) and (I2) are the same properties as (I1) and (I2) in Theorem 4'.

This tree-decomposition is not difficult to construct: Start from the tree-decomposition $(T, \mathcal{V}'_{\mathrm{NT}})$ described in the introduction. By contracting all edges of the decomposition tree whose corresponding adhesion sets are not critical vertex sets of G, one obtains a tree-decomposition of G that satisfies all properties required for the tree-decomposition in Theorem 6'' except that it might not display the critical vertex sets. We then describe how one can turn this tree-decomposition into one that additionally displays all critical vertex sets and their tight components cofinitely.

We first collect two lemmas that describe how the critical vertex sets of a graph interact with a tight, componental, rooted tree-decomposition.

Lemma 3.5.1. *Let $(T, \mathcal{V})$ be a rooted tree-decomposition of a graph G of finite adhesion. Then every critical vertex set of G is contained in some bag of $(T, \mathcal{V})$.*

Moreover, if $(T, \mathcal{V})$ is tight, componental and its torsos are tough, then for every $X \in \mathrm{crit}(G)$ cofinitely many tight components of $G - X$ are of the form $G\mathring{\uparrow}e$ for an edge $e = t_X t \in T$ with $t_X <_T t$ where t_X is the (unique) $\leqslant_T$-minimal node of T whose corresponding bag contains X.

Proof. Since critical vertex sets are infinitely connected and $(T, \mathcal{V})$ has finite adhesion, every critical vertex set is contained in some bag of $(T, \mathcal{V})$. Thus, the nodes whose corresponding bags contain a fixed critical vertex set X form a subtree of T by (T2) which thus has a unique $\leqslant_T$-minimal node t_X. Let $e_0 <_T t_X$ be the unique edge of T incident with t_X. Since $(T, \mathcal{V})$ is componental, every tight component of $G - X$ which does not meet $G{\downarrow}e_0$ is of the form $G\mathring{\uparrow}e$ for an edge $e = t_X t \in T$ with $t_X <_T t$. By the choice of t_X, every tight component that meets $G\mathring{\downarrow}e_0$ also meets $V_{e_0} \subseteq V_{t_X}$. Lemma 3.2.5 (ii) yields that only finitely many tight components of $G - X$ meet V_{t_X}, as $(T, \mathcal{V})$ is tight and its torsos are tough. Thus, cofinitely many tight components of $G - X$ are of the desired form. □

The following lemma is kind of a converse of Lemma 3.5.1:

Lemma 3.5.2. *Let $(T, \mathcal{V})$ be a tight, rooted tree-decomposition of a graph G such that for every $X \in \mathrm{crit}(G)$ there is a node $t_X \in T$ such that cofinitely many tight components of $G - X$ are of the form $G\mathring{\uparrow}e$ for an edge $e \in t_X t \in T$ with $t_X <_T t$. Then all torsos of $(T, \mathcal{V})$ are tough.*

Proof. Let $t \in T$ be arbitrary, and suppose for a contradiction that $\mathrm{torso}(\sigma_t)$ is not tough, that is, there is some $X \in \mathrm{crit}(\mathrm{torso}(\sigma_t))$. Since $(T, \mathcal{V})$ is tight by assumption, we may apply Lemma 3.2.5 (i) to the star σ_t, which yields that X is also a critical vertex set of G; moreover, infinitely many tight components of $G - X$ meet V_t. In particular, we have $t \neq t_X$ by assumption on t_X. So let $e \in T$ be the edge incident with t_X on the unique t–t_X path in T, and let $(A, B) \in \sigma_{t_X}$ be the separation induced by the orientation of e towards t_X; in particular, $V_t \subseteq A$. Then by the assumption on t_X at most finitely many tight components of $G - X$ are contained in A, or equivalently meet A, which is a contradiction as $V_t \subseteq A$. □

We now use the previous two lemmas to show that every graph of finite tree-width admits a tree-decomposition which satisfies all properties required for the tree-decomposition in Theorem 6'' except that it might not display the critical vertex sets.

Lemma 3.5.3. *Every graph of finite tree-width admits a tight, componental, rooted tree-decomposition $(T, \mathcal{V})$ whose adhesion sets are critical vertex sets and whose torsos are tough. Moreover,*

(i) *$V_t \setminus V_e$ is non-empty for every node $t \in T$ and the unique edge $e = st \in T$ with $s <_T t$.*

Proof. We may assume that G is connected. Indeed, if G is not connected, we may apply Lemma 3.5.3 to every component C of G to obtain a tree-decomposition $(T^C, \mathcal{V}^C)$. Let T be the disjoint union of all T^C and a new node r to which we join the root of every T^C. We assign $V_r := \emptyset$ and $V_t := V_t^C$ for the respective component C of G. Now $(T, \mathcal{V})$ is as desired.

Since the connected graph G has finite tree-width, we may fix a normal spanning tree T of G by Theorem 3.2.2. Let $(T, \mathcal{V})$ ne the tree-decomposition into finite parts introduced unter the name tree-decomposition $(T, \mathcal{V}'_{\mathrm{NT}})$ in the introduction; i.e. the tree-decomposition with decomposition tree T whose bags are given by $V_t := \{t\} \cup N_G(\lfloor t \rfloor) \subseteq \lceil t \rceil$. By construction, for an edge $e = st \in T$ with $s <_T t$, we have $G\mathring{\uparrow}e = \lfloor t \rfloor$ and $V_e = \lceil s \rceil \cap N_G(\lfloor t \rfloor)$; in particular, the rooted tree-decomposition $(T, (V_t)_{t \in T})$ is componental and tight. Moreover, its bags are all finite and hence tough, so by Lemma 3.5.1 there exists for every $X \in \mathrm{crit}(G)$ a node $t_X \in T$ such that cofinitely many tight components of $G - X$ are of the form $G\mathring{\uparrow}e$ for an edge $e = t_X t \in T$ with $t_X <_T t$.

Let $(T', \mathcal{V}')$ be the tree-decomposition induced by $(T, \mathcal{V})$ given by contracting all edges $e \in T$ with $V_e \notin \mathrm{crit}(G)$. We claim that $(T', \mathcal{V}')$ is as desired. It is immediate from the construction that $(T', \mathcal{V}')$ is still tight and componental and that all its adhesion sets are critical vertex sets of G. So we are left to show that all the torsos of $(T', \mathcal{V}')$ are tough. To this end, for $X \in \mathrm{crit}(G)$, let t'_X be the node of T' whose branch set in T contains t_X. Since we only contracted edges of T whose adhesion set is not a critical vertex set of G, we still have that cofinitely many tight components of $G - X$ are of the form $G\mathring{\uparrow}e$ for an edge $e = t'_X s \in T'$ with $t'_X <_{T'} s$. Hence, we may apply Lemma 3.5.2, which concludes the proof.

To verify that all bags are distinct, let nodes $t' \neq s' \in T'$ be given, and pick nodes $t \in t'$ and $s \in s'$. In particular, $t \in V_t \subseteq V'_{t'}$ and $s \in V_s \subseteq V'_{s'}$ by the definition of $(T, \mathcal{V})$ and $(T', \mathcal{V}')$. If either $t \notin V'_{s'}$ or $s \notin V'_{t'}$, then $V'_{t'} \neq V'_{s'}$ and we are done; so suppose otherwise. Then, again by construction, there are nodes $y \in t'$ and $x \in s'$ such that $y \geqslant_T s$ and $x \geqslant_T t$. But since the branch sets t' and s' are connected in T and disjoint, this contradicts that T is a tree.

Finally, $(T, \mathcal{V})$ satisfies (i) by construction, as we have $t \in V_t \setminus V_e$ for every node $t \in T$ and its unique edge $e = st \in T$ with $s <_T t$. As $(T', \mathcal{V}')$ was induced by edge-contractions on T, also $(T', \mathcal{V}')$ satisfies (i). □

Next, we describe how one can turn the tree-decomposition from Lemma 3.5.3 into one that additionally displays all critical vertex sets and their tight components cofinitely.

Construction 3.5.4. Let $(T, \mathcal{V})$ be a rooted tree-decomposition of a graph G of finite adhesion. For every $X \in \text{crit}(G)$, let t_X be the $\leqslant_T$-minimal node t with $X \subseteq V_t$ (which exists by Lemma 3.5.1). We obtain the tree T' from T by simultaneously adding for each $X \in \text{crit}(G)$ the node t'_X, the edge $t_X t'_X$ and rerouting each edge $t_X t$ of T with $V_{t_X t} = X$ to $t'_X t$ in T'. The tree-decomposition $(T', \mathcal{V}')$ is given by $V'_t := V_t$ for all $t \in T$ and $V'_{t'_X} := X$ for every $X \in \text{crit}(G)$.

To show that the tree-decomposition $(T', \mathcal{V}')$ from Construction 3.5.4 displays the critical vertex sets of G we need the following auxiliary lemma.

Lemma 3.5.5. *Let $(T, \mathcal{V})$ be a tight, rooted tree-decomposition of a graph G such that for every $X \in \text{crit}(G)$ there is a unique node $t_X \in T$ with $V_{t_X} = X$ and cofinitely many tight components of $G - X$ are some $G\mathring{\uparrow}e$ for an edge $e = t_X t \in T$ with $t_X <_T t$. Then $(T, \mathcal{V})$ displays all critical vertex sets of G.*

Proof. We need to show for every infinite-degree node $t \in T$ with finite V_t that $V_t \in \text{crit}(G)$. So let $t \in T$ be a node with finite V_t and $V_t \notin \text{crit}(G)$. Since V_t is finite, some subset $X \subseteq V_t$ must be the adhesion set corresponding to infinitely many edges e incident with t in T. Since $(T, \mathcal{V})$ is tight, X is a critical vertex set of G. By assumption on $(T, \mathcal{V})$ there is a unique node $t_X \in T$ with $V_{t_X} = X$. If $t_X = t$ we are done so suppose otherwise. Since by assumption only finitely many tight components of $G - X$ can meet $G\mathring{\downarrow}f$ for the unique edge $f = st_X$ with $s <_T t_X$, we have that $t_X <_T t$. But then, again by the assumption on $(T, \mathcal{V})$, at most finitely many edges $e = ts$ with $t <_T s$ can contain a tight component of $G - X$ contradicting that $(T, \mathcal{V})$ is tight. □

We can now show that the tree-decomposition $(T', \mathcal{V}')$ from Construction 3.5.4 is as desired for Theorem 6'' if we start with a tree-decomposition $(T, \mathcal{V})$ that is tight and componental.

Lemma 3.5.6. *Let $(T, \mathcal{V})$ be a tight, componental, rooted tree-decomposition of a graph G of finite adhesion whose torsos are all tough. Then the tree-decomposition $(T', \mathcal{V}')$ from Construction 3.5.4 displays all critical vertex sets and their tight components cofinitely. Additionally,*

(i) *$(T', \mathcal{V}')$ is fully tight,*

(ii) *$(T', \mathcal{V}')$ is cofinally componental; moreover, if $G\mathring{\uparrow}e$ is disconnected for $e = st \in E(T')$ with $s <_{T'} t$ then $V_s \supseteq V_t \in \operatorname{crit}(G)$ and $\deg(t) = \infty$,*

(iii) *the torsos of $(T', \mathcal{V}')$ are tough.*

Proof. Let $(T', \mathcal{V}')$ be the tree-decomposition from Construction 3.5.4. Then (i), (ii) and (iii) hold for $(T', \mathcal{V}')$. Lemma 3.5.1 applied to $(T, \mathcal{V})$ ensures that for $X \in \operatorname{crit}(G)$ cofinitely many tight components of $G - X$ are of the form $G\mathring{\uparrow}e$ for an edge $e = t'_X t \in T'$ with $t'_X <_{T'} t$. Thus, by Lemma 3.5.5, $(T', \mathcal{V}')$ displays all critical vertex sets of G. Moreover, by construction, $(T', \mathcal{V}')$ displays the tight components of all critical vertex sets cofinitely. □

Proof of Theorem 6''. Apply Construction 3.5.4 to the tree-decomposition $(T, \mathcal{V})$ from Lemma 3.5.3. By Lemma 3.5.6 this tree-decomposition $(T', \mathcal{V}')$ is as desired; in particular, it satisfies (I1). Moreover, since $(T, \mathcal{V})$ satisfies (i) from Lemma 3.5.3, $(T, \mathcal{V})$ still satisfies it at all $t \in T'$ which have not been not added in Construction 3.5.4, i.e. $(T', \mathcal{V}')$ satisfies (I2). □

We remark that a theorem similar to Theorem 6'' was proven by Elm and Kurkofka [65, Theorem 2]. However, their result cannot be used in place of our Theorem 6'', as their method in general produces only a nested set of separations, and not a tree-decomposition (see the appendix of [12] for a detailed discussion).

3.6 Lifting paths and rays from torso

From Theorem 6' we obtain a tree-decomposition $(T, \mathcal{V})$ whose torsos are tough. Following our overall proof strategy for Theorem 1 we want to further decompose each of its torsos $\operatorname{torso}(\sigma_t)$ with a tree-decomposition $(T^t, \mathcal{V}^t)$ which is linked and rayless. We have already indicated in the beginning of Section 3.4.1 that by using the infinite connectivity of critical vertex sets the arising tree-decomposition $(T', \mathcal{V}')$ is linked, if $(T^t, \mathcal{V}^t)$ is linked. For this, we need to extend the path families in $\operatorname{torso}(\sigma_t)$ which witness the linkedness of $(T^t, \mathcal{V}^t)$ to path families in G joining up the same sets of vertices. In this section we show that such an extension of (finite) path families is possible, as long as the separations in the star σ_t at t are 'left-well-linked' (Lemma 3.6.3). Moreover, we prove similar results, Proposition 3.6.1 and Lemma 3.6.2, for extending rays and infinite path families in torsos at stars whose separations are just left-tight.

Proposition 3.6.1. *Let σ be a star of left-tight, finite-order separations of a graph G. Let P be a path or ray in* $\operatorname{torso}(\sigma)$*. Then there exists a path or ray P' in G, respectively, with $P' \cap G[\operatorname{int}(\sigma)] \subseteq P$ which starts in the same vertex as P and, if P is a path, also ends in the same vertex as P, and such that P' meets $V(P)$ infinitely often if P is ray.*

Proof. Fix for every torso edge $e = uv$ of P a separation $(A_e, B_e) \in \sigma$ such that both $u, v \in A_e \cap B_e$. Since each (A_e, B_e) is left-tight, we may further fix for every torso edge $e = uv$ of P a u–v path P_e through $A_e \setminus B_e$. Note that since all $(A, B) \in \sigma$ have finite order, we have $(A_e, B_e) = (A_f, B_f)$ for at most finitely many torso edges f of P. As the strict left sides $A \setminus B$ of the separations (A, B) in the star σ are pairwise disjoint, every P_e meets at most finitely many P_f for the torso edges f of P.

Then the graph H obtained from $P \cap G$ by adding all the P_e for torso edges of P is a connected subgraph of G which contains the startvertex and, if P is a path, the endvertex of P. Moreover, H is locally finite: by construction, all vertices in $P \cap G$ have degree 1 or 2 in H since P is a path or ray, and all vertices in $H - V(P)$ are contained in at most finitely many P_e by the argument above, and hence also have finite degree in H. Now if P is a path, then $H \subseteq G$ contains a path P' whose endvertices are the same as P, and if P is a ray, then $H \subseteq G$ contains a ray P' that starts in the same vertex as P (cf. [41, Proposition 8.2.1]). In particular, if P is a ray, then P' meets $V(P)$ infinitely often since each component of $H - V(P)$ is finite.

Since $H \cap G[\operatorname{int}(\sigma)] = P \cap G[\operatorname{int}(\sigma)]$ by construction, and because $P' \subseteq H$, we have that $P' \cap G[\operatorname{int}(\sigma)] \subseteq P$, so P' is as desired. □

We remark that later in this chapter we sometimes want to apply Proposition 3.6.1 to a ray P in the torso at a star σ of finite-order separations of a graph G in which all but one separation (A, B) are left-tight but we are not interested in keeping the startvertex of P. Hence, we may apply Proposition 3.6.1 instead to the tail of P which avoids the finite set $A \cap B$ and the star $\{(C \cap B, D \cap B) \mid (C, D) \in \sigma \setminus \{(A, B)\}\}$ induced by $\sigma \setminus \{(A, B)\}$ on $G[B]$. This still yields a ray P' in G with $P' \cap G[\operatorname{int}(\sigma)] \subseteq P$ which meets $V(P)$ infinitely often.

Lemma 3.6.2. *Let σ be a star of left-tight, finite-order separations of a graph G. Let $X, Y \subseteq \operatorname{int}(\sigma)$ such that there are infinitely many disjoint X–Y paths in* $\operatorname{torso}(\sigma)$*. Then there are infinitely many X–Y paths in G.*

Proof. Let $\mathcal{P}$ be an infinite family of disjoint X–Y paths in $\operatorname{torso}(\sigma)$. We define an infinite family $\mathcal{P}'$ of disjoint X–Y paths in G recursively. For this, set $P'_0 := \emptyset$, let $n \in \mathbb{N}$, and assume that we have already constructed a family $\mathcal{P}'_n$ of n pairwise disjoint X–Y paths in G.

Then $V(\mathcal{P}'_n)$ is finite, and hence meets the separators of at most finitely many separations in σ. Since all separations in σ have finite order, and because $\mathcal{P}$ is an infinite family of disjoint paths, there exists a path $P \in \mathcal{P}$ that avoids both the finite set $V(\mathcal{P}'_n)$ and the finitely many finite separators of separations in σ that meet $V(\mathcal{P}'_n)$. By Proposition 3.6.1, there exists a path P' in G with the same endvertices as P and with $P' \cap G[\mathrm{int}(\sigma)] \subseteq P$. In particular, P' is an X–Y path in G. Moreover, P' is disjoint from the paths in $\mathcal{P}'_n$ by the choice of P and because $P' \cap G[\mathrm{int}(\sigma)] \subseteq P$. Hence, we may define $\mathcal{P}'_{n+1} := \mathcal{P}'_n \cup \{P'\}$. Then $\mathcal{P}' := \bigcup_{n\in\mathbb{N}} \mathcal{P}'_n$ is as desired. □

Let us call a finite-order separation (A, B) of a graph G *left-well-linked* if, for every two disjoint sets $X, Y \subseteq A \cap B$, there is a family of $\min\{|X|, |Y|\}$ disjoint X–Y paths in G through $A \setminus B$.

Lemma 3.6.3. *Let σ be a star of left-well-linked, finite-order separations of a graph G. Then the following assertions hold:*

(i) *For every countable family $\mathcal{P}$ of disjoint rays in* $\mathrm{torso}(\sigma)$ *there exists a family $\mathcal{P}'$ of disjoint rays in G with the same set of startvertices, and such that each ray in $\mathcal{P}'$ meets $V(\mathcal{P})$ infinitely often.*

(ii) *For every $k \in \mathbb{N}$ and every family $\mathcal{P}$ of k disjoint paths in* $\mathrm{torso}(\sigma)$ *there exists a family $\mathcal{P}'$ of k disjoint paths in G with the same set of startvertices and the same set of endvertices.*

Proof. We prove (i) and (ii) simultaneously; so let $\mathcal{P}$ be as in (i) or as in (ii). We define the family $\mathcal{P}'$ recursively, going through the countably many torso edges on $\mathcal{P}$. For this, fix an enumeration $e_1, e_2, \dots$ of the torso edges in $\mathcal{P}$ such that if $e_i \in P \in \mathcal{P}$, then all torso edges occurring on P before e_i are enumerated as e_j for some $j < i$. Further, for every e_i, we fix some $(A_i, B_i) \in \sigma$ such that both endvertices of e_i are contained in $A_i \cap B_i$ where we choose $(A_i, B_i) = (A_j, B_j)$ for some $j < i$ if possible.

We start the recursion with $\mathcal{P}_0 := \mathcal{P}$. At step $i \in \mathbb{N}$, we assume that we have already constructed a family $\mathcal{P}_{i-1}$ of disjoint paths/rays in $\mathrm{torso}(\sigma) \cup \bigcup_{j<i} G[A_j]$ with the same startvertices, and, if $\mathcal{P}$ is as in (ii), with the same endvertices, as $\mathcal{P}$, and without all torso edges whose endvertices are both contained in $A_j \cap B_j$ for some $j < i$. We now consider the torso edge e_i. If e_i is not contained in any path/ray in $\mathcal{P}_{i-1}$, then we set $\mathcal{P}_i := \mathcal{P}_{i-1}$. Otherwise, for each $P \in \mathcal{P}_{i-1}$, let x_P be the first vertex on P such that the subsequent edge on P is a torso edge with both endvertices in $A_i \cap B_i$, and let y_P be the respective last vertex on P such that its previous edge on P is a torso edge with both endvertices in $A_i \cap B_i$. Now set $X_i := \{x_P \mid P \in \mathcal{P}_i\}$ and $Y_i := \{y_P \mid P \in \mathcal{P}_i\}$. Since the paths/rays in $\mathcal{P}_{i-1}$ are disjoint, X_i and Y_i are disjoint subsets of $A_i \cap B_i$ with $k_i := |X_i| = |Y_i|$. We can thus use that (A_i, B_i) is left-well-linked to find a family $\mathcal{Q}$ of k_i disjoint X_i–Y_i paths in $G[(A_i \setminus B_i) \cup X_i \cup Y_i]$; for $Q \in \mathcal{Q}$, write x_Q for its first and y_Q for its last vertex.

For each $P \in \mathcal{P}_i$, we now define $P^* := P x_P x_Q Q y_Q y_{P'} P'$ where Q is the unique path in $\mathcal{Q}$ with $x_Q = x_P$ and P' is the unique path/ray in $\mathcal{P}'$ with $y_{P'} = y_Q$. By construction, the set $\mathcal{P}_i := \{P^* \mid P \in \mathcal{P}_{i-1}\}$ is a family of disjoint paths/rays in $\mathrm{torso}(\sigma) \cup \bigcup_{j\leqslant i} G[A_j]$ with the same startvertices, and, if $\mathcal{P}$ is as in (ii), with the same endvertices, as $\mathcal{P}_{i-1}$ and thus as $\mathcal{P}$. Moreover, $\mathcal{P}_i$ not only avoids all torso edges e_j with $j < i$ but also all torso edges whose endvertices are both contained in $A_j \cap B_j$ for some $j < i$. This completes step i.

To define $\mathcal{P}'$ from the $\mathcal{P}_i$, let X be the set of startvertices of paths/rays in $\mathcal{P}$ and let Y be the set of endvertices of paths in $\mathcal{P}$. By construction, for each vertex $x \in X$ there is a (unique) path/ray in $\mathcal{P}_i$ starting in x and we denote it by P_i^x. We let P^x be the limit $\liminf_{i \in \mathbb{N}} P_i^x$ of the P_i^x, and define $\mathcal{P}' := \{P^x \mid x \in X\}$. By construction, all the P^x are disjoint. Note that if $\mathcal{P}$ is as in (ii), then $\mathcal{P}' = \mathcal{P}_n$ for some $n \in \mathbb{N}$, since the construction yields $P_i^x = P_j^x$ for all $i, j \geqslant |E(\mathcal{P})|$; thus $\mathcal{P}'$ is as desired for (ii). We now prove that if $\mathcal{P}$ is as in (i), then all P^x are indeed rays which meet $V(\mathcal{P})$ infinitely often.

If an initial segment of some P_i^x is contained in G, then it contains no torso edge and it hence remains untouched by the above construction in all steps $j \geqslant i$; in other words, this initial segment of P_i^x in G is also an initial segment of all P_j^x with $j \geqslant i$. Moreover, if P_i^x still contains some torso edge and we let e_j be the first such one occurring along P_i^x, then the construction at step j implies that the maximal initial segment of P_j^x in G is strictly longer than the one in P_i^x. Moreover, it contains a vertex in $V(\mathcal{P})$, the respective $y_{(P_j^x)'}$ in the above construction step, that has not been contained in the maximal initial segment of P_i^x in G.

Now if P_i^x contains no torso edge at some step i, then it is a ray in G starting in x, and thus some tail of P_i^x equals the tail of some ray in $\mathcal{P}$ by construction; in particular, P_i^x meets $V(\mathcal{P})$ infinitely often. Otherwise, the length of the initial segment of P_i^x in G strictly increases infinitely often, and hence the limit P^x of the P_i^x is a ray in G starting in x that meets infinitely many vertices of $V(\mathcal{P})$, which is both witnessed by the respective $y_{(P_j^x)'}$ described above. Thus, $\mathcal{P}'$ is as desired if $\mathcal{P}$ is as in (i). This concludes the proof. □

3.7 Linked tree-decompositions into rayless parts

In this section we prove Theorem 7', which we restate here for convenience.

Theorem 7'' (Detailed version of Theorem 7)**.** *Let G be a graph, and let σ be a star of left-well-linked, left-fully-tight, finite-order separations of G such that* torso(σ) *has finite tree-width. Further, let $X \subseteq \operatorname{int}(\sigma)$ be a prescribed finite set of vertices of G.*

Then G admits a linked X-linked, fully tight, rooted tree-decomposition $(T, \mathcal{V})$ of finite adhesion such that

(R1) *its torsos at non-leaves are rayless and its leaf separations are precisely* $\{(B, A) \mid (A, B) \in \sigma\}$,

(R2) *for all edges e of T, the adhesion set V_e is either linked to an end living in $G{\uparrow}e$ or linked to a set $A \cap B \subseteq G{\uparrow}e$ with $(A, B) \in \sigma$,* (R3) *for every* $e <_T e' \in E(T)$ *with* $|V_e| \leqslant |V_{e'}|$*, each vertex of* $V_e \cap V_{e'}$ *either dominates some end of* G *that lives in* $G{\uparrow}e'$ *or is contained in* $A \cap B \subseteq V(G{\uparrow}e')$ *for some* $(A, B) \in \sigma$*,*

(R4) *for all edges* $e = st \in T$ *with* $s <_T t$ *and* $s \neq r := \mathrm{root}(T)$ *we have* $V_s \supsetneq V_e \subsetneq V_t$*. Moreover, if* $X \subsetneq \mathrm{int}(\sigma)$*,* $G - X$ *is connected and* $N_G(G - X) = X$*, then* $X \subsetneq V_r$ *and also* $V_r \supsetneq V_e \subsetneq V_t$ *for all edges* $e = rt \in T$*, and*

(R5) *if* $G\mathring{\uparrow}e$ *is disconnected for some edge* $e \in T$*, then* e *is incident with a leaf of* T*.*

We remark that in Theorem 7'' we explicitly allow the case $\sigma = \emptyset$, where we go with the convention that the interior of the empty star is $V(G)$.

First we reduce Theorem 7'' to the following statement, which differs slightly from Theorem 7''. A separation (A, B) of a graph G is *left-connected* if $G[A \setminus B]$ is connected, and it is *left-end-linked* if $A \cap B$ is linked to some end ε of G with $C(A \cap B, \varepsilon) \subset A$.

Theorem 3.7.1. *Let* G *be a graph of finite tree-width, let* σ *be a star of left-end-linked left-connected finite-order separations of* G*, and let* $X \subseteq \mathrm{int}(\sigma)$ *be a finite set of vertices of* G*. Assume further that for every* $(A, B) \in \sigma$*, the graph* $G[A \cap B]$ *is complete and* $A \cap B \subseteq \mathrm{Dom}(G[A \setminus B])$*.*

Then G *admits a linked,* X*-linked, tight, componental, rooted tree-decomposition* $(T, \mathcal{V})$ *of finite adhesion such that*

(R1') *its torsos at non-leaves are rayless, its leaf separations are precisely* $\{(B, A) \mid (A, B) \in \sigma\}$ *and no other separation induced by an edge of* T *is* (B, A) *for* $(A, B) \in \sigma$*,*

(R2') $(T, \mathcal{V})$ *is end-linked,*

(R3') *for every* $e <_T e' \in E(T)$ *with* $|V_e| \leqslant |V_{e'}|$ *each vertex of* $V_e \cap V_{e'}$ *dominates some end of* G *that lives in* $G{\uparrow}e'$*, and*

(R4') $(T, \mathcal{V})$ *satisfies* (R4) *from Theorem 7'.*

Proof of Theorem 7'' given Theorem 3.7.1. Let H be the graph obtained from $\mathrm{torso}(\sigma)$ by adding, for every separation $(A, B) \in \sigma$ a disjoint ray $R_{A,B}$ as well as an edge from every vertex $v \in A \cap B$ to every vertex of $R_{A,B}$. We aim to apply Theorem 3.7.1 to H with X and the star

$$\sigma' := \{(V(R_{A,B}) \cup (A \cap B), V(H - R_{A,B})) \mid (A, B) \in \sigma\};$$

but to be able to do so we first have to show that H has finite tree-width.

By assumption in Theorem 7'', $\mathrm{torso}(\sigma)$ admits a tree-decomposition $(T^\sigma, \mathcal{V}^\sigma)$ into finite parts. Any subgraph of H of the form $H[(A \cap B) \cup V(R_{A,B})]$ clearly also admits such a tree-decomposition $(T^{A,B}, \mathcal{V}^{A,B})$. Since $H[A \cap B]$ is complete for all $(A, B) \in \sigma$, it is contained in some part of $(T^\sigma, \mathcal{V}^\sigma)$ and also it is contained in some part of $(T^{A,B}, \mathcal{V}^{A,B})$. We then obtain the desired decomposition tree from the disjoint union of the decomposition trees by adding for each $(A, B) \in \sigma$ an edge between the nodes corresponding to the respective parts containing $H[A \cap B]$. Keeping the parts yields the desired tree-decomposition of H into finite parts.

So by construction and the previous argument, H, X and σ' are as required for Theorem 3.7.1, which then yields a rooted tree-decomposition $(T, \mathcal{V}')$ of H. In particular, this $(T, \mathcal{V}')$ has precisely $((B, A) \mid (A, B) \in \sigma')$ as its leaf separations. Thus, we obtain a tree-decomposition $(T, \mathcal{V})$ of G by letting $V_t := V'_t$ for all non-leaves of T and $V_t := B$ for all leafs of T whose bag is of the form $(A \cap B) \cup V(R_{A,B})$. In particular, its leaf separations are precisely $\{(B, A) \mid (A, B) \in \sigma\}$.

We claim that the tree-decomposition $(T, \mathcal{V})$ of G is as desired. We remark that the adhesion sets corresponding to an edge of the decomposition tree are unchanged. So, $(T, \mathcal{V})$ still has finite adhesion, as $(T, \mathcal{V}')$ has finite adhesion. Also, $(T, \mathcal{V})$ is linked and X-linked: as the $H[A \cap B]$ are complete, this follows from the (X-)linkedness of $(T, \mathcal{V}')$ and Lemma 3.6.3 (ii). Further, $(T, \mathcal{V})$ is fully tight, since $(T, \mathcal{V}')$ is fully tight and all separations in σ are left-fully-tight. Also $(T, \mathcal{V})$ satisfies (R5), since $(T, \mathcal{V}')$ is componental and the separations in σ are fully tight and no separation of H induced by an edge of T which is not incident with a leaf is (B, A) for some $(A, B) \in \sigma'$ by (R1'). By construction of $(T, \mathcal{V})$, property (R1') of $(T, \mathcal{V}')$ immediately implies that $(T, \mathcal{V})$ satisfies (R1). Additionally, for all edges e of T, the adhesion set V'_e is linked to an end of H by (R2'). If that end corresponds to some ray $R_{A,B}$, then $V_e = V'_e \subseteq \operatorname{int}(\sigma)$ is linked to $A \cap B \subseteq V(G{\uparrow}e)$ since $A \cap B$ separates $\operatorname{int}(\sigma)$ from $R_{A,B}$ and because of Lemma 3.6.3 (ii) as the $H[A \cap B]$ are complete. Otherwise, Lemma 3.6.3 (i) and Lemma 3.6.2 ensure that V_e is linked to an end of G that lives in $G{\uparrow}e$. Indeed, let $\mathcal{R}$ be a family of $|V_e|$ equivalent, disjoint rays in $\operatorname{torso}(\sigma)$ that start in V_e and that witness that V_e is end-linked. Then Lemma 3.6.3 (i) yields a family $\mathcal{R}'$ of $|V_e|$ disjoint rays in $G{\uparrow}e$ that start in V_e and that each meet $V(\mathcal{R})$ infinitely often. In particular, since $\mathcal{R}$ is finite, for every $R' \in \mathcal{R}'$ there is $R \in \mathcal{R}$ such that R' meets $V(R)$ infinitely often. To see that the rays in $\mathcal{R}'$ are equivalent, let $R'_0, R'_1 \in \mathcal{R}'$ be given. Since R_0 and R_1 are equivalent in $\operatorname{torso}(\sigma)$, the infinite sets $V(R'_0) \cap V(R_0)$ and $V(R'_1) \cap V(R_1)$ cannot be separated by finitely many vertices. Hence, we may greedily pick infinitely many disjoint $V(R'_0) \cap V(R_0)$–$V(R'_1) \cap V(R_1)$ paths in $\operatorname{torso}(\sigma)$. Since these paths are in fact R'_0–R'_1 paths in $\operatorname{torso}(\sigma)$, Lemma 3.6.2 yields that there are infinitely many disjoint R'_0–R'_1 paths in G, which concludes the proof that V_e is end-linked in G, and that $(T, \mathcal{V})$ satisfies (R2).

Also $(T, \mathcal{V})$ satisfies (R3) because $(T, \mathcal{V}')$ satisfies (R3'). Finally, $(T, \mathcal{V})$ satisfies (R4) since $(T, \mathcal{V}')$ satisfies (R4') too. □

Let us briefly sketch the proof of Theorem 3.7.1. We will construct the tree-decomposition inductively. We start with the trivial tree-decomposition whose decomposition tree is a single vertex whose bag is the whole vertex set of G. In the induction step, we assume that we have already constructed a linked, X-linked, tight, componental, rooted tree-decomposition $(T^n, \mathcal{V}^n)$

of G of finite adhesion which satisfies (R2'), (R3') but it may not satisfy (R1'). Instead its torsos at non-leaves are rayless and A is contained in a bag at a leaf of T^n for every $(A, B) \in \sigma$. Then, for every leaf ℓ of T^n whose (unique) incident edge e does not induce a separation in σ, we define a set $Y \subseteq V(G{\uparrow}e)$ such that the adhesion set V_e^n is contained in Y. We then replace the bag V_ℓ^n with Y and add for each component C of $(G{\uparrow}e) - Y$ a new leaf to ℓ and associate the bag $V(C) \cup N(C)$ to it. By carefully choosing these sets Y for each such leaf we will ensure that the arising tree-decomposition $(T^{n+1}, \mathcal{V}^{n+1})$ again satisfies all the properties as assumed for $(T^n, \mathcal{V}^n)$. We then show that the pair $(T, \mathcal{V})$ which arises as the limit for $n \to \infty$ is indeed a tree-decomposition of G and that it satisfies all the desired properties.

This section is organised as follows. First, we describe in Section 3.7.1 the Algorithm 3.7.2 that will construct the bags of the desired tree-decomposition for Theorem 3.7.1, that is, the sets Y mentioned above. This algorithm is simple to state, but we require some tools to properly analyse it. In Section 3.7.2 we build a set of tools centred around 'regions', which we then use in Section 3.7.3 to prove some properties of Algorithm 3.7.2 that will later on ensure that the resulting tree-decomposition is as desired. Finally, in Section 3.7.4, we follow the above described approach to construct a tree-decomposition by inductively applying Algorithm 3.7.2 and then prove with the help of the main result from Section 3.7.3 that this tree-decomposition is as desired.

3.7.1 Building the bags of the tree-decomposition

In this section we describe a transfinite recursion, Algorithm 3.7.2, that will construct the bags of the tree-decomposition for Theorem 3.7.1. For this, we need the following definitions.

Let G be a graph. A *region* C of G is a connected subgraph of G. By $\bar{C}$ we denote the *closure* $G[V(C) \cup N_G(C)]$ of C. A *k-region* for $k \in \mathbb{N}$ is a region C whose neighbourhood $N(C)$ has size k. A $({<}\,k)$- or $({\leqslant}\,k)$*-region* for $k \in \mathbb{N}$ is a k'-region C for some $k' < k$ or $k' \leqslant k$, respectively. Similarly, an $({<}\,\aleph_0)$*-region* is a k-region for some $k \in \mathbb{N}$. Two regions C and D of G *touch*, if they have a vertex in common or G contains an edge between them. Two regions C and D of a graph G are *nested* if they do not touch, or $C \subseteq D$, or $D \subseteq C$. Note that the set of all regions $G\mathring{\uparrow}e$ given by a rooted componental tree-decomposition of a graph G is nested.

A region C is *ε-linked* for an end ε of G if ε lives in C and the neighbourhood of C is linked to ε. A region C is *end-linked* if it is ε-linked for some end ε of G. We emphasise that a region C is ε-linked if $N(C)$ is linked to the end ε (and not $V(C)$). To distinguish both cases, we say a region is ε-linked and a set of vertices is linked to ε.

Algorithm 3.7.2. (Construction of a bag)

Input: a connected graph H; a finite set X of $k \in \mathbb{N}$ vertices of H; a set $\mathcal{D}$ of pairwise non-touching end-linked $({<}\,\aleph_0)$-regions D that are disjoint from X and satisfy $X \cap N_H(D) \subseteq \mathrm{Dom}(D)$;

a vertex $x \in V(G) \setminus X$ that lies in no $D \in \mathcal{D}$.

Output: a transfinite sequence $C_0, C_1, \ldots, C_i, \ldots$, indexed by some ordinal $< |H|^+$, of distinct end-linked $(< \aleph_0)$-regions which are disjoint from X and pairwise nested and which are also nested with all $D \in \mathcal{D}$; a set $Y := V(H) \setminus (\bigcup_i C_i)$

Recursion: Iterate the following step:

Case A: If there is a $(< k)$-region C_i of H that is disjoint from X, is ε-linked for some end ε which lives in no C_j for $j < i$, nested with the C_j for $j < i$ and with all $D \in \mathcal{D}$, then choose a *nicest* such region C_i. Here, nicest[12] means that

- (N1) C_i is such an ℓ_i-region where $\ell_i \in \mathbb{N}$ is minimum among such regions, and
- (N2) C_i is an inclusion-wise maximal such region subject to (N1).

Case B: If there is no region as in Case A, but there is a $(< \aleph_0)$-region C_i that is disjoint from $X' := X \cup \{x\}$, is ε-linked for some end ε which lives in no C_j for $j < i$, nested with the C_j for $j < i$ and with all $D \in \mathcal{D}$ such that $X' \cap N_H(C_i) \subseteq \mathrm{Dom}(C_i)$, then choose a nicest such region C_i.

Case C: If there is no region as in Case A or Case B, then terminate the recursion.

Recall that we will construct the tree-decomposition for Theorem 3.7.1 recursively, where in each step we replace the leaves of the previous tree-decomposition with stars whose torsos at the centre vertices are rayless. So we can think of the graph H in Algorithm 3.7.2 as the part $G{\uparrow}\ell$ above a leaf ℓ of the decomposition tree T^n constructed so far and the set X as the adhesion set V_e that corresponds to the unique edge e incident with that leaf in T^n. The set Y which Algorithm 3.7.2 outputs will then be the bag at the centre of the newly added star, and the parts at the new leaves will be the components of $H - Y$ (which will be the $\supseteq$-maximal elements of the C_i, see Theorem 3.7.9 below) together with their boundaries.

Let us briefly give some intuition on why the set Y is a good candidate for the new bag. Theorem 3.7.1 requires the torsos of the tree-decomposition to be rayless, so Y should not contain any end of H: for this, Algorithm 3.7.2 iteratively cuts off all ends of H by choosing $(< \aleph_0)$-regions C_i around them. Additionally, we have to make sure that in the limit of our construction we do end up with a tree-decomposition which in particular must satisfy (T1). The specified vertex x will ensure that we make the appropriate progress when defining the tree-decompositions. In fact, in the final construction of the tree-decomposition we will carefully specify the vertex x in order to ensure that in the end every vertex will lie in some bag of the pair $(T, \mathcal{V})$ arising as the limit for $n \to \infty$. But this is not the only point where we need to be careful. Since the tree-decomposition should be linked, we cannot just choose the regions C_i arbitrarily; instead, we need to choose those first whose neighbourhood has size less than $|X|$. This is encoded in Case A, while Case B then

[12] Note that there might be several nicest regions.

will cut off all the remaining ends of G. We will justify these intuitions in Theorem 3.7.9.

We begin by argueing that Algorithm 3.7.2 is well-defined: Whenever there exists a region in Case A or Case B, then there obviously exists such a region as in Case A or Case B satisfying (N1), and then Zorn's Lemma and the following lemma (specifically: item (ii)) ensure that there is also such a region satisfying (N2).

Lemma 3.7.3. *Let H be a graph. Let $\mathcal{C}$ be a chain of regions of H. Then $C = \bigcup \mathcal{C}$ is a region, and $\mathcal{C}' := \{V(H - C) \cap N_H(C') \mid C' \in \mathcal{C}\}$ is a chain with $N_H(C) = \bigcup \mathcal{C}'$.*

Moreover, let $X \subseteq V(H)$. Then the following statements hold:

(i) *If all $C' \in \mathcal{C}$ are disjoint from X, then C is disjoint from X.*

(ii) *If, for some $k \in \mathbb{N}$, every $C' \in \mathcal{C}$ is a $(\leqslant k)$-region, then C is a $(\leqslant k)$-region.*

(iii) *If, for every $C' \in \mathcal{C}$, $X \cap N_H(C') \subseteq \mathrm{Dom}(C')$, then $X \cap N_H(C) \subseteq \mathrm{Dom}(C)$.*

(iv) *If $N_H(C)$ is finite and every $C' \in \mathcal{C}$ is end-linked, then C is end-linked.*

(v) *If every $C' \in \mathcal{C}$ is nested with a region D, then also C is nested with every D.*

Proof. Since $\mathcal{C}$ is a chain and all $C' \in \mathcal{C}$ are connected, C is connected, too. By definition of neighbourhood, $N_H(C) = V(H - C) \cap \bigcup_{C' \in \mathcal{C}} N_H(C')$ and, since $\mathcal{C}$ is a chain, also the set $\mathcal{C}' = \{V(H - C) \cap N_H(C') \mid C' \in \mathcal{C}\}$ forms a chain. This ensures that (ii) and (iii) hold. Next, (i) follows immediately from the definition of C as the union of the $C' \in \mathcal{C}$.

(iv): If the neighbourhood of C is finite, the fact that $\mathcal{C}'$ is a chain with $\bigcup \mathcal{C}' = N_H(C)$ yields that there is a $C' \in \mathcal{C}$ such that $N_H(C') \supseteq N_H(C)$. Now, since $C' \subseteq C$, the end-linkedness of C' yields the end-linkedness of C.

(v): Since every $C' \in \mathcal{C}$ is nested with D, the regions C' and D either do not touch or one is contained in the other. If D is contained in some $C' \in \mathcal{C}$, then D is also contained in $C \supseteq C'$, as desired. So we may assume that D is contained in no $C' \in \mathcal{C}$. Note that whenever $C'_0 \in \mathcal{C}$ is contained in D or does not touch D, every $C' \in \mathcal{C}$ with $C' \subseteq C'_0$ is contained in D or does not touch D, respectively. Thus, either all $C' \in \mathcal{C}$ are contained in D or they all do not touch D. This yields that their union C either is contained in D or does not touch D, respectively. □

3.7.2 Regions

In this section we collect some statements about regions which we then use in Section 3.7.3 to analyse the regions C_i and the set Y which Algorithm 3.7.2 outputs. Recall that we want Algorithm 3.7.2 to cut off all ends of H in that every end of H lives in some C_i. In order to prove that Algorithm 3.7.2 actually achieves this (at least in the case where every end of H has countable combined degree), we will work in this section towards Lemmas 3.7.6 and 3.7.7, which will later on ensure that if at step i of the recursion in Algorithm 3.7.2 there is still an end of H that does not live in any C_j for $j < i$, then there is a region that is a 'candidate for C_i' in Case A or Case B. Then Algorithm 3.7.2 will not terminate as long as there is still some such uncovered end of G that lives in no C_j.

We say that a region C is a *candidate for* C_i if C_i was chosen in Case A or Case B and C satisfied all properties of the regions in Case A or Case B at step i, respectively, except that C may have not been a nicest such region. We start by showing for that every end of H there is a region C_i as in in Case A (Lemma 3.7.4 (i)) or Case B (Lemma 3.7.4 (ii)) – except that it might not be nested with $\mathcal{D}$ and with the previously chosen C_j for $j < i$.

Lemma 3.7.4. *Let H be a graph, let $X \subseteq V(H)$ be finite, and let ε be an end of H with countable combined degree. Then the following statements hold:*

(i) *For every $|\cdot|$-minimal X–ε separator S the component of $H - S$ in which ε lives is ε-linked and disjoint from X.*

(ii) *There is an ε-linked $(<\aleph_0)$-region C disjoint from X which satisfies $X \cap N_H(C) \subseteq \mathrm{Dom}(\varepsilon)$.*

Proof. (i): Since S separates X and ε, the component $C(S, \varepsilon)$ is disjoint from X. It remains to show that C is ε-linked. Because ε has countable combined degree, [77, Lemma 5.1] yields an ε-*defining* sequence $(S_n)_{n\in\mathbb{N}}$, that is a sequence of finite sets $S_n \subseteq V(G)$ such that $C(S_n, \varepsilon) \supseteq C(S_{n+1}, \varepsilon)$, $S_n \cap S_{n+1} \subseteq \mathrm{Dom}(\varepsilon)$ and $\bigcap_{n\in\mathbb{N}} C(S_n, \varepsilon) = \emptyset$. It suffices to find an ε-defining sequence $(S'_n)_{n\in\mathbb{N}}$ with $S'_0 = S$ such that there are $|S'_n|$ pairwise disjoint S'_n–S'_{n+1} paths for every $n \in \mathbb{N}$.

So set $S'_0 := S$. Assume that we have constructed $S'_0, \dots, S'_n$ for some $n \in \mathbb{N}$. Since the $C(S_n, \varepsilon)$ are $\subseteq$-decreasing and their intersection is empty, we may choose $N \in \mathbb{N}$ sufficiently large such that $S'_n \cap C(S_N, \varepsilon) = \emptyset$. Then we choose S'_{n+1} as a $|\cdot|$-minimal S_{N+1}–ε separator; in particular, $C(S'_n, \varepsilon) \supseteq C(S_N, \varepsilon) \supseteq C(S_{N+1}, \varepsilon) \supseteq C(S'_{n+1}, \varepsilon)$, and thus $S'_n \cap S'_{n+1} \subseteq S_N \cap S_{N+1} \subseteq \mathrm{Dom}(\varepsilon)$. Note that every S'_n–S'_{n+1} separator would have been a suitable choice for S'_n and thus has size at least $|S'_n|$. By Menger's Theorem (see for example [41, Proposition 8.4.1]), the desired paths exist.

(ii): [77, Lemma 5.1] yields a region C' disjoint from X in which ε lives and whose finite neighbourhood $N(C')$ shares with X only vertices in $\mathrm{Dom}(\varepsilon)$. Now applying (i) to a $|\cdot|$-minimal $N(C')$–ε separator yields the desired ε-linked region $C := C(S, \varepsilon) \subseteq C'$. □

In order to get a candidate for C_i we will use the region from Lemma 3.7.4 to obtain one which is additionally nested with $\mathcal{D}$ and with the C_j for $j < i$. For this, we first need one auxiliary lemma.

A region C of a graph G is *well linked* if, for every two disjoint finite $X, Y \subseteq N_G(C)$, there is a family of $\min\{|X|, |Y|\}$ pairwise disjoint X–Y paths in G through C (i.e. all internal vertices contained in C).

Lemma 3.7.5. *Let H be a graph and let C be an end-linked region of H. Then C is well linked.*

Proof. Let $X, Y \subseteq N(C)$ be disjoint and finite, and suppose for a contradiction that there is no family of $\min\{|X|, |Y|\}$ pairwise disjoint X–Y paths through C. By Menger's theorem (see for example [41, Proposition 8.4.1]), there is an X–Y separator S of size less than $\min\{|X|, |Y|\}$ in $H[V(C) \cup X \cup Y]$. Since C is end-linked, there is some end ε of H which lives in C and a family $\{R_v \mid v \in N(C)\}$ of pairwise disjoint $N(C)$–ε paths and rays with $v \in R_v$. Since $|S| < \min\{|X|, |Y|\}$ and the R_v are pairwise disjoint, there is $x \in X$ and $y \in Y$ such that R_x and R_y both avoid S. Since R_x and R_y are $N(C)$–ε paths or rays, there is some R_x–R_y path P in C which avoids the finite set S. Hence, $R_x + P + R_y$ is a connected subgraph of $H[V(C) \cup X \cup Y]$ which meets X and Y but avoids S. This contradicts the fact that S is an X–Y separator in $H[V(C) \cup X \cup Y]$. □

Let ε be an end of H that does not live in any C_j for $j < i$ or in any $D \in \mathcal{D}$ and which can be separated from X by fewer than $|X|$ vertices. Then our next lemma yields an ε-linked region which is a candidate for the region C_i in Case A.

Lemma 3.7.6. *Let H be a graph and let $\mathcal{E}$ be a set of pairwise non-touching well-linked $(< \aleph_0)$-regions. If C is a k-region with $k \in \mathbb{N}$ which is ε-linked for some end ε that only lives in $D \in \mathcal{E}$ which are contained in C, then there exists an ε-linked $(\leqslant k)$-region C' that is nested with all $D \in \mathcal{E}$.*

Moreover, the set $\mathcal{E}'$ of regions $D \in \mathcal{E}$ which are not nested with C is finite, the region C' is not only nested with all $D \in \mathcal{E}$ but contains D or does not touch D, and C' can be chosen such that $C' \subseteq C \cup \bigcup_{D \in \mathcal{E}'} D$ and $N_H(C') \subseteq N_H(C) \cup \bigcup_{D \in \mathcal{E}'} N_H(D)$.

Proof. We first show that $\mathcal{E}'$ is finite. For this, it suffices to show that each of the pairwise disjoint regions in $\mathcal{E}'$ meets the finite set $N(C)$. So let $D \in \mathcal{E}'$ be given. Because D and C are not nested, they touch, that is, the closure $\bar{C}$ meets D. Moreover, D is not contained in C, so $D - C$ is non-empty. Thus, since D is connected, there is a $(D - C)$–$\bar{C}$ path in D. Its endvertex in $\bar{C}$ is in $N(C)$, as $D - C$ and C are disjoint. Thus, $D \in \mathcal{D}$ meets $N(C)$.

Let $\mathcal{E}'_<$ consist of all $D \in \mathcal{D}'$ with $|V(D) \cap N(C)| < |V(C) \cap N(D)|$. Now $C^* := G[V(C) \cup \bigcup_{D \in \mathcal{E}'_<} V(D)]$ has neighbourhood $(N(C) \smallsetminus \bigcup_{D \in \mathcal{E}'_<} V(D)) \cup \bigcup_{D \in \mathcal{E}'_<} (N(D) \smallsetminus V(\bar{C}))$, since the $D \in \mathcal{E}$ are pairwise non-touching. We claim that C^* is a ε-linked $(\leqslant k)$-region.

The subgraph C^* is connected and thus a region since every $D \in \mathcal{E}' \supseteq \mathcal{E}'_<$ touches C. For $D \in \mathcal{E}'_<$, the set $V(D) \cap N(C)$, which separates $N(D) \smallsetminus V(\bar{C})$ and $V(C) \cap N(D)$ in $\bar{D}$, must have at least size $|N(D) \smallsetminus V(\bar{C})|$, since D is well linked by assumption and $|V(D) \cap N(C)| < |V(C) \cap N(D)|$ by definition of $\mathcal{E}'_<$. Hence, $|N(D) \smallsetminus V(\bar{C})| \leqslant |V(D) \cap N(C)| < |V(C) \cap N(D)|$ for every $D \in \mathcal{E}'_<$ yields that the size of the neigbourhood of C^* is at most k and a family $\mathcal{P}_D$ of $|N(D) \smallsetminus V(\bar{C})|$ pairwise disjoint $(N(D) \smallsetminus V(\bar{C}))$–$(V(C) \cap N(D))$ paths through D. We claim that C^* is ε-linked. Indeed, since the $D \in \mathcal{E}$ are pairwise non-touching, all these paths in $\mathcal{P}_D$ with $D \in \mathcal{E}$ and the trivial paths in $N(C) \smallsetminus (\bigcup_{D \in \mathcal{E}'_<} V(D))$ are pairwise disjoint. Hence, we obtain the desired $N(C^*)$–ε paths and rays by extending the collection of all these paths and the trivial paths in $N(C) \setminus (\bigcup_{D \in \mathcal{E}'_<} V(D))$ via the original $N(C)$–ε paths or rays witnessing the ε-linkedness of C.

By definition of C^* and since the $D \in \mathcal{E}$ are pairwise non-touching, we have $C^* \cap \bar{D} = C \cap \bar{D}$, $V(C^*) \cap N(D) = V(C) \cap N(D)$ and $N(C^*) \cap V(D) = N(C) \cap V(D)$ for every $D \in \mathcal{E} \setminus \mathcal{E}'_<$. Hence, the regions $D \in \mathcal{E}$ which are not nested with C^* are precisely those in $\mathcal{E}^* := \mathcal{E}' \setminus \mathcal{E}'_<$ and every $D \in \mathcal{E}^*$ satisfies $|V(D) \cap N(C^*)| = |V(D) \cap N(C)| \geqslant |V(C) \cap N(D)| = |V(C^*) \cap N(D)|$. Hence, the size of

$$B := \Big(N(C^*) \setminus (\bigcup_{D \in \mathcal{E}^*} V(D))\Big) \cup \bigcup_{D \in \mathcal{E}^*} (V(C^*) \cap N(D)),$$

which includes the neighbourhood of $C^* - (\bigcup_{D \in \mathcal{E}^*})\bar{D}$, is at most $|N(C^*)| \leqslant k$. Since ε lives only in $D \in \mathcal{E}$ which are contained in $C \subseteq C^*$ and thus such $D \notin \mathcal{E}' \supseteq \mathcal{E}^*$, the end ε lives in $C^* - (\bigcup_{D \in \mathcal{E}^*} \bar{D})$. So the ε-linkedness of C^* and $|B| \leqslant |N(C^*)|$ yields that B is linked to ε. All in all, the component $C' := C_H(B, \varepsilon)$ is the desired ε-linked ($\leqslant k$)-region with $C' \subseteq (C \cup \bigcup_{D \in \mathcal{E}'_<} D) - (\bigcup_{D \in \mathcal{E}^*} \bar{D}) \subseteq C \cup \bigcup_{D \in \mathcal{E}'} D$ and $N(C') = B \subseteq N(C) \cup \bigcup_{D \in \mathcal{E}'} N(D)$. Note that by definition C' does not touch $D \in \mathcal{E}^*$. Since every other $D \in \mathcal{E} \setminus \mathcal{E}^*$ either does not touch C^* or is contained in C^* and also does not touch any $D' \in \mathcal{E}^*$, the region D does not touch C' or is contained in C'. □

The next lemma yields for every end ε that does not already live in some C_j for $j < i$ or in some $D \in \mathcal{D}$ an ε-linked region which is a candidate for the region C_i in Case B.

Lemma 3.7.7. *Let H be a graph, let $Z \subseteq V(H)$ be finite, and let $\mathcal{E}$ be a set of pairwise non-touching well-linked ($< \aleph_0$)-regions of H.*

If ε is an end of H of countable combined degree that lives in no $D \in \mathcal{E}$, then there exists an ε-linked ($< \aleph_0$)-region C that is disjoint from Z, satisfies $Z \cap N_H(C) \subseteq \mathrm{Dom}(\varepsilon)$ and is nested with all $D \in \mathcal{E}$. Moreover, every region in $\mathcal{E}$ that touches C is contained in C.

Proof. By Lemma 3.7.4 (ii), there exists an ε-linked ($< \aleph_0$)-region C_0 of H that is disjoint from Z and satisfies $Z \cap N(C_0) \subseteq \mathrm{Dom}(\varepsilon)$. Let $\mathcal{E}'$ be the set of all regions in $\mathcal{E}$ that are not nested with C_0. The 'moreover'-part of Lemma 3.7.6 ensures that $\mathcal{E}'$ is finite. Thus, Lemma 3.7.4 (ii) applied to the finite set $Z' := Z \cup N(C_0) \cup \bigcup_{D \in \mathcal{E}'} N(D)$ and the end ε yields an ε-linked region C_1 that is disjoint from Z' and satisfies $Z \cap N(C_1) \subseteq \mathrm{Dom}(\varepsilon)$. Since ε lives in no $D \in \mathcal{E} \supseteq \mathcal{E}'$, the fact that C_1 is disjoint from Z' yields that $C_1 \subseteq C_0 - (\bigcup_{D \in \mathcal{E}'} \bar{D})$.

Note that, since C_1 is nested with all regions in D that were not nested with C_0, every $D \in \mathcal{E}$ which touches C_1 is contained in C_0. In particular, $Z \cap N(D) \subseteq Z \cap N(C_0) \subseteq \mathrm{Dom}(\varepsilon)$ for all such $D \in \mathcal{D}$ as C_0 avoids Z. Now Lemma 3.7.6 yields a ε-linked ($< \aleph_0$)-region C_2 that is disjoint from Z, is nested with $\mathcal{E}$ and satisfies $Z \cap N(C_2) \subseteq Z \cap N(C_0) \subseteq \mathrm{Dom}(\varepsilon)$. Thus $C := C_2$ is the desired region.

For the ‘moreover’-part we note that ε lives in C but in no $D \in \mathcal{E}$, and thus every $D \in \mathcal{E}$ that touches C and that is nested with C is contained in C. Hence, the ‘moreover’-part follows since C is nested with all $D \in \mathcal{E}$. □

3.7.3 Analysis of Algorithm 3.7.2

In this section we analyse Algorithm 3.7.2 and provide in Theorem 3.7.9 the properties of the regions C_i and the set Y obtained from Algorithm 3.7.2 which we need for the proof of Theorem 3.7.1. First let us take note of some basic properties of Algorithm 3.7.2 that follow easily from its definition:

Observation 3.7.8. *In the setting of Algorithm 3.7.2, we have that*

(O1) *every C_i with $|N_H(C_i)| < k$ was chosen in Case A and every C_i with $|N_H(C_i)| \geqslant k$ was chosen in Case B,*

(O2) *if $\ell_i \geqslant k$, then $X \cap N_H(C_i) \subseteq \mathrm{Dom}(C_i)$,*

(O3) *the ℓ_i are increasing,*

(O4) *a C_i is never contained in C_j with $j < i$,*

(O5) *if $C_j \subseteq C_i$ with $j < i$, then $\ell_j < \ell_i$,*

(O6) *Algorithm 3.7.2 terminates at some ordinal $< |H|^+$.*

Proof. (O1): It is immediate from Algorithm 3.7.2 that every $(\geqslant k)$-region was chosen in Case B. Every $(\leqslant k)$-region as in Case B is already a region as in Case A; thus, every $(< k)$-region C_i was chosen in Case A.

(O2): This is immediate from Algorithm 3.7.2 and (O1).

(O3): By (O1), every C_i with $|N(C_i)| < k$ was chosen in Case A. Moreover, every C_i that was chosen in Case A was already a candidate for C_j in earlier steps $j < i$ where C_j was chosen due to Case A, but was not the chosen nicest such region (and similar for Case B). Thus (N1) ensures that the ℓ_i are increasing.

(O4): A C_i will never be a subgraph of any C_j with $j < i$ since the end to which $N(C_i)$ is linked lives in C_i but not in C_j by the choice of C_i in Algorithm 3.7.2.

(O5): By (O3) we have $\ell_j \leqslant \ell_i$. So suppose for a contradiction that $\ell_j = \ell_i$. By Algorithm 3.7.2, C_i was already a candidate for C_j. Since $\ell_i = \ell_j$, the region C_i satisfied (N1) in step j. But then $C_j = C_i$ by (N2) of C_j, which contradicts that the end to which $N(C_i)$ is linked does not live in C_j by Algorithm 3.7.2.

(O6): It suffices to show that C_i contains some vertex which is in no other C_j for $j < i$. If all C_j for $j < i$ are disjoint from C_i, then we are done. So we may assume that there is some $C_j \subseteq C_i$ with $j < i$. Consider such a $\subseteq$-maximal C_{j^*}; it exists, since chains of such regions $C_j \subseteq C_i$ are finite by (O5). If $N(C_{j^*}) \subseteq N(C_i)$, then the connectedness of C_i and C_{j^*} yields $C_i = C_{j^*}$ which contradicts their distinctness. Thus, there exists a vertex $v \in V(C_i) \cap N(C_{j^*})$. Since the C_j for $j < i$ are nested and by the $\subseteq$-maximality of C_{j^*}, the region C_{j^*} does not touch C_j or contains C_j for all $j \neq j^*$ with $C_j \subseteq C_i$ and $j < i$; in particular, in both cases $V(C_j) \cap N(C_{j^*}) = \emptyset$. Thus, $v \in C_i \setminus \bigcup_{j<i} C_j$ as desired. □

Using the following properties of Algorithm 3.7.2, we show in Section 3.7.4 that the iterative application of Algorithm 3.7.2 yields the desired tree-decomposition for Theorem 3.7.1:

Theorem 3.7.9. *In the setting of Algorithm 3.7.2 the following statements hold:*

(A1) *Either $x \in Y$ or $|X| > |N_H(C)|$ for the component C of $H - Y$ containing x.*

(A2) *Every $D \in \mathcal{D}$ is contained in a component of $H - Y$.*

Moreover, if every end of H has countable combined degree, then also the following hold:

(A3) *Every end of H lives in some C_i.*

(A4) *Every component of $H - Y$ is a C_i.*

(A5) *If $|N_H(C)| \geqslant k$ for a component C of $H - Y$, then $X \cap N_H(C) \subseteq \mathrm{Dom}(C)$.*

(A6) *Set $H^0 := H$ and $Y^0 := Y$. Let $C'_1 \supseteq C'_2 \supseteq \dots \supseteq C'_n$ be a sequence of regions of H such that, for every $m \in \{0, \dots, n-1\}$, the region C'_{m+1} is a component of $H^m - Y^m$ where, for $m \in \{1, \dots, n-1\}$, Y^m is given by Algorithm 3.7.2 applied to $H^m := \bar{C}'_m$, the finite set $X^m := N_H(C'_m)$, the set $\mathcal{D}^m := \{D \in \mathcal{D} \mid D \subseteq C'_m\}$ and an arbitrary vertex $x_m \in C'_m - (\bigcup \mathcal{D})$. Then there are $\min\{|N_H(C'_m)| \mid m \in \{1, \dots, n\}\}$ pairwise disjoint X–$N_H(C'_n)$ paths in H.*

(A7) *If $H - X$ is connected and $N_H(H - X) = X$, then $N_H(C) \not\subseteq X$ for all components C of $H - Y$; in particular, $X \subsetneq Y$.*

Proof. (A1): Assume that $x \notin Y$. Let $\mathcal{C}_x$ be the collection of all those C_i that contain x. Then $\mathcal{C}_x$ is non-empty since by the definition of Y at least one C_i contains x. We claim that $C := \bigcup \mathcal{C}_x$ is a component of $H - Y$ with $|N(C)| < k = |X|$. Indeed, it is immediate from Algorithm 3.7.2 that every $C_i \in \mathcal{C}_x$ is a region as in Case A, since the regions that where chosen in Case B avoid x; in particular, every $C_i \in \mathcal{C}_x$ is a $(< k)$-region. Since all the $C_i \in \mathcal{C}_x$ meet in x and are nested, $\mathcal{C}_x$ is a chain (with respect to inclusion). Thus, Lemma 3.7.3 ensures that C is a region of H with $|N(C)| < k = |X|$.

To finish the proof of (A1), it thus remains to show that C is a component of $H - Y$. For this, it suffices to prove that $N(C) \subseteq Y$ since C is a region and hence connected. As $C = \bigcup \mathcal{C}_x$, we have $N(C) \subseteq \bigcup \{N(C_i) \mid C_i \in \mathcal{C}_x\}$. Since all the C_i are nested, every C_j which meets the neighbourhood of some $C_i \in \mathcal{C}_x$ already contains C_i; in particular, such C_j are also in $\mathcal{C}_x$ as $x \in C_i \subseteq C_j$. Thus $N(C)$ meets no C_i, and so the definition of Y yields that $N(C) \subseteq Y$.

(A2): By the assumptions on $\mathcal{D}$ and the choice of the C_i, in every step i of Algorithm 3.7.2 every $D \in \mathcal{D}$ is a candidate for the region C_i in Case B as long as D is ε-linked for some end ε of H

which lives in no C_j for $j < i$. Again by the assumptions on $\mathcal{D}$, the region $D \in \mathcal{D}$ is ε-linked for some end ε of H. So, since Algorithm 3.7.2 terminates because of Case C by Observation 3.7.8 (O6), the end ε lives is some C_i. As ε lives in both C_i and D, they touch. Thus, $C_i \subseteq D$ or $D \subseteq C_i$ since they are nested by Algorithm 3.7.2. Hence, we finish the proof of (A2) by showing that if $C_i \subseteq D$, then $C_i = D$. For this, it suffices to show that in step i of Algorithm 3.7.2, the region D was a candidate for C_i and satisfies (N1), because then $C_i \subseteq D$ together with (N2) yields $C_i = D$.

It is immediate from Algorithm 3.7.2 and the assumptions on $\mathcal{D}$ that in step i, the region $D \in \mathcal{D}$ is nested with all C_j for $j < i$, disjoint from X and that $X \cap N(D) \subseteq \mathrm{Dom}(D)$. Since D is ε-linked and ε lives in $C_i \subseteq D$, there are $|N(D)|$ disjoint $N(D)$–$N(C_i)$ paths in $\bar{D} - C_i$. These paths can be extended to disjoint $N(D)$–ε' paths and rays where ε' is the end so that C_i is ε'-linked. In particular, D is end-linked to the end ε' which lives in no C_j for $j < i$, and $\ell_i = |N(C_i)| \geqslant |N(D)|$. Thus, D was a candidate for C_i and satisfies (N1), as desired.

(A3): Suppose towards a contradiction that there is an end ε of H that lives in no C_i. For every $\ell \in \mathbb{N}$, let $\mathcal{C}_{\max}^{<\ell}$ be the set of all C_i with $\ell_i < \ell$ that are $\subseteq$-maximal among all $(<\ell)$-regions C_j. By Observation 3.7.8 (O5) every C_i with $\ell_i < \ell$ is included in some $C_j \in \mathcal{C}_{\max}^{<\ell}$. Moreover, since all C_i are nested, the regions in $\mathcal{C}_{\max}^{<\ell}$ are pairwise non-touching. Hence, by Lemma 3.7.7 applied to $Z := X \cup \{x\}$ and the set $\mathcal{E}$ of all $\subseteq$-maximal elements in $\mathcal{D} \cup \mathcal{C}_{\max}^{<k}$, there exists a ε-linked $(<\aleph_0)$-region C that is disjoint from $X \cup \{x\}$, satisfies $(X \cup \{x\}) \cap N(C) \subseteq \mathrm{Dom}(\varepsilon)$ and is nested with all $\subseteq$-maximal regions in $\mathcal{D} \cup \mathcal{C}_{\max}^{<k}$. Set $\ell := |N(C)|$. Then Lemma 3.7.6 applied to C and the set $\mathcal{E}$ of all $\subseteq$-maximal regions in $\mathcal{D} \cup \mathcal{C}_{\max}^{<\ell+1}$ yields a ε-linked $(\leqslant \ell)$-region C' which is nested with all regions in $\mathcal{D} \cup \mathcal{C}_{\max}^{<\ell+1}$ and such that $(X \cup \{x\}) \cap N(C') \subseteq \mathrm{Dom}(C')$, where we have used that $(X \cup \{x\}) \cap N(D) \subseteq \mathrm{Dom}(D)$ for all $D \in \mathcal{C}_{\max}^{<\ell+1} \smallsetminus \mathcal{C}_{\max}^{<k}$ since any such region D was chosen in Case B. Moreover, by Lemma 3.7.6, C' is not strictly contained in any region in $\mathcal{D} \cup \mathcal{C}_{\max}^{<\ell+1}$. In particular, C' is nested with all $D \in \mathcal{D}$ and all C_i with $\ell_i \leqslant \ell$. But this contradicts Algorithm 3.7.2: If there is no C_i with $\ell_i > \ell$, then C' contradicts that Algorithm 3.7.2 terminated because of Case C by Observation 3.7.8 (O6). Otherwise, if there exists some C_i with $\ell_i > \ell$, then C' witnesses that the first such C_i did not satisfy (N1).

(A4): Let C be a component of $H - Y$. By Zorn's Lemma there exists a $\subseteq$-maximal index set I such that $C_j \subseteq C_i \subseteq C$ for all $j \leqslant i \in I$. Then $\bigcup_{i \in I} C_i = C$. Indeed, if $C' := \bigcup_{i \in I} C_i \subsetneq C$, then $N(C') \cap C \neq \emptyset$ since C is connected. So by the definition of Y, there is some C_j such that $N(C') \cap C_j \neq \emptyset$. In particular, $C_j \subseteq C$ since C_j is connected. But since $N(C') \subseteq \bigcup_{i \in I} N(C_i)$, the $(C_i \mid i \in I)$ are $\subseteq$-increasing and all the C_i are nested, we have $C_i \subseteq C_j$ for all $i \in I$, which contradicts that I is $\subseteq$-maximal. By Observation 3.7.8 (O3), the sequence $(\ell_i)_{i \in I}$ is strictly increasing. In particular, I is either finite or of the same order type as $\mathbb{N}$. In the former case we are done as $C = C_i$ for $i = \max(I)$. So assuming the latter, we now aim towards a contradiction. Enumerate $I = \{i_n : n \in \mathbb{N}\}$ so that $i_n < i_m$ for all $n < m \in \mathbb{N}$.

Consider the auxiliary graph A that arises from C by contracting C_{i_1} and, for every $n \in \mathbb{N}$, all components of $C_{i_{n+1}} - C_{i_n}$. The graph A is obviously infinite and connected. Since the $N(C_{i_n})$ are finite and the C_{i_n} are connected, there are at most finitely many components of $C_{i_{n+1}} - C_{i_n}$ for every $n \in \mathbb{N}$. Together with the fact that the $N(C_{i_n})$ are finite this yields that A is also locally finite. Hence, the minor A of C contains a ray by [41, Proposition 8.2.1]. We lift this ray to a ray R in C by choosing suitable paths in each branch set connecting the endvertices of the incident edges. Then by construction, the end of H that contains R lives in no C_i, contradicting (A3).

(A5): This follows directly from (A4) and (O2).

(A6): Set $k' := \min\{|N(C'_m)| \mid m \in \{1, \dots, n\}\}$. By Menger's Theorem (see for example [41, Proposition 8.4.1]), we either find the desired k' pairwise disjoint X–$N(C'_n)$ paths or there is an X–$N(C'_n)$ separator S with $|S| < k'$. We may assume the latter and let S be a $|\cdot|$-minimal such separator. Further, let ε be an end such that C'_n is ε-linked, which exists by Algorithm 3.7.2 and (A4). Since C'_n is ε-linked, the separator S is also a $|\cdot|$-minimal X–ε separator. Then the component $\tilde{C}$ of $H - S$ in which ε lives is ε-linked and disjoint from X by Lemma 3.7.4 (i). Moreover, we have $C'_n \subseteq \tilde{C}$ since S avoids C'_n by its choice as a minimal X–$N(C'_n)$ separator.

Now let C be a ε-linked ($\leqslant |S|$)-region that is disjoint from X, contains C'_n and is contained in as many of the C'_i as possible. Note that $\tilde{C}$ is a candidate for C. Let $N \geqslant 0$ be the largest index i such that $C \subseteq C'_N$, and let C^N_i be the regions obtained from the application of Algorithm 3.7.2 to H^N and $N(C'_N)$ (or H and X if $N = 0$). Note that $N < n$ since otherwise we have $C = C'_n$ and hence $|N(C'_n)| = |N(C)| < k'$, a contradiction. Let $\mathcal{C}^{\leqslant |S|}_{\max}$ be the $\subseteq$-maximal elements of the C^N_i with $\ell^N_i \leqslant |S|$. Note that by Observation 3.7.8 (O5) every C^N_i with $\ell^N_i < |S|$ is contained in some $C^N_j \in \mathcal{C}^{\leqslant |S|}_{\max}$. Since all regions in $\mathcal{C}^{\leqslant |S|}_{\max} \cup \mathcal{D}$ are end-linked, they are well linked by Lemma 3.7.5. Moreover, if ε lives in $D \in \mathcal{D}$, then $D \subseteq C'_n \subseteq C$ by (A2). To be able to apply Lemma 3.7.6 to C and the set $\mathcal{E}$ of $\subseteq$-maximal regions in $\mathcal{C}^{\leqslant |S|}_{\max} \cup \mathcal{D}$ it suffices that ε lives in no C^N_j with $\ell^N_j \leqslant |S|$.

So suppose for a contradiction that ε lives in some C^N_j with $\ell^N_j \leqslant |S|$. We claim that then $C'' := C^N_j \cup C'_n$ is an ($\leqslant |S|$)-region. Indeed, C'' is a region because C^N_j and C'_n are both connected and intersect as ε lives in both of them. Moreover, $N(C'') = (N(C^N_j) \cup N(C'_n)) \setminus V(C'') = (N(C^N_j) \setminus V(C'_n)) \dot\cup (N(C'_n) \setminus V(\bar{C}^N_j))$ has size at most $|N(C^N_j)| = \ell^N_j \leqslant |S|$: if not, one easily checks by doubling counting that $Z := ((N(C^N_j) \cup N(C'_n)) \cap V(C'')) \cup (N(C^N_j) \cap N(C'_n)) = (N(C'_n) \cap V(C^N_j)) \dot\cup (N(C^N_j) \cap V(C'_n)) \dot\cup (N(C^N_j) \cap N(C'_n))$ has size less than $|N(C'_n)|$. But this contradicts that C'_n is ε-linked as ε lives in a component of $C^N_j \cap C'_n$, whose neighbourhood then is contained in Z, and so all $N(C'_n)$–ε rays and paths have to meet Z. Hence, C'' is a ($\leqslant |S|$)-region with $C'_n \subseteq C''$. Moreover, C^N_j and C'_{N+1} both contain C'_n by assumption and thus touch. So since C^N_j and C'_{N+1} are nested by Algorithm 3.7.2 (because C'_{N+1} is some C^N_i by (A4)), Observation 3.7.8 (O4) implies that $C^N_j \subseteq C'_{N+1}$, and therefore $C'' \subseteq C'_{N+1}$. But this contradicts our choice of C. Thus, ε lives in no region in $\mathcal{C}^{\leqslant |S|}_{\max} \cup \mathcal{D}$.

So we may apply Lemma 3.7.6 to C and the set $\mathcal{E}$ of $\subseteq$-maximal separations in $\mathcal{C}_{\max}^{\leqslant|S|} \cup \mathcal{D}$ to obtain an ε-linked $(\leqslant|S|)$-region C' which is disjoint from X and which, for every region in $\mathcal{C}_{\max}^{\leqslant|S|} \cup \mathcal{D}$, either contains that region or does not touch it. Thus, C' is nested with all C_i^N with $\ell_i^N \leqslant |S|$ and with all $D \in \mathcal{D}$. This concludes the proof since the existence of C' then contradicts that C_I^N satisfies (N1) where $I = \min\{i \mid \ell_i^n > |S|\}$. Here we used that C'_{N+1} is some C_i^N and that $|N(C'_{N+1})| \geqslant k' > |S|$, so the minimum is not taken over the empty set.

(A7): Let C be a component of $H - Y$. By (A4), we have $C = C_i$ for some suitable i. Suppose first that C_i that was chosen in Case A. Then $|N_H(C)| < |X|$; since $H - X$ is connected and $N_H(H - X) = X$, $N_H(C)$ contains a vertex in $V(H) \setminus X$, implying $N_H(C) \not\subseteq X$. Now suppose that C_i that was chosen in Case B. Then $N_H(C)$ separates x and C. Since $H - X$ is connected and $x \notin X$, we have $N_H(C) \not\subseteq X$. The in-particular part is now also clear. □

3.7.4 Proof of Theorem 3.7.1

Using Theorem 3.7.9 we now show that the tree-decomposition obtained from iteratively applying Algorithm 3.7.2 is as desired for Theorem 3.7.1.

Proof of Theorem 3.7.1. First, we define the desired tree-decomposition $(T, \mathcal{V})$ in the case where $X = \operatorname{int}(\sigma)$. The decomposition tree T is a star whose edges are in a bijective correspondence to σ. We assign to the centre the bag X, and to each leaf ℓ the bag A where $(A, B) \in \sigma$ corresponds to the edge incident with ℓ. It is immediate to see that this tree-decomposition is as desired. Let us assume from now on that $X \subsetneq \operatorname{int}(\sigma)$.

Suppose first that G is connected. By Theorem 3.2.2, we may fix a normal spanning tree T_{NST} of G whose root is in X. We denote by $\mathcal{D}$ the set of subgraphs $G[A \setminus B]$ of G with $(A, B) \in \sigma$. The assumptions on σ ensure that $\mathcal{D}$ satisfies the assumptions in Algorithm 3.7.2. Since G admits a normal spanning tree, all its ends have countable combined degree. Hence, G satisfies the assumptions of Theorem 3.7.9. We now define the desired tree-decomposition recursively as follows.

Let $(T^0, \mathcal{V}^0)$ be the trivial tree-decomposition of G where T^0 is the tree on a single vertex r which is also its root and $V_r^0 := V(G)$. Now let $n \geqslant 0$ and suppose that we have already constructed linked, X-linked, tight, componental, rooted tree-decompositions $(T^m, \mathcal{V}^m)$ of finite adhesion such that for all $m \leqslant n$

(i) $T^m \subseteq T^n$ and $V_t^m = V_t^n$ for all $t \in T^m$ that are not at height m, the decomposition tree T^m has height m, all their leaves are on height m except possibly those whose corresponding leaf separation is a (B, A) with $(A, B) \in \sigma$,

(ii) the torsos of $(T^m, \mathcal{V}^m)$ at non-leaves are rayless, and $(T^m, \mathcal{V}^m)$ satisfies (R2') and (R3').

Note that $(T^0, \mathcal{V}^0)$ satisfies all these properties immediately, where we remark that we treat the unique node of T^0 as a leaf. Let L be the set of leaves of T^n whose corresponding leaf separation is not some (B, A) with $(A, B) \in \sigma$. Note that, if $n = 0$, then $L = \{r\}$. For every leaf $\ell \in L$, we construct a tree-decomposition $(T^\ell, \mathcal{V}^\ell)$ of $G[V^n_\ell]$ as follows. Set $X' := X$ if $n = 0$ and ℓ is the unique node of T^0, and $X' := V^n_f$ otherwise where f is the unique edge of T^n that is incident with ℓ. Further, let x be a $(\leqslant_{T_{NST}})$-minimal vertex in $V^n_\ell \setminus (X' \cup \bigcup \mathcal{D})$. Note that, if $n = 0$, x exists (but might not be unique) as $X \subsetneq \operatorname{int}(\sigma)$. If $n > 0$, the vertex x exists and is unique, since T_{NST} is normal and because $(T^n, \mathcal{V}^n)$ is componental by construction and so $G[V^n_\ell] - (X' \cup \bigcup \mathcal{D})$ is connected by the assumption on σ.

Now apply Algorithm 3.7.2 to $H := G[V^n_\ell]$ with the finite set X', the vertex x and the set $\{D \in \mathcal{D} \mid D \subseteq H\}$ to obtain $Y \subseteq V(H)$. Then let T^ℓ be the star with centre ℓ and whose set of leaves is the set $\mathcal{C}$ of all components of $H - Y$. Further, set $V^\ell_\ell := Y$ and $V^\ell_C := \bar{C}$ for every $C \in \mathcal{C}$. Then Theorem 3.7.9 (A4) yields that for every edge $e = \ell C$ of T^ℓ, the graph $H \mathring{\uparrow} e$ is a C_i in Algorithm 3.7.2 applied to H. Thus, $(T^\ell, \mathcal{V}^\ell)$ is componental and end-linked; thus, $(T^\ell, \mathcal{V}^\ell)$ satisfies (R2'). Moreover, the torso at ℓ is rayless: since $(T^\ell, \mathcal{V}^\ell)$ is tight, we can obtain from any ray in the torso at ℓ a ray in G that meets V^ℓ_ℓ infinitely often by (the comment after) Proposition 3.6.1. But this contradicts Theorem 3.7.9 (A3). Furthermore, it satisfies the following (which will ensure (R3') of $(T^{n+1}, \mathcal{V}^{n+1})$) by Theorem 3.7.9 (A5): if $|V^\ell_e| \geqslant |V^n_f|$ for an edge $e \in T^\ell$, then $V^\ell_e \cap V^n_f \subseteq \operatorname{Dom}(H \uparrow e)$.

By construction, the T^ℓ are pairwise disjoint and only share their centre ℓ with T^n. We set $T^{n+1} := T^n \cup \bigcup_{\ell \in L} T^{n+1}_\ell$ with $\operatorname{root}(T^{n+1}) := \operatorname{root}(T^n) (= r)$, and bags $V^{n+1}_t := V^n_t$ for every node $t \in T^n - \ell$ and $V^{n+1}_t := V^\ell_t$ for every node $t \in T^{n+1} \setminus T^n$ where ℓ is the unique leaf of T^n such that $t \in T^\ell$. In particular, the decomposition tree T^{n+1} has height $n+1$, all its leaves are on height $n+1$, except possibly those whose corresponding leaf separation is (B, A) for some $(A, B) \in \sigma$, and $(T^{n+1}, \mathcal{V}^{n+1})$ satisfies all properties that we demanded from the $(T^m, \mathcal{V}^m)$: All properties but the (X-)linkedness of $(T^{n+1}, \mathcal{V}^{n+1})$ follows immediately from the construction and the discussion of the properties of the $(T^\ell, \mathcal{V}^\ell)$. Its ($X$-)linkedness is ensured by Theorem 3.7.9 (A6).

Let $(T, \mathcal{V})$ be the limit of $(T^n, \mathcal{V}^n)$ for $n \to \infty$, that is $T = \bigcup_{n \in \mathbb{N}} T^n$, $\operatorname{root}(T) := r = \operatorname{root}(T^n)$ for all $n \in \mathbb{N}$ and $V_t := V^N_t = V^n_t$ for all $n \geqslant N$ where N is the minimal number such that $t \in T^N$ and either t is not a leaf of T^N or its leaf separation is a (B, A) with $(A, B) \in \sigma$.

We first show that $(T, \mathcal{V})$ is a tree-decomposition of G. For this, by (i), the definition of $(T, \mathcal{V})$ and since all $(T^n, \mathcal{V}^n)$ are tree-decompositions of G, it suffices to show that every vertex of G is contained in some bag V_t. In other words, we need to show that each $v \in V(G)$ satisfies $(*)$: there are $n \in N$ and $t \in T^n$ such that $v \in V^n_t$ and t is either a non-leaf of T^n are its leaf separation is a (B, A) with $(A, B) \in \sigma$. Let us first assume that $v \notin D$ for all $D \in \mathcal{D}$. Then $(*)$ is ensured by considering Theorem 3.7.9 (A1) along $\operatorname{root}(T_{NST}) T_{NST} v = v_0 v_1 \dots v_m$: Suppose for a contradiction that $(*)$ does not hold for v. Then let i be the minimal index such that v_i does not satisfy $(*)$; in particular, $v_i \in G \mathring{\uparrow} e^n$ for a unique edge e^n incident with a leaf of T^n for every $n \in \mathbb{N}$. We have $i \geqslant 1$, since $\text{root}(T_{NST}) \in X \subseteq V_r^1$ by assumption on T_{NST} and construction of V_r^1. Thus, v_{i-1} satisfies $(*)$; let N be a sufficiently large integer given by $(*)$ of v_{i-1}. Then $v_i = \min_{\leqslant_S}(V(G\mathring{\uparrow}e^n) \smallsetminus (\bigcup V(\mathcal{D})))$ for every $n \geqslant N$. Hence, Theorem 3.7.9 (A1) ensures that the finite sets $V_{e^n}^n$ strictly decrease in their size for $n \geqslant N$, which is a contradiction.

Second, assume that $v \in D$ for some $D \in \mathcal{D}$. By Theorem 3.7.9 (A2), we find a path P ending in a leaf ℓ of T or a ray P in T starting in the root r of T such that $D \subseteq G\mathring{\uparrow}e$ for every $e \in P$. If P is a path, then we find $v \in V_\ell$ as desired. In particular, by construction of $(T, \mathcal{V})$ and since the regions in $\mathcal{D}$ are nested, the leaf separation at ℓ is $(V(G) \smallsetminus V(D), V(D) \cup N(D))$.

But P cannot be a ray. Indeed, applying $(*)$ to the finite set $N(D)$ yields $N \in \mathbb{N}$ such that all $u \in N(D)$ are contained in bags V_t^N at nodes $t \in T^N$ which are either not leaves or their corresponding leaf separations are some (B, A) with $(A, B) \in \sigma$. Then $D \subseteq G\mathring{\uparrow}e_N$ for the N-th edge of P by the definition of P, and $N(D) \subseteq G{\downarrow}e_N$ by the choice of N. Since $(T^N, \mathcal{V}^N)$ is componental, this implies that $V(D) \cup N(D) = V_{\bar{D}}^N$ is a bag at a leaf of $(T^N, \mathcal{V}^N)$, and thus, by the construction of $(T, \mathcal{V})$ also appears as a bag at a leaf of $(T, \mathcal{V})$. In particular, the argument above implies that all (B, A) with $(A, B) \in \sigma$ are leaf separations of $(T, \mathcal{V})$. Since $(T, \mathcal{V})$ has no other leaf separations by construction and (i), it follows that the leaf separations of $(T, \mathcal{V})$ are precisely $\{(B, A) \mid (A, B) \in \sigma\}$.

Now it is immediate from the construction that $(T, \mathcal{V})$ is linked, X-linked, tight, componental, has finite adhesion and satisfies (R1'), (R2') and (R3') because the $(T^n, \mathcal{V}^n)$ are linked, X-linked, tight, componental, have finite adhesion and satisfy (ii).

We are thus left to show the 'moreover'-part. For this, recall that $(T, \mathcal{V})$ is tight and componental, so every $G[V_\ell^n]$ considered in the construction of $(T, \mathcal{V})$, except possibly $G[V_r^0]$, satisfies the premise of (A7). It follows for all edges $e = ts \in T$ with $t \leqslant_T s$ and $t \neq r$ that $V_t \supsetneq V_e \subsetneq V_s$. In particular, if $G - X$ is connected and $N_G(G - X) = X$, then also $G[V_r^0]$ satisfies the premise of (A7), so $V_r \supsetneq V_e \subsetneq t$ for all edges $e = rt \in T$ and $X \subsetneq V_r$.

The proof of the case that G is disconnected is analogous by choosing for each component C of G a normal spanning tree T_{NST}^C and considering the partial order $\leqslant_{NST}$ as the disjoint union of their tree orders.[13] □

3.8 Lean tree-decompositions

In this section we prove Theorem 3, the strengthening of Theorem 1 for graphs without half-grid minor in which we replace the linkedness of the tree-decomposition with leanness. We restate it here in its more detailed version:

[13]The only adaption one has to do is to prove $(*)$ for the roots of every T_{NST}^C.

Theorem 3' (Detailed version of Theorem 3)**.** *Every graph G without half-grid minor admits a lean, cofinally componental, rooted tree-decomposition into finite parts which displays the infinities.*

Moreover, if the tree-width of G is finitely bounded, then the tree-decomposition can be chosen to have width $\mathrm{tw}(G)$*.*

The proof of Theorem 3' is structured as follows. We start our construction with the tree-decomposition $(T, \mathcal{V})$ of the graph G from Theorem 4', and we then aim to refine its (finite) parts $G[V_t]$ via lean tree-decompositions $(T^t, \mathcal{V}^t)$ given by the finite version of Theorem 3', that is Thomas's Theorem 3.1.4. If V_t is a critical vertex set, then we may choose $(T^t, \mathcal{V}^t)$ as the trivial tree-decomposition into one bag. In order to combine such refinements the tree-decompositions $(T^t, \mathcal{V}^t)$ along $(T, \mathcal{V})$, we apply the finite result not to the parts themselves, but to the torsos. Torsos, however, need not have the same tree-width as G. So our second ingredient to the proof of Theorem 3' provides a sufficient condition on the separations induced by the edges incident with a node $t \in T$ which allows us to transfer the tree-width bound from G to the torso of $(T, \mathcal{V})$ at t via Corollary 5.7.3 (see Lemma 3.8.1 below). Moreover, it also yields the well-linkedness of the separations on their left side (see Lemma 3.8.2 below). This property is crucial to the proof of Theorem 3' as it ensures that the tree-decomposition $(T', \mathcal{V}')$ arising from $(T, \mathcal{V})$ and the $(T^t, \mathcal{V}^t)$ by refinement is lean: First, if all separations induced by edges at some node t are left-well-linked, then we can transfer families of disjoint paths from the torso at t to G via Lemma 3.6.3 (ii), which yields the paths families required for lean between bags that belong to the same $(T^t, \mathcal{V}^t)$. Second, whenever the separation induced by an edge $e = t_0t_1$ of T is well linked on both sides, this ensures that the leanness of the two $(T^{t_i}, \mathcal{V}^{t_i})$ combines to the leanness of the tree-decomposition resulting from gluing them together along V_e. This will ensure that we obtain disjoint paths families, as required for lean, also between bags of $(T', \mathcal{V}')$ that belong to distinct tree-decompositions $(T^t, \mathcal{V}^t)$ and $(T^s, \mathcal{V}^s)$. Our third, and last, ingredient to the proof of Theorem 3', then, is a pre-processing step: We first contract all edges which neither satisfy the well-linked condition on both sides nor are incident with a node whose bag is a critical vertex set. This will ensure that the sufficient condition mentioned above is met by all separations induced by edges of T that are incident with a node of T whose torso we need to refine (recall that we do not need to refine those torsos whose bags are critical vertex sets). Combining these three ingredients then yields a tree-decomposition of G, which we prove to be as desired.

We call a finite-order separation (A, B) of a graph G *left-ℓ-robust*[14] for $\ell \in \mathbb{N}$ if there exist a set $U \subseteq A$ of size ℓ and a family $\{P_x \mid x \in A \cap B\}$ of pairwise disjoint paths in $G[A]$ such that P_x ends in x and for each $x \in A \cap B$ there are ℓ many U–P_x paths in $G[(A \setminus B) \cup \{x\}]$ that do not meet outside P_x. Analogously, (A, B) is *right-ℓ-robust* for $\ell \in \mathbb{N}$ if (B, A) is left-ℓ-robust. We call

[14] See Section 5.5 for a motivation and a more intuitive explanation of this property.

(A, B) *ℓ-robust* if it is both left- and right-ℓ-robust.

In Chapter 5 (Corollary 5.7.3) we prove the following:

Lemma 3.8.1. *Let G be a graph of tree-width at most $w \in \mathbb{N}$, and let σ be a finite star of separations of G of order at most $w+1$ whose interior is finite. Suppose that all separations in σ are left-ℓ-robust for $\ell = (w+1)^2(w+2)+w+1$. Then* $\mathrm{torso}(\sigma)$ *has tree-width at most w.*

Lemma 3.8.2. *If a separation of a graph G has order k and is left-$(2k+1)$-robust, then it is left-well-linked.*

Proof. Consider a left-$(2k+1)$-robust separation (A, B) of G, and let $X, Y \subseteq A \cap B$ be disjoint. Suppose towards a contradiction that there is no family of $\min\{|X|, |Y|\}$ disjoint X–Y paths through $A \setminus B$ in G. By Menger's theorem (see for example [41, Proposition 8.4.1]), there then is an X–Y separator S of size less than $\min\{|X|, |Y|\}$ in $G[(A \setminus B) \cup X \cup Y]$.

Now fix a set U and a path family $\{P_x \mid x \in A \cap B\}$ which witness that (A, B) is left-$(2k+1)$-robust, where k is the order of (A, B). Since $|S| < \min\{|X|, |Y|\}$ and the P_x are pairwise disjoint, there are $x \in X$ and $y \in Y$ such that P_x and P_y avoid S. For $z \in \{x, y\}$, at least $k+1$ of the $2k+1$ P_z–U paths in $G[(A \setminus B) \cup \{z\}]$ given by the left-$(2k+1)$-robustness avoid S. So since U has size $2k+1$, there are such a P_x–U path Q_x and such a P_y–U path Q_y that both avoid S and end in the same vertex in U. Hence, $P_x + Q_x + Q_y + P_y$ is a connected subgraph of $G[(A \setminus B) \cup X \cup Y]$ which meets X and Y but avoids S. This contradicts that S is an X–Y separator in $G[(A \setminus B) \cup X \cup Y]$, which completes the proof. □

Let us now turn to the third ingredient for our proof of Theorem 3': the pre-processing step in which we contract certain edges of the tree-decomposition from Theorem 4. This ensures that the assumptions of both Lemmas 3.8.1 and 3.8.2 are met at all edges incident with nodes whose (finite) torso we later aim to refine using Thomas's Theorem 3.1.4.

To simplify the wording, let us thus make the following definitions. Given some $m \in \mathbb{N}_0$, we call a separation of G of order k *left-m-good* if it is ℓ-left-robust for $\ell := \max\{m, 2k+1\}$. Analogously, (A, B) is *right-m-good* for $m \in \mathbb{N}$ if (B, A) is left-m-good. We call (A, B) *m-good* if it is left- and right-m-good. An left-m-good separation of G is *left-good* if $m = (w+1)^2(w+2)+w+1$ for G with $w = \mathrm{tw}(G) \in \mathbb{N}$ or $m = 0$ for graphs G whose tree-width is not finitely bounded. Analogously, we define *right-good.* We call (A, B) *good* if it is both left- and right-good. The bounds in the definition of (left-)good are exactly the ones sufficient to apply Lemmas 3.8.1 and 3.8.2 in the respective contexts.

Setting out from Theorem 4', our pre-processing step yields the following result:

Lemma 3.8.3. *Let G be a graph without half-grid minor, let $m \in \mathbb{N}$ and let $(T, \mathcal{V})$ be a fully tight, rooted tree-decomposition into finite parts which displays the infinities and satisfies* (I1) *from Theorem 4'. Then the tree-decomposition $(T', \mathcal{V}')$ of G induced by contracting every edge st of T with $\deg(s), \deg(t) < \infty$ whose induced separation is not m-good has finite parts, is cofinally componental, and displays the infinities. Moreover,*

(i) *if $\deg(t) < \infty$ for $t \in T'$, then all separations in σ'_t at t are left-m-good.*

We first show three auxiliary lemmas.

Lemma 3.8.4. *Let ε be a finitely dominated end of a graph G, and $\ell \in \mathbb{N}$ arbitrary. For every collection $R_1, \dots, R_d$ of disjoint rays in ε that avoid $\mathrm{Dom}(\varepsilon)$, there exists a set $U \subseteq V(R_1)$ of size ℓ such that*

- *for every $i \in \{2, \dots d\}$, there are ℓ pairwise disjoint U–R_i paths $P_1^i, \dots, P_\ell^i$ in G avoiding $\mathrm{Dom}(\varepsilon)$, and*
- *for every $v \in \mathrm{Dom}(\varepsilon)$, there are ℓ U–v paths $P_1^v, \dots, P_\ell^v$ in G avoiding $\mathrm{Dom}(\varepsilon) \smallsetminus \{v\}$ that only meet in v.*

Proof. Since $R_1, \dots, R_d$ belong to the same end ε of G and $\mathrm{Dom}(\varepsilon)$ is finite, there are, for every $i \in \{2, \dots, d\}$, infinitely many pairwise disjoint R_1–R_i paths $Q_1^i, Q_2^i, \dots$ avoiding $\mathrm{Dom}(\varepsilon)$, and we write q_k^i for the endvertex of Q_k^i in R_1. Similarly, for every $v \in \mathrm{Dom}(\varepsilon)$ there are infinitely many v–R_1 paths $Q_1^v, Q_2^v, \dots$ in G avoiding $\mathrm{Dom}(\varepsilon) \smallsetminus \{v\}$ that only meet in v, and we write q_k^v for the endvertex of Q_k^v in R_1.

We now find the elements $u_1, \dots, u_\ell$ of U one by one: Let u_1 be the first vertex of R_1. Given $u_1, \dots, u_{j-1}$ for some $2 \leqslant j \leqslant \ell$, we then choose u_j as the first vertex on R_1 such that $u_{j-1} R_1 u_j$ contains as inner vertices some q_i^k for every $i \in \{2, \dots, d\}$ and some q_v^k for every $v \in \mathrm{Dom}(\varepsilon)$. Then $U := \{u_1, \dots, u_\ell\}$ is as desired, as witnessed by the $P_j^i := u_j R_1 q_k^i Q_k^i$ and $P_j^v := u_j R_1 q_k^v Q_k^v$ for the respective k as given by the choice of u_j. □

Lemma 3.8.5. *Let $m \in \mathbb{N}$, let $(T, \mathcal{V})$ be a rooted tree-decomposition of finite adhesion of a graph G without half-grid minor, and let ε be an end of G. Assume that ε gives rise to a ray $R = r_0 r_1 \dots$ in T such that $\liminf_{e \in R} |V_e| = \Delta(\varepsilon)$. Then cofinitely many edges e of R with $|V_e| = \Delta(\varepsilon)$ induce m-good separations.*

Proof. Since G has no half-grid minor, the combined degree $\Delta(\varepsilon)$ of ε is finite. Thus, as $(T, \mathcal{V})$ has finite adhesion and ε gives rise to R, we have $\liminf_{e \in R} V_e \supseteq \mathrm{Dom}(\varepsilon)$. In fact, $\liminf_{e \in R} V_e = \mathrm{Dom}(\varepsilon)$ since $\liminf_{e \in R} |V_e| = \Delta(\varepsilon)$, so the V_e eventually have to meet each ray in a family of $\deg(\varepsilon)$ disjoint ε-rays avoiding $\mathrm{Dom}(\varepsilon)$ precisely once. Hence, for some $N_0 \in \mathbb{N}$ all indices $i \geqslant N_0$ satisfy $\mathrm{Dom}(\varepsilon) \subseteq V_{e_i}$ where $e_i := \{r_i, r_{i+1}\}$. Let us denote the set of all indices $i \geqslant N_0$ with $|V_{e_i}| = \Delta_G(\varepsilon) =: k$ by I. Note that I is infinite since $\liminf_{e \in R} |V_e| = \Delta(\varepsilon)$.

Write (A_i, B_i) for the separation induced by the edge $\vec{e}_i = (r_i, r_{i+1})$ for all $i \in \mathbb{N}$. We now show that cofinitely many of the separations (A_i, B_i) with $i \in I$ are m-good, which clearly implies the assertion. So let $\ell := \max\{2k+1, (w+1)^2(w+2)+w+1\}$ if $w := \mathrm{tw}(G) \in \mathbb{N}$, and let $\ell := 2k+1$ if the tree-width of G is not finitely bounded.

Fix a collection $R_1, \dots, R_d$ of $d := \deg(\varepsilon)$ disjoint rays in ε that avoid $\mathrm{Dom}(\varepsilon)$, and consider a set $U \subseteq V(R_1)$ and corresponding paths $P_1^i, \dots, P_\ell^i$ and $P_1^v, \dots, P_\ell^v$ as given by Lemma 3.8.4. Let Z be the union of U and all the $V(P_j^i)$ and $V(P_j^v)$. Then Z is finite. So since $\bigcap_{i \in \mathbb{N}}(B_i \setminus A_i) = \bigcap_{i \in \mathbb{N}} G \mathring{\uparrow} e_i = \emptyset$ and $\bigcap_{i \in \mathbb{N}}(A_i \cap B_i) = \liminf_{i \in \mathbb{N}} V_{e_i} = \mathrm{Dom}(\varepsilon)$, there exists $N_1 \in \mathbb{N}$ such that, for all $i \geqslant N_1$, we have $Z \subseteq (A_i \setminus B_i) \cup \mathrm{Dom}(\varepsilon)$. Consider $i \in I$ with $i \geqslant N_1$. We claim that (A_i, B_i) is left-ℓ-robust. Indeed, we have $|V_{e_i}| = |A_i \cap B_i| = \Delta_G(\varepsilon)$ and $\mathrm{Dom}(\varepsilon) \subseteq A_i \cap B_i$. So the set U and the initial segments $R_j \cap G[A_i]$ together with the trivial paths in $\mathrm{Dom}(\varepsilon)$ are as required in the definition of left-ℓ-robust, as witnessed by the paths $P_\ell^i, \dots, P_\ell^i$ and $P_1^v, \dots, P_\ell^v$ (and because $Z \subseteq (A_i \setminus B_i) \cup \mathrm{Dom}(\varepsilon)$). Hence, (A_i, B_i) is left-ℓ-robust. To show that (A_i, B_i) is also right-ℓ-robust for $i \in I$, we apply Lemma 3.8.4 in $G[B_i]$ to the rays $R_j \cap G[B_i]$ and take the P_x as suitable finite initial segments of those rays and the trivial paths on $\mathrm{Dom}(\varepsilon)$. Altogether, we obtain that all (A_i, B_i) with $i \in I$ and $i \geqslant N_1$ are ℓ-robust and hence m-good. □

Lemma 3.8.6. *Let $(T, \mathcal{V})$ be a tight, rooted tree-decomposition of a graph G into finite parts which displays the infinities of G. Assume that F is some set of edges of T which are not incident with any infinite-degree node of T and such that, for every end ε of G, the set F avoids cofinitely many edges e of the arising ray R_ε in T with $|V_e| = \Delta(\varepsilon)$. Then the tree-decomposition $(T', \mathcal{V}')$ obtained from $(T, \mathcal{V})$ by contracting all edges in F has finite parts and displays the infinities.*

Proof. We first note that $(T', \mathcal{V}')$ displays the infinities: Since $(T, \mathcal{V})$ displays the ends homeomorphically, their combined degrees and their dominating vertices, so does $(T', \mathcal{V}')$, since by the assumptions on F every end ε of G still gives rise to a ray in T', and this ray still contains infinitely many edges e with $|V_e| = \Delta(\varepsilon)$. By assumption, F contains no edges incident with nodes of infinite degree; thus, $(T', \mathcal{V}')$ displays the critical vertex sets and their tight components cofinitely, as $(T, \mathcal{V})$ does so.

To prove that all bags of $(T', \mathcal{V}')$ are finite, it suffices to show by Theorem 2.5.2 that its torsos are tough and rayless. We first show that the torsos are tough. By construction of $(T', \mathcal{V}')$, we did not contract edges of T which are incident with some node $t \in T$ whose corresponding bag V_t is critical. In particular, for every $X \in \mathrm{crit}(G)$, the unique node $t_X \in T$ with $V_t = X$ is also in T' with $V'_{t_X} = X$, and cofinitely many tight components of $G - X$ are some $G \mathring{\uparrow} e$ for an edge $e = t_X t \in T'$ with $t_X <_{T'} t$, since $(T, \mathcal{V})$ displays all critical vertex sets and their tight components cofinitely. Thus, Lemma 3.5.2 ensures that all torsos of the tight rooted tree-decomposition $(T', \mathcal{V}')$ are tough.

It remains to show that the torsos are rayless. Note that the torso of every node $t \in T'$ with $V'_t \in \mathrm{crit}(G)$ is rayless, as it is finite. Thus, it remains to consider $t \in T'$ with $V'_t \notin \mathrm{crit}(G)$. So suppose that there is such a node t whose torso contains a ray R'. Since $(T', \mathcal{V}')$ is tight, we may apply (the comment after) Proposition 3.6.1 to R' and the star σ'_t at t to obtain a ray of G that meets the part V'_t infinitely often, which contradicts the fact that $(T', \mathcal{V}')$ displays the ends of G. □

Proof of Lemma 3.8.3. Let F be the set of edges in T that we contracted. By Lemma 3.8.5, the set F satisfies the premise of Lemma 3.8.6, and hence $(T', \mathcal{V}')$ has finite parts and displays the infinities.

Further, $(T', \mathcal{V}')$ is cofinally componental: Let R' be a rooted ray in T'. It suffices to show that $G \mathring{\uparrow} e$ is disconnected for at most finitely many consecutive edges e of T'. Let $e = rs, f = st \in R'$ with $r <_{T'} s <_{T'} t$ be two successive edges such that $G \mathring{\uparrow} e$ and $G \mathring{\uparrow} f$ are disconnected. Since $(T', \mathcal{V}')$ is obtained from $(T, \mathcal{V})$ by edge-contractions, e and f are also edges of T. It follows from (I1) from Theorem 4' of $(T, \mathcal{V})$ together with the construction of $(T', \mathcal{V}')$ from $(T, \mathcal{V})$ that s and t were already nodes of T and $V_t, V_s \in \mathrm{crit}(G)$ as well as $\deg(t), \deg(s) = \infty$ and $V_s \supseteq V_t$. Since $(T, \mathcal{V})$ displays the critical vertex sets of G, this implies that $V_s \neq V_t$, and thus $V_s \supsetneq V_t$. Hence, this can only happen finitely many times consecutively, as critical vertex sets are finite.

It remains to show (i). For this, let $t \in T'$ be a node with finite degree, and denote with T_t the subtree T_t of T whose contraction yields t. Suppose for a contradiction that some separation $(A, B) \in \sigma'_t$ is not left-m-good. We remark that as $(T', \mathcal{V}')$ is obtained from $(T, \mathcal{V})$ by edge-contractions, the separations induced by the edges of T with precisely one endvertex in T_t are the same as the separations induced by the edges incident with t in T', i.e. the ones in σ'_t; so let $e = t's$ be such an edge with $t' \in T_t$ which induces (A, B). The construction of T' yields that either (A, B) is m-good or one of t', s has infinite degree in T. If T_t is a singleton, then $\deg(t') < \infty$ by assumption on t. If T_t contains at least one edge, then all nodes of T_t, in particular t', have finite degree, as the edges in T_t have been contracted. In both cases, $\deg(s) = \infty$ and thus $V_s =: X \in \mathrm{crit}(G)$. Since $(T, \mathcal{V})$ displays the critical vertex sets and their tight components cofinitely and because $\deg(s) = \infty$, cofinitely, and thus infinitely, many of the tight components of $G - X$ are contained in $G[A]$. Since $X = V_s$ and thus $X \supseteq A \cap B$, this shows that (A, B) is left-m-good as witnessed by the trivial paths in $A \cap B$ and a set U consisting of m vertices that lie in pairwise distinct tight components of $G - X$ contained in $G[A]$. □

With the three ingredients at hand, we are ready to prove the main result of this section.

Proof of Theorem 3'. Let G be a graph without half-grid minor; in particular, G has no $K_{\aleph_0}$ minor. So G has finite tree-width, as it has a normal spanning tree by [84]. Let $(T, \mathcal{V})$ be the (rooted) tree-decomposition of G from Theorem 4', and let $(T', \mathcal{V}')$ be the rooted tree-decomposition obtained from $(T, \mathcal{V})$ by applying Lemma 3.8.3 with $m = (w+1)^2(w+2)+w+1$ if $w := \mathrm{tw}(G) \in \mathbb{N}$

for finitely bounded tree-width G or with $m = 0$ for other G. Then Lemma 3.8.3 yields that all the bags of $(T', \mathcal{V}')$ are finite and that, moreover, for every node $t \in T'$ either $V'_t \in \operatorname{crit}(G)$ or each separation in the star σ'_t at t is left-good. If V'_t is critical in G, then $|V'_t| \leqslant \operatorname{tw}(G) + 1$ since critical vertex sets are infinitely connected and thus every tree-decomposition of G, in particular those witnessing $\operatorname{tw}(G)$, contains V'_t in one of its bags. If V'_t is not critical in G, then all separations in σ'_t are left-good, and we can apply Lemma 3.8.1 to find that the torso of $(T', \mathcal{V}')$ at t has again at most the tree-width of G. Altogether, every torso of $(T', \mathcal{V}')$ is finite and has tree-width at most $\operatorname{tw}(G)$.

We may thus apply Thomas's result [120, Theorem 5] (cf. Theorem 3.1.1 for finite graphs) to the torsos of $(T', \mathcal{V}')$ and obtain, for every node $t \in T'$, an (unrooted) lean tree-decomposition $(T^t, \mathcal{V}^t)$ of $\operatorname{torso}(\sigma_t)$ where T^t is a finite tree[15] and whose parts have size at most the tree-width of G; in particular, all bags in $\mathcal{V}^t$ are finite, as the torso at t is finite. Furthermore, we may assume that for $t \in T'$ with $V'_t \in \operatorname{crit}(G)$, the tree-decomposition $(T^t, \mathcal{V}^t)$ is the trivial tree-decomposition consisting of a single node-tree, since the torso of $(T', \mathcal{V}')$ at t is complete.

Now for every edge $e = st \in T'$, the adhesion set V'_e induces a complete subgraph in both $\operatorname{torso}(\sigma'_s)$ and $\operatorname{torso}(\sigma'_t)$. Thus, we may fix $u \in T^s$ and $w \in T^t$ with $V'_e \subseteq V^s_u, V^t_w$ for every edge $e = st \in T'$. We now build a tree $\tilde{T}$ from the disjoint union of the T^t by joining the corresponding u and w for every $e \in E(T')$; we say that these new edges uv of $\tilde{T}$ *belong to* T' and *correspond to* the respective $st \in T'$. Keeping the respective parts from the $(T^t, \mathcal{V}^t)$, we obtain a tree-decomposition $(\tilde{T}, \tilde{\mathcal{V}})$ of G which has finite parts of size at most the tree-width of G. We remark that the separation induced by an edge $uv \in \tilde{T}$ which belongs to T' is the same as the separation induced by the corresponding edge in T'.

We claim that $(\tilde{T}, \tilde{\mathcal{V}})$ is as desired. For this, let us first note that $(\tilde{T}, \tilde{\mathcal{V}})$ is cofinally componental since $(T', \mathcal{V}')$ is cofinally componental by Lemma 3.8.3 and because the T^t are finite. Moreover, since $(T', \mathcal{V}')$ displays the infinities by Lemma 3.8.3, it follows that $(\tilde{T}, \tilde{\mathcal{V}})$ also does so, as the T^t are finite and consist of a single node, if $V'_t \in \operatorname{crit}(G)$ and $\deg_{T'}(t) = \infty$.

It remains to show that $(\tilde{T}, \tilde{\mathcal{V}})$ is lean. For this, fix nodes $t_1, t_2 \in \tilde{T}$ and sets $Z_1 \subseteq \tilde{V}_{t_1}$ and $Z_2 \subseteq \tilde{V}_{t_2}$ with $|Z_1| = |Z_2| =: k$. We prove the claim by induction on the number of edges on $t_1 \tilde{T} t_2$ that belong to T'.

If there is no such edge, then there exists a node $t \in T'$ with $t_1, t_2 \in T^t$. Since $(T^t, \mathcal{V}^t)$ is a lean tree-decomposition of $\operatorname{torso}(\sigma'_t)$, either there exists an edge $e \in t_1 T^t t_2$ with $|V^t_e| < k$ or there are k disjoint Z_1–Z_2 paths in $\operatorname{torso}(\sigma'_t)$. In the first case, the construction of $(\tilde{T}, \tilde{V})$ yields $t_1 \tilde{T} t_2 = t_1 T^t t_2$ and $\tilde{V}_e = V^t_e$, so e is as desired. In the second case, we distinguish between the case whether V'_t is critical or not. If V'_t is critical, then $t_1 = t_2$ and $\tilde{V}_{t_1} = V'_t$ as T^t has a single node by construction. We then use the infinitely many tight components of $G - V'_t$ to find the desired disjoint Z_1–Z_2 paths.

[15] We remark that this is not explicitly stated in [120, Theorem 5] but follows directly from its proof.

And if V'_t is not critical, then all separations in σ'_t are left-good by (i) from Lemma 3.8.3 and because $(T', \mathcal{V}')$ displays the critical vertex sets of G. Hence, they are left-well-linked by Lemma 3.8.2. This allows us to apply Lemma 3.6.3 (ii) to lift the k pairwise disjoint Z_1–Z_2 paths in torso(σ'_t) to G.

Now suppose that there is an edge $f = s_1 s_2$ on $t_1 \tilde{T} t_2$ which belongs to T'. We may assume by renaming that s_1 appears before s_2 on $t_1 \tilde{T} t_2$. If $|\tilde{V}_f| < k$, then f is the desired edge e; so suppose otherwise. For $i \in \{1, 2\}$, the induction hypothesis then yields either an edge $e \in t_i \tilde{T} s_i$ with $|\tilde{V}_e| < k$ or a family $\mathcal{P}_i$ of k pairwise disjoint Z_i–$\tilde{V}_f$ paths in G. The first case already yields the desired edge e; so we may assume that the second case holds for both $i \in \{1, 2\}$.

If $\tilde{V}_{s_1} \in \mathrm{crit}(G)$, Lemma 3.2.4 yields the desired family of k disjoint Z_1–Z_2 paths, as $\tilde{V}_f \subseteq \tilde{V}_{s_1}$. The same argument applies if $\tilde{V}_{s_2} \in \mathrm{crit}(G)$.

Suppose now that $\tilde{V}_{s_1}, \tilde{V}_{s_2} \notin \mathrm{crit}(G)$. Since there are no k disjoint Z_1–Z_2 paths in G, Menger's theorem (see for example [41, Proposition 8.4.1]) yields a separation (C, D) of G of order less than k with $Z_1 \subseteq C$ and $Z_2 \subseteq D$. We now use (C, D) to find a Z_1–$\tilde{V}_f$ or a Z_2–$\tilde{V}_f$-separator of size less than k, which gives the desired contradiction.

Let (A, B) be the separation of G induced by $\vec{f} = (s_1, s_2)$, and let $X := (A \cap B) \cap (C \setminus D)$ and $Y := (A \cap B) \cap (D \setminus C)$. By symmetry on the assumptions up to this point, we may assume $|Y| \leqslant |X|$. Since $\tilde{V}_{s_1} \notin \mathrm{crit}(G)$, (A, B) is left-good by the construction of $(\tilde{T}, \tilde{\mathcal{V}})$ and (i) from Lemma 3.8.3. Hence, (A, B) is left-well-linked by Lemma 3.8.2. Thus, there exist $|Y|$ pairwise disjoint X–Y paths through $A \setminus B$ in G. All these paths meet $(C \cap D) \cap (A \setminus B)$, so $|(C \cap D) \cap (A \setminus B)| \geqslant |Y|$. But then

$$\begin{aligned} |(A \cap C) \cap (B \cup D)| &= |(A \cap B) \cap (C \setminus D)| + |(A \cap B) \cap (C \cap D)| + |(C \cap D) \cap (A \setminus B)| \\ &\geqslant |X| + |(A \cap B) \cap (C \cap D)| + |Y| = |A \cap B|, \end{aligned}$$

which in turn yields by double counting that

$$|(A \cup C) \cap (B \cap D)| = |A \cap B| + |C \cap D| - |(A \cap C) \cap (B \cup D)| \leqslant |C \cap D| < k.$$

But $(A \cup C) \cap (B \cap D)$ is a Z_2–$\tilde{V}_f$-separator (or a Z_1–$\tilde{V}_f$ separator in the symmetric case), a contradiction. □

3.9 Tree-decompositions distinguishing infinite tangles

By [40, Theorem 3], for every infinite tangle τ of a graph G there is either an end ε of G such that a finite-order separation (A, B) lies in τ if and only if B contains a tail of every, or equivalently some, ray in ε, or there is a critical vertex set X of G such that $(V(C) \cup X, V(G - C)) \in \tau$ for all

$C \in \mathcal{C}(G-X)$. If the first case holds for τ, then we say that τ is *induced* by ε.

Following [65], we say that two infinite tangles τ, τ' of G are *combinatorially distinguishable* if at least one of them is induced by an end or there exists a finite set $X \subseteq V(G)$ such that $(V(C) \cup X, V(G-C)) \in \tau$ for all $C \in \mathcal{C}_X$ and such that $(V(G-C), V(C) \cup X) \in \tau'$ for a component $C \in \mathcal{C}_X$.

With the above characterization of infinite tangles the following observation is immediate:

Observation 3.9.1. *Every tree-decomposition of a graph G that displays its infinities distinguishes all combinatorially distinguishable infinite tangles of G.* □

We emphasise though that such a tree-decomposition need not distinguish those infinite tangles *efficiently*. In fact, we remark that the graph in Construction 4.2.1 (in Chapter 4) has finite tree-width and thus by Theorem 4 a tree-decomposition that displays its infinities, but no tree-decomposition of that graph efficiently distinguishes all its infinite tangles that are induced by ends (see Lemma 4.2.2).

However, Elm and Kurkofka [65, Theorem 1] showed that every graph G has a nested set of separations that efficiently distinguishes all the combinatorially distinguishable infinite tangles of G. They also showed that this is best possible in the following sense: every graph G that has at least two combinatorially indistinguishable infinite tangles has no nested set of separations that efficiently distinguishes all infinite tangles of G [65, Corollary 3.4]. In the special case where G has no half-grid minor we obtain the following strengthening of their result:

Corollary 5'. *Every graph G without half-grid minor has a tree-decomposition that efficiently distinguishes all the combinatorially distinguishable infinite tangles of G.*

Proof. We show that the tree-decomposition $(T, \mathcal{V})$ from Theorem 3' is as desired. For this, let any pair τ_1, τ_2 of combinatorially distinguishable infinite tangle be given. Then by Observation 3.9.1, there is an edge $e \in E(T)$ such that the separation $\{A_e, B_e\}$ induced by e distinguishes τ_1 and τ_2 (though not necessarily efficiently). For $i = 1, 2$, if there is a critical vertex set X_i of G such that $(V(C) \cup X_i, V(G-C)) \in \tau_i$ for all $C \in \mathcal{C}(G-X_i)$, then there is unique infinite-degree node $t_i \in T$ with $V_{t_i} = X_i$, and we then let f_i be the unique edge of the $t_i T e$ path incident with t_i. Otherwise, τ_i is induced by an end ε_i of G, and we then let f_i be any edge on the unique ray R of T starting in e and arising from ε_i such that $|V_f| \geqslant |V_{f_i}|$ for all edges $f >_T f_i$ on R.

Now using the fact that $(T, \mathcal{V})$ is lean, we find an edge e' and a family $\{P_x \mid x \in V_{e'}\}$ of pairwise disjoint V_{f_1}–V_{f_2} paths in G. Since for $i = 1, 2$ either V_{f_i} is linked to ε_i or $V_{f_i} = X_i$, these paths P_x can be extended to paths or rays witnessing that τ_1 and τ_2 cannot be distinguished by a separation of order less than $V_{e'}$, which shows that $(T, \mathcal{V})$ distinguishes τ_1 and τ_2 efficiently. □

Proof of Corollary 5. Apply Corollary 5'. □

4 Counterexamples regarding linked and lean tree-decompositions of infinite graphs

Kříž and Thomas showed that every (finite or infinite) graph of tree-width $k \in \mathbb{N}$ admits a lean tree-decomposition of width k. We discuss a number of counterexamples demonstrating the limits of possible generalisations of their result to arbitrary infinite tree-width.

In particular, we construct a locally finite, planar, connected graph that has no lean tree-decomposition.

This chapter is based on [11] and joint work with Raphael W. Jacobs, Paul Knappe and Max Pitz.

4.1 Introduction

All graphs in this chapter may be infinite, unless otherwise stated.

4.1.1 Lean tree-decompositions

A cornerstone in both Robertson and Seymour's work [113] on well-quasi-ordering finite graphs, and in Thomas's result [119] that the class of infinite graphs of tree-width $< k$ is well-quasi-ordered under the minor relation for all $k \in \mathbb{N}$, is Kříž and Thomas's result on lean tree-decompositions. Recall that a tree-decomposition $(T, (V_t)_{t \in T})$ is *lean* if for every two (not necessarily distinct) nodes $s, t \in T$ and sets of vertices $Z_s \subseteq V_s$ and $Z_t \subseteq V_t$ with $|Z_s| = |Z_t| =: \ell \in \mathbb{N}$, either G contains ℓ pairwise disjoint Z_s–Z_t paths or there exists an edge $e = xy \in sTt$ whose corresponding adhesion set $V_e := V_x \cap V_y$ has size less than ℓ.

Theorem 4.1.1 (Thomas 1990 [120], Kříž and Thomas 1991 [95]). *For every $k \in \mathbb{N}$, every (finite or infinite) graph of tree-width $< k$ has a lean tree-decomposition of width $< k$.*

Is it possible to generalise Theorem 4.1.1 from finite k to arbitrary infinite cardinalities? In what follows let κ be an infinite cardinal. A graph G has *tree-width* $< \kappa$ if it admits a tree-decomposition of width $< \kappa$, i.e. one into parts of size $< \kappa$. A graph G of tree-width $< \aleph_0$, i.e. with a tree-decomposition into finite parts, is said to have *finite tree-width*. The following questions arise naturally:

(i) Does every graph of tree-width $< \kappa$ admit a lean tree-decomposition of width $< \kappa$? In particular, does every graph of finite tree-width admit a lean tree-decomposition into finite parts?

(ii) If not, does every infinite graph at least admit a lean tree-decomposition?

Note that even without the width restriction, Question (ii) remains non-trivial as the leanness-property has to be satisfied within each bag, which means that we cannot take the trivial tree-decomposition into a single part, unless the graph is infinitely connected. Still, our main example shows that the answers to these questions are in the negative:

Example 8. *There is a planar, locally finite, connected graph that has no lean tree-decomposition.*

Every locally finite, connected graph is countable, and thus has tree-width $< \aleph_0$: given an arbitrary enumeration $\{v_i : i \in \mathbb{N}\}$ of a countable graph G, assigning to each vertex r_i of a ray $R = r_0 r_1 \dots$ the bag $V_{r_i} := \{v_0, \dots, v_i\}$ yields a ray- and thus also tree-decomposition $(R, \mathcal{V})$ of G into finite parts. Hence, the graph from Example 8 witnesses that the answers to both questions (i) and (ii) are in the negative.

On the positive side, recall that we provided in Chapter 3 a sufficient criterion that guarantees the existence of a lean tree-decomposition into finite parts:

Theorem 3. *Every graph without half-grid minor has a lean tree-decomposition into finite parts.*

Note that excluding the half-grid as a minor is sufficient but not necessary for the existence of lean tree-decompositions into finite parts: The countably infinite clique $K_{\aleph_0}$ contains the half-grid even as a subgraph but also the above described ray-decomposition of any given countable graph G into finite parts is lean for $G = K_{\aleph_0}$.

As the graph from Example 8 is planar, it has no K_5 minor, and thus no $K_{\aleph_0}$ minor. Hence, in terms of excluded minors, the gap between our positive result Theorem 3 and our negative result Example 8 is quite narrow. Nevertheless, it remains open to exactly characterise the graphs which admit a lean tree-decomposition (into finite parts, or more generally, of width $< \kappa$).

4.1.2 Linked tree-decompositions

Since the answers to questions (i) and (ii) are in the negative, it is natural to ask what happens if we weaken the condition that the tree-decomposition be lean. One possible such weakening is suggested by the 'linkedness'-property, which was extensively studied in [16, 67]: We say a tree-decomposition $(T, \mathcal{V})$ of a graph G is

- *strongly linked* if for every two nodes $s \neq t$ of T there are $\min\{|V_e| : e \in E(sTt)\}$ pairwise disjoint V_s–V_t paths in G.

One may further weaken 'strongly linked' by requiring the existence of disjoint V_s–V_t paths only between nodes $s, t \in T$ that are 'comparable'. For this, recall that given a tree T rooted at a node r, its *tree-order* is given by $s \leqslant t$ for nodes s, t of T if s lies on the (unique) path rTt from r to t. Thomas [119] defined a rooted tree-decomposition $(T, \mathcal{V})$ of a graph G to be

- *linked* if for every two comparable nodes $s < t$ of the rooted tree T there are $\min\{|V_e| : e \in E(sTt)\}$ pairwise disjoint V_s–V_t paths in G.

These weakenings of the 'leanness'-property are motivated by the fact that for the aforementioned applications of Theorem 4.1.1 by Robertson and Seymour [113] and by Thomas [119] it is only important that the rooted tree-decomposition is linked.

So what happens if we replace the condition *lean* in (i) and (ii) by *strongly linked* or *linked*? As the trivial tree-decomposition into a single part is always strongly linked, and thus every graph has a strongly linked tree-decomposition, only the weakened version of (i), but not of (ii), is interesting:

(iii) Does every graph of tree-width $< \kappa$ admit a strongly linked tree-decomposition of width $< \kappa$?

However, question (iii) is trivially true for all infinite cardinalities κ: By definition, every graph G of tree-width $< \kappa$ admits a tree-decomposition $(T, \mathcal{V})$ of width $< \kappa$. Choose an arbitrary root r of T. By assigning to each of the nodes t of T the bag $V'_t := \bigcup_{s \in rTt} V_s$ we obtain a (rooted) tree-decomposition $(T, \mathcal{V}')$ of width $< \kappa$ that is strongly linked; albeit for the trivial reason that $s \neq t \in T$ implies $V'_s \cap V'_t = V'_u$ where u is the $\leqslant$-minimal node in sTt.

But this tree-decomposition is not useful in practice, and so one would like to have some additional properties making the tree-decomposition less redundant. Especially linked rooted tree-decompositions into finite parts which are additionally 'tight' and 'componental' turned out to be a powerful tool (see Sections 3.1.2–3.1.4 and Chapter 5 for details; in the paragraph after Theorem 1 below we also provide a brief summary). Given a rooted tree-decomposition $(T, \mathcal{V})$ of a graph G and an edge e of T, we call the subgraph $G \mathring{\uparrow} e$ of G induced on $\bigcup_{t \in T_e} V_t \setminus V_e$ the *part strictly above* e, where T_e is the unique component of $T - e$ that does not contain the root of T. Then $(T, \mathcal{V})$ is

- *componental* if all the parts $G \mathring{\uparrow} e$ strictly above edges $e \in E(T)$ are connected, and
- *tight* if for every edge $e \in E(T)$ there is some component C of $G \mathring{\uparrow} e$ such that $N_G(C) = V_e$.

We remark that the above strongly linked (rooted) tree-decomposition $(T, \mathcal{V}')$ is componental if $(T, \mathcal{V})$ was componental; but even if $(T, \mathcal{V})$ was tight, $(T, \mathcal{V}')$ may no longer be tight.

The property *tight* ensures that the adhesion sets contain no 'unnecessary' vertices. Any given componental rooted tree-decomposition can easily be transformed into a tight and componental rooted tree-decomposition by deleting for every edge e of T the non-neighbours of $G \mathring{\uparrow} e$ in V_e from every V_t with $t \in T_e$. While this construction obviously does not increase the width, it does not necessarily maintain the property *(strongly) linked*. So it is natural to strengthen (iii) to ask whether there exists a tree-decomposition which has all three properties:

(iv) Does every graph of tree-width $< \kappa$ admit a tight and componental rooted tree-decomposition of width $< \kappa$ that is strongly linked?

(v) If not, does every graph of tree-width $< \kappa$ admit a tight and componental rooted tree-decomposition of width $< \kappa$ that is at least linked?

The answer to Question (iv) is in the negative already for $\kappa = \aleph_0$, as witnessed by the same graph that we constructed for Example 8:

Example 9. *There is a planar, locally finite, connected graph which admits no tight, componental rooted tree-decomposition into finite parts that is strongly linked.*

Question (v), however, has an affirmative answer for $\kappa = \aleph_0$ as we have shown in Chapter 3:

Theorem 1. *Every graph of finite tree-width admits a rooted tree-decomposition into finite parts that is linked, tight and componental.*

We remark that (v) remains open for uncountable cardinalities $\kappa > \aleph_0$. If the answer is positive, one might, as a next step, also strengthen the notion of *linked* from just considering sizes to a structural notion that encapsulates the typical desired behaviour of infinite path families between two sets, as it is given by Menger's theorem for infinite graphs [2, Theorem 1.6] proven by Aharoni and Berger.

It turns out that linked rooted tree-decompositions into finite parts which are additionally tight and componental, as given by Theorem 1, are particularly useful (cf. Sections 3.1.2–3.1.4 and Chapter 5): In Section 3.3, we have shown that rooted tree-decompositions into finite parts which are linked, tight and componental display the end structure of the underlying graph. This not only resolves a question of Halin [83, §6] but also allowed us to deduce from Theorem 1, by means of short and unified proofs, the characterisations due to Robertson, Seymour and Thomas of graphs without half-grid minor [115, Theorem 2.6], and of graphs without binary tree subdivision [116, (1.5)]. Also the proof of Theorem 3 in Section 3.8 heavily relied on post-processing the tree-decomposition from Theorem 1. Beside these, there are more applications of rooted tree-decomposition into finite parts which are linked, tight and componental in Section 3.1.4 and Chapter 5.

In fact, we have shown in Section 3.4.2 (Theorem 1') a more detailed version of Theorem 1 which, among others, yields that the adhesion sets of the tree-decomposition intersect 'not more than necessary'. We also give an example which proves that this property is best possible even for locally finite graphs (see Section 4.4 for details).

In the light of Question (v) being true for $\kappa = \aleph_0$, one may ask whether a similar modification could rescue (i) for $\kappa = \aleph_0$: What happens if we relax the condition *lean* in (i) to a corresponding 'rooted' version?

(vi) Does every graph of finite tree-width admit a rooted tree-decomposition into finite parts that satisfies the property of being lean for all *comparable* nodes $s \leqslant t$ of T?

However, the answer to this question is again in the negative, as there is a graph of finite tree-width such that all its tree-decompositions into finite parts violate the 'leanness'-property within a single bag:

Example 10. *There is a countable graph G such that every tree-decomposition of G into finite parts has a bag V_t which violates the property of being lean for $s = t$.*

We remark that we do not know whether every tree-decomposition which satisfies the 'leanness'-property for every two comparable nodes must already be lean.

4.1.3 How this chapter is organised

In Section 4.2 we prove Examples 8 and 9, and in Section 4.3 we prove Example 10. Finally, in Section 4.4, we discuss whether it is possible to strengthen Theorem 1 so that the adhesion sets of the tree-decomposition are 'upwards' disjoint.

We refer the reader to Section 3.2 for all definitions that were not presented in Chapter 2.

4.2 A graph with no lean tree-decomposition

In this section we explain Examples 8 and 9, which we restate here for convenience:

Example 8. *There is a planar, locally finite, connected graph that has no lean tree-decomposition.*

Example 9. *There is a planar, locally finite, connected graph which admits no tight, componental rooted tree-decomposition into finite parts that is strongly linked.*

For our proofs of Examples 8 and 9 we construct a graph G in Construction 4.2.1 below, and then show that G is already as desired for both Examples 8 and 9. The graph in this construction is inspired by [31, Example 7.4].[1]

Construction 4.2.1. Let G be the graph depicted in Figure 4.1. Formally, let G' be the graph on the vertex set $V(G') := \{(i/2^j, j) \mid j \in \mathbb{N}, 0 \leqslant i \leqslant 2^{j+1}\}$ and with edges between $(i/2^j, j)$ and $((i+1)/2^j, j)$, between $(i/2^j, j)$ and $(i/2^j, j+1)$, and also between $(i/2^j, j)$ and $((2i-1)/2^{j+1}, j+1)$

[1]The example is only presented in the arXiv version of [31].

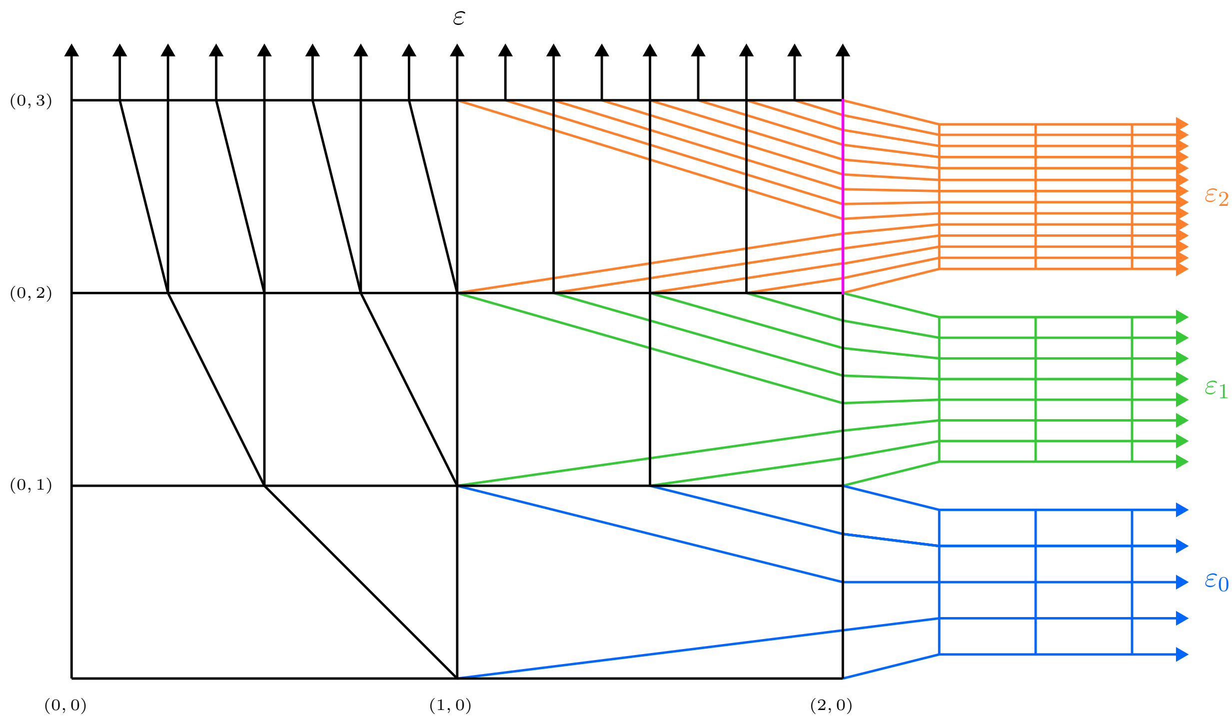

FIGURE 4.1: Depicted is the graph G from Construction 4.2.1. The subgraph induced by the orange edges is G_2 together with the extensions of its horizontal rays.

for $i \leqslant 2^j$ (this is the black subgraph in Figure 4.1). Note that G' has a unique end, which we denote by ε.

For every $n \in \mathbb{N}_{\geqslant 1}$, let $G_n := \mathbb{N}_{\geqslant 1} \boxtimes P_{2^{n+1}+2^n+1}$ be a grid with $2^{n+1} + 2^n + 2$ rows and infinitely many columns. Then G_n is one-ended and its end ε_n has degree $2^{n+1} + 2^n + 2$. Now the graph G is obtained from $G' \sqcup \bigsqcup_{n \in \mathbb{N}} G_n$ by deleting for every $n \in \mathbb{N}$ the edge $\{(2, n), (2, n + 1)\}$ of G', identifying the vertices $(2, n), (2, n+1)$ of G' with the respective vertices $(1, 1)$ and $(1, 2^{n+1}+2^n+2)$ of G_n, and by extending the horizontal rays in G_n as indicated in Figure 4.1. In particular, the extended horizontal rays of each G_n are still disjoint and their initial vertices are precisely $\{(i/2^j, j) \mid j \in \{n, n+1\}, 2^j \leqslant i \leqslant 2^{j+1}\}$. To get a graph that is not only locally finite but also planar, we subdivide the edges between $(i/2^n, n)$ and $(i/2^n, n + 1)$ with $i > 2^n$ in G' to obtain the extended rays. For later use, we refer to the vertex set $\{(i, j) \mid 1 \leqslant j \leqslant 2^{n+1} + 2^n + 2\}$ in G_n as the *i-th column of* G_n (indicated in pink in Figure 4.1 is the first column of G_2) and also set

$$S_n := \{(i/2^{n+1}, n + 1) \mid 2^{n+1} < i \leqslant 2^{n+2}\} \cup \{(1, j) \mid 0 \leqslant j \leqslant n + 1\}$$

for all $n \in \mathbb{N}$ (indicated in purple in Figure 4.2.) This completes the construction.

In the remainder of this section we prove that the graph G from Construction 4.2.1 is as desired for Examples 8 and 9. For this, we first show two auxiliary lemmas. The first says that G does not admit a tree-decomposition which 'efficiently distinguishes' all ends of G. Recall that in a tree-decomposition $(T, \mathcal{V})$ of G every edge $e = t_0t_1$ of T induces a separation as follows: For $i = 0, 1$ write T_i for the component of $T - e$ that contains t_i. Then $\{\bigcup_{s \in T_0} V_s, \bigcup_{s \in T_1} V_s\}$ is a separation of G [41, Lemma 12.3.1]. A separation $\{A, B\}$ of G *efficiently distinguishes* two ends $\varepsilon, \varepsilon'$ of G if ε and ε' live in components of $G - (A \cap B)$ on different sides of $\{A, B\}$, and there is no separation $\{C, D\}$ of G of smaller order $|C \cap D|$ with this property. A tree-decomposition $(T, \mathcal{V})$ *distinguishes* two ends $\varepsilon, \varepsilon'$ of G *efficiently* if some edge e of T induces a separation which efficiently distinguishes ε and ε'.

Lemma 4.2.2. *Let $(T, \mathcal{V})$ be a tree-decomposition of the graph from Construction 4.2.1. Then there exist $n \in \mathbb{N}$ such that $(T, \mathcal{V})$ does not efficiently distinguish ε_n and ε.*

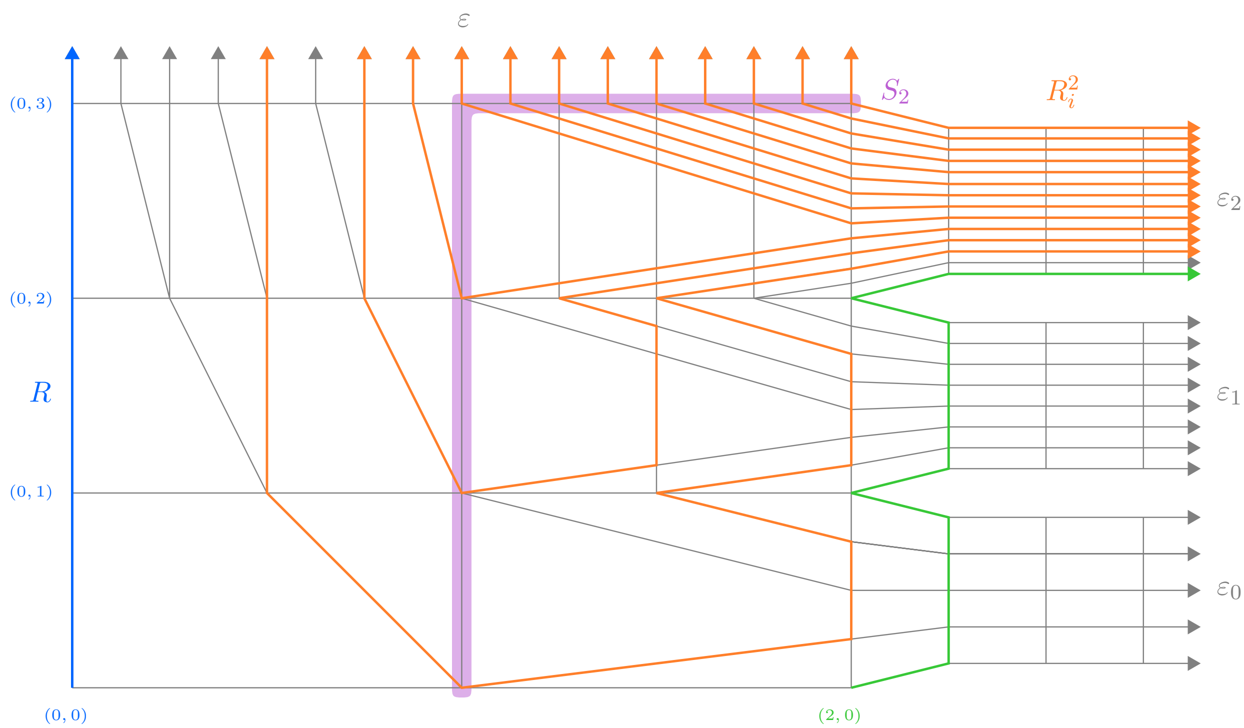

FIGURE 4.2: Depicted is the situation in the proof of Lemma 4.2.2 for $n = 2$.

Proof. Suppose towards a contradiction that G admits a tree-decomposition $(T, \mathcal{V})$ such that, for every end ε_n of G, there exists an edge f_n such that the separation induced by f_n distinguishes ε_n and ε efficiently. Then V_{f_n} has size at most $2^{n+1} + n + 2$ as witnessed by S_n (indicated in purple in Figure 4.2). By the definition of G', there are in fact $|S_n|$ disjoint ε–ε_n double rays R_i^n in G

(indicated in orange in Figure 4.2), so every sizewise-minimal ε–ε_n separator, and in particular V_{f_n}, has to meet each R_i^n precisely once. In particular, $|V_{f_n}| = |S_n|$ and hence $f_n \neq f_m$ for all $n \neq m \in \mathbb{N}$. Moreover, V_{f_n} avoids the ε-ray $R = (0,0)(0,1)\dots$ (indicated in blue in Figure 4.2) and the vertex $(2,0)$, as they are both disjoint from all the ε–ε_n double rays R_i^n. Thus, $G - V_{f_n}$ has a component C_n that contains R, and a component C_n' that contains $(2,0)$. In particular, ε lives in C_n. Since there is an ε_n-ray (indicated in green in Figure 4.2) that starts in $(2,0)$ and is disjoint from the R_i^n, the end ε_n lives in C_n'. As V_{f_n} separates ε and ε_n, we have $C_n \neq C_n'$. Thus, $(0,0)$ and $(2,0)$ lie on different sides of the separation induced by f_n.

Let V_t and V_s be some bags of $(T,\mathcal{V})$ which contain $(0,0)$ and $(2,0)$, respectively. Since the separations induced by every f_n separate $(0,0)$ and $(2,0)$, all the infinitely many distinct edges f_n lie on the finite path tTs, which is a contradiction. □

The next lemma essentially says that every tree-decomposition of G that displays all ends of G and their combined degrees cannot be strongly linked.

Lemma 4.2.3. *Let $(T,\mathcal{V})$ be a rooted tree-decomposition of the graph G from Construction 4.2.1. Suppose that every end ω of G gives rise to a rooted ray R_ω in T with $\liminf_{e\in R_\omega} V_e = \emptyset$ and $\liminf_{e\in R_\omega} |V_e| = \deg(\omega)$. Then $(T,\mathcal{V})$ is not strongly linked.*

Proof. Suppose towards a contradiction that $(T,\mathcal{V})$ is strongly linked. Let $n \in \mathbb{N}$ be arbitrary. We write R_n for the rooted ray R_{ε_n} in T, which arose from ε_n. Since $(T,\mathcal{V})$ is a tree-decomposition, for every ray R of T, we have $\bigcap_{e\in R} V(G\mathring{\uparrow}e) = \emptyset$, and thus for every finite vertex set W of G, all but finitely many edges e of R have the property that $G\mathring{\uparrow}e$ avoids W. In particular, for all but finitely many edges e of R_n, the part $G\mathring{\uparrow}e$ avoids the first and second column of G_n. As $\liminf_{e\in R_n} |V_e| = \deg(\varepsilon_n)$, we may choose such an edge $e_n \in R_n$ so that $|V_{e_n}| = \deg(\varepsilon_n)$ and every later edge e on R_n satisfies $|V_e| \geqslant |V_{e_n}|$. Since $G\mathring{\uparrow}e_n$ avoids the first column of G_n, the component D_n of $G\mathring{\uparrow}e_n$ in which ε_n lives is contained in G_n.

We claim that $G\mathring{\uparrow}e_n$ contains all but finitely many vertices of G_n, and $N_G(D_n) = V_{e_n}$. Indeed, the rows of G_n are pairwise disjoint ε_n-rays. Since ε_n lives in D_n, the component D_n contains a tail of each row of G_n. As the rows of G_n cover the entire G_n, the part $G\mathring{\uparrow}e_n \supseteq D_n$ contains all but finitely many vertices of G_n. Since $D_n \subseteq G\mathring{\uparrow}e_n$ avoids the first column of G_n, its neighbourhood $N_G(D_n) \subseteq V_{e_n}$ contains at least one vertex of each row. Since there are $\deg(\varepsilon_n)$ rows, we have $|N_G(D_n)| \geqslant \deg(\varepsilon_n) = |V_{e_n}|$, and thus $N_G(D_n) = V_{e_n}$.

As $N_G(D_n) = V_{e_n}$ and D_n is connected, $G\uparrow e_n$ is connected. It follows that $G\uparrow e_n$ is a subgraph of G_n and avoids the first column of G_n, since $G\mathring{\uparrow}e_n$ meets G_n and avoids the first and second column of G_n.

Let H_n be the component of $G - S_n$ that contains G_0 (where S_n is the set indicated in purple in Figure 4.2). We claim that there is an edge e_n' of the rooted ray R_ε in T, which arose from ε, such that $|V_{e'_n}| \geqslant \deg(\varepsilon_n)$ and such that $G{\uparrow}e$ avoids H_n. For this, we note that the rays R_m are distinct from R_ε, since $\liminf_{e\in R_m}|V_e| = \deg(\varepsilon_m)$ is finite but $\liminf_{e\in R_\varepsilon}|V_e| = \deg(\varepsilon)$ is infinite. Hence, we may choose a tail $R^n_\varepsilon \subseteq R_\varepsilon$ of R_ε that avoids all rays R_m for $m \leqslant n$. Since all R_m and also R_ε are rooted at the same vertex of T, every edge e of R^n_ε satisfies $(G{\uparrow}e) \cap (G\mathring{\uparrow}e_m) = \emptyset$. As $G\mathring{\uparrow}e_m$ contains all but finitely many vertices of G_m for $m \leqslant n$, the set $W_n := V(H_n) \setminus (\bigcup_{m\leqslant n} V(G\mathring{\uparrow}e_n))$ is finite. Since $(T,\mathcal{V})$ is a tree-decomposition, for all but finitely many edges e of R^n_ε, the part $G\mathring{\uparrow}e$ avoids W_n and hence H_n. Note that, for each edge e of R^n_ε, its adhesion set V_e avoids $G\mathring{\uparrow}e_m$ for every $m \leqslant n$. Since $\liminf_{e\in R^n_\varepsilon} V_e = \emptyset$ and W_n is finite, for all but finitely many edges e of R^n_ε, the adhesion set V_e avoids the finite set W_n. Thus, for all but finitely many edges e of R_ε, the part $G{\uparrow}e$ avoids H_n. Finally, since $\liminf_{e\in R_\varepsilon}|V_e| = \deg(\varepsilon)$ is infinite, we may choose an edge e'_n of R_ε so that $G{\uparrow}e'_n$ avoids H_n and $|V_{e'_n}| \geqslant \deg(\varepsilon_n)$.

Recall that the adhesion set V_{e_n} is contained in G_n but avoids the first column of G_n, and thus $V_{e_n} \subseteq V(H_n)$. Since also $V_{e'_n} \cap V(H_n) = \emptyset$, there are at most $2^{n+1} + n + 2$ pairwise disjoint V_{e_n}–$V_{e'_n}$ paths in G, as witnessed by S_n. As $(T,\mathcal{V})$ is strongly linked by assumption, there is an edge f_n on the unique e_n–e'_n path in T such that $|V_{f_n}| \leqslant |S_n|$. Moreover, V_{f_n} separates V_{e_n} and $V_{e'_n}$, and thus also ε_n and ε. By the definition of G', there are in fact $|S_n|$ disjoint ε–ε_n double rays R^n_i in G (indicated in orange in Figure 4.2), and hence the separation induced by f_n efficiently distinguishes ε and ε_n. Since $n \in \mathbb{N}$ was chosen arbitrarily, Lemma 4.2.2 yields the desired contradiction. □

Proof of Example 9. The graph G from Construction 4.2.1 is planar, locally finite and connected. By Lemma 3.3.4, every linked, tight, componental rooted tree-decomposition of G into finite parts displays all the ends of G, their dominating vertices and their (combined) degrees. Thus, Lemma 4.2.3 ensures that those tree-decompositionsare not strongly linked. □

We now turn to our proof of Example 8, i.e. that the graph G from Construction 4.2.1 does not admit a lean tree-decomposition. The proof consists of two steps. First, we show in Lemma 4.2.6 that G does not admit a lean tree-decomposition into finite parts. Then, we show in Lemma 4.2.7 that G neither admits a lean tree-decomposition that has an infinite part.

By definition, every lean tree-decomposition is in particular strongly linked. For the first step, it remains to show that every lean tree-decomposition of G into finite parts satisfies the premise of Lemma 4.2.3: it displays all ends of G and their combined degrees, up to the fact that maybe some rays of the decomposition tree do not arise from an end of G. In fact, the following Lemma 4.2.5 together with Lemma 3.3.3 implies that this holds true for *all* graphs H and lean tree-decompositions of them.

To prove Lemma 4.2.5, we first show that even though lean tree-decompositions may not be componental, they are not far away from it.

Lemma 4.2.4. *Let $(T, \mathcal{V})$ be a rooted tree-decomposition of finite adhesion of a graph H. Suppose there is an edge $e = st \in E(T)$ with $s <_T t$ and a set $Y \subseteq V_e$ such that $(H{\uparrow}e) - Y$ has at least two components C_1, C_2. Suppose further that $V(C_1) \cap V_e \neq \emptyset$ and $V(C_2) \cap (V_t \setminus V_e) \neq \emptyset$. Then $(T, \mathcal{V})$ is not lean.*

Proof. By assumption, we may pick vertices $v_1 \in V(C_1) \cap V_e$ and $v_2 \in V(C_2) \cap (V_t \setminus V_e)$. Set $Z_1 := V_e$ and $Z_2 := (V_e \setminus \{v_1\}) \cup \{v_2\}$. By construction, $|Z_1| = |V_e| = |Z_2| =: k$. Since $(T, \mathcal{V})$ is lean and $Z_1, Z_2 \subseteq V_t$, there is a family $\mathcal{P}$ of k disjoint Z_1–Z_2 paths in H. Note that every path in $\mathcal{P}$ that starts in $Z_1 \cap Z_2 = V_e \setminus \{v_1\}$ is trivial. Hence, there is a path $P \in \mathcal{P}$ that starts in $\{v_1\} = Z_1 \setminus Z_2$ and ends in $\{v_2\} = Z_2 \setminus Z_1$. But $v_1 \in V(C_1)$ is separated from $v_2 \in V(C_2) \setminus V_e \subseteq V(H\mathring{\uparrow}e)$ by $Y \cup (V(C_2) \cap V_e) \subseteq V_e \setminus \{v_1\}$. This contradicts that the paths in $\mathcal{P}$ are pairwise disjoint. ☐

Lemma 4.2.5. *Let $(T, \mathcal{V})$ be a rooted tree-decomposition of a graph H into finite parts which is lean as an unrooted tree-decomposition. Then every ray in T arises from at most one end of H. Moreover, if a ray R in T arises from an end ε of H, then* $\liminf_{e \in R} V_e = \mathrm{Dom}(\varepsilon)$.

Proof. For the first assertion, suppose towards a contradiction that there are two distinct ends $\varepsilon_1, \varepsilon_2$ of H that give rise to the same rooted ray R of T. As $\varepsilon_1, \varepsilon_2$ are distinct, there is a finite set X of vertices of H such that $\varepsilon_1, \varepsilon_2$ live in distinct components C_1, C_2 of $G - X$.

Now pick a vertex $v_1 \in V(C_1)$. Since $(T, \mathcal{V})$ is a tree-decomposition, $\bigcap_{e \in R} V(H\mathring{\uparrow}e) = \emptyset$, and thus all but finitely many edges e of R have the property that $H{\downarrow}e$ contains the finite set $X \cup \{v_1\}$. As ε_2 gives rise to R and lives in C_2, there is an edge $e = st$ on R with $s <_T t$ such that $X \cup \{v_1\} \subseteq H{\downarrow}e$ and $V(C_2) \cap (V_t \setminus V_e)$ is non-empty. Note that $V(C_1) \cap V_e$ is non-empty as C_1 is connected and meets both $H{\downarrow}e$ (in the vertex v_1) and $H\mathring{\uparrow}e$ (as ε_1 gives rise to R and lives in C_1). Set $Y := X \cap V_e$ and, for $i = 1, 2$, let C_i' be a component of $C_i \cap (H{\uparrow}e - Y)$ which contains a vertex from $V(C_1) \cap V_e$ or $V(C_2) \cap (V_t \setminus V_e)$, respectively. Then Y, C_1', C_2' are as in Lemma 4.2.4. It follows that $(T, \mathcal{V})$ is not lean, which is a contradiction.

To show the moreover statement, let ε be an end of H that gives rise to a ray R in T. Now suppose towards a contradiction that there is a vertex $w \in \liminf_{e \in R} V_e$ that does not dominate ε. Then there is a finite set $X \subseteq V(H)$ and distinct components C_1, C_2 of $G - X$ such that $w \in V(C_1)$ and ε lives in C_2.

As above, there is an edge $e = st$ on R with $s <_T t$ such that $H{\downarrow}e$ contains $X \cup \{w\}$ and $V(C_2) \cap (V_t \setminus V_e)$ is non-empty. Since w also lies in $\liminf_{f \in R} V_f$, we have $w \in V_e$, and thus $V(C_1) \cap V_e$ is non-empty. As above we obtain a contradiction by applying Lemma 4.2.4. ☐

Lemma 4.2.6. *The graph G from Construction 4.2.1 has no lean tree-decomposition into finite parts.*

Proof. Suppose towards a contradiction that G admits a lean tree-decomposition $(T, \mathcal{V})$ into finite parts. It follows from Lemma 4.2.5 that every end ω of G gives rise to a ray R_ω in T which arises from no other end of G. Moreover, we have $\liminf_{e \in R_\omega} V_e = \mathrm{Dom}(\omega) = \emptyset$, as locally finite graphs have no dominating vertices. So since $(T, \mathcal{V})$ is lean and hence strongly linked, Lemma 3.3.3 implies that $\liminf_{e \in R_\omega} |V_e| = \Delta(\omega) = \deg(\omega)$ (because G is locally finite). Thus, by Lemma 4.2.3, $(T, \mathcal{V})$ is not strongly linked, which is a contradiction. □

For the proof of Example 8 it remains to show that the graph from Construction 4.2.1 neither admits a lean tree-decomposition with possibly infinite parts.

Lemma 4.2.7. *The graph G from Construction 4.2.1 admits no lean tree-decomposition.*

Proof. Suppose for a contradiction that G has a lean tree-decomposition $(T, \mathcal{V})$. Let T be rooted in an arbitrary node. We first show that every end ε_n gives rise to a ray R_n in T, that is every bag of $(T, \mathcal{V})$ meets every ε_n-ray at most finitely often. For this it suffices to show that every bag meets G_n at most finitely often. Let $n \in \mathbb{N}$ be given, and suppose there is a bag V_t that contains infinitely many vertices of G_n. Then let $Z_1 \subseteq V_t$ be a set of $2^{n+1} + 2^n + 3$ vertices of G_n, and let $i \in \mathbb{N}$ such that Z_1 is contained in the first i columns of G_n. Now let $Z_2 \subseteq V_t$ be a set of $2^{n+1} + 2^n + 3$ vertices of G_n that avoids the first i columns of G_n. Since the i-th column of G_n has size $2^{n+1} + 2^n + 2$ and separates Z_1 and Z_2, this contradicts that $(T, \mathcal{V})$ is lean.

Hence, each ε_n-ray meets every bag of $(T, \mathcal{V})$ at most finitely often, which implies that ε_n gives rise to a ray R_n in T. In particular, then, there exists a node t of R_n such that V_t meets every row of G_n: If not, then every bag V_t of R_n avoids some row of G_n. In fact, since ε_n gives rise to R_n and thus every V_t meets every row that meets V_s for some node $s <_T t$ of R_n, all V_t with $t \in R_n$ avoid the same row of G_n. Since ε_n gives rise to R_n, it follows that this row is contained in $G\mathring{\uparrow}e$ for all edges e of R_n, which contradicts that $(T, \mathcal{V})$ is a tree-decomposition. Hence, there is a node t_n of R_n such that V_{t_n} meets every row of G_n. Let $X_n \subseteq V_{t_n}$ be a set consisting of precisely one vertex of each row of G_n. In particular, X_n has size $2^{n+1} + 2^n + 2$ and is linked to ε_n.

Since G admits no lean tree-decomposition into finite parts by Lemma 4.2.6, $(T, \mathcal{V})$ contains an infinite bag V_s. As shown above, V_s contains at most finitely many vertices of each G_n. Hence, for every $n \in \mathbb{N}$, the bag V_s must contain infinitely many vertices of the subgraph $G'_n := G'[\{(i,j) \in V(G') : j > n\}] \cup \bigcup_{m>n} G_m$. Let $Y_n \subseteq V_s$ consist of $2^{n+1} + 2^n + 2$ vertices of G'_n. Since G is locally finite and connected, it follows from the Star-Comb Lemma (see Lemma 2.5.1) that there is a comb C in G with teeth in V_s. Recall that the comb C is the union of a ray R, its spine, and infinitely many disjoint (possibly trivial) paths with precisely their first vertex in R and their last vertex in V_s. As $V_s \cap V(G_n)$ is finite for all $n \in \mathbb{N}$, the spine R of C is an ε-ray.

Since $X_n \subseteq V(G_n)$ and $Y_n \subseteq V(G'_n)$, there are at most $2^{n+1} + n + 2$ pairwise disjoint X_n–Y_n paths in G, as witnessed by S_n (indicated in purple in Figure 4.2). As $(T, \mathcal{V})$ is lean by assumption, there is an edge f_n on the unique t_n–s path in T such that $|V_{f_n}| \leqslant |S_n|$. Moreover, V_{f_n} separates V_{t_n} and V_s. We claim that V_{f_n} also separates ε_n and ε. Indeed, since C has teeth in V_s, the component D of $G - V_{f_n}$ in which ε lives meets $V_s \setminus V_{f_n}$. Because X_n is linked to ε_n and $|V_{f_n}| \leqslant |S_n| < |X_n|$, the component D_n of $G - V_{f_n}$ in which ε_n lives meets $X_n \setminus V_{f_n} \subseteq V_{t_n} \setminus V_{f_n}$. Since V_{f_n} separates V_{t_n} and V_s, the components D and D_n are distinct, i.e. V_{f_n} separates ε and ε_n. By the definition of G, there are $|S_n|$ disjoint ε–ε_n double rays R_i^n in G, and hence the separation induced by f_n distinguishes ε and ε_n efficiently. Now Lemma 4.2.2 yields the desired contradiction. □

Proof of Example 8. The graph G from Construction 4.2.1 is planar, locally finite and connected. The assertion thus follows from Lemma 4.2.7. □

4.3 A further counterexample regarding the 'leanness'-property

In this section we construct Example 10, which we restate here for convenience:

Example 10. *There is a countable graph G such that every tree-decomposition of G into finite parts has a bag V_t which violates the property of being lean for $s = t$.*

The graph in this example is essentially the same as [26, Example 3.7].

Proof. Let G' be the $\mathbb{N} \times \{0,1,2\}$ grid, that is, $V(G') = \{(i,j) \mid i \in \mathbb{N},\ j \in \{0,1,2\}\}$, and there is an edge in G' between (i,j) and (i',j') whenever $|i-i'| + |j-j'| = 1$. For every $n \in \mathbb{N}_{\geqslant 1}$, let $U_n := \{(i,1) \mid i \leqslant n\} \cup \{(n-1,0),(n,0)\}$ (see Figure 4.3). Now the graph G is obtained from G' by making the sets U_n complete. We claim that G is as desired.

Let $(T, \mathcal{V})$ be a rooted tree-decomposition of G into finite parts. Let $R \subseteq T$ be the rooted ray arising from the unique end ε of G. Since $(T,\mathcal{V})$ is a tree-decomposition, we have $\bigcap_{e \in R} V(G\mathring{\uparrow} e) = \emptyset$, and hence there is an edge e of R such that $(0,0),(0,1),(0,2) \in V(G{\downarrow}e)$. As ε gives rise to R, the ray $(0,0)(1,0)(2,0)\dots$ through the bottom row has a tail in $G\mathring{\uparrow}e$; i.e. there is $n \in \mathbb{N}$ such that $(n',0) \in V(G\mathring{\uparrow}e)$ for all $n' \geqslant n$. Since $G[U_{n'}]$ is complete, it follows that $U_{n'} \subseteq V(G{\uparrow}e)$ for all $n' \geqslant n$.

Next, observe that there is an edge $e' > e$ of R such that $U_n \subseteq V(G{\downarrow}e')$. Now let $f = t_1t_2$ with $t_1 <_T t_2$ be the $\leqslant_T$-minimal edge of R such that there exists $m \in \mathbb{N}_{\geqslant n}$ with $U_m \subseteq V(G{\downarrow}f)$. Note that e' is a candidate for f, and observe that $e < f$. Let m be maximal such that $U_m \subseteq V(G{\downarrow}f)$; in particular, $m \geqslant n$. To see that this maximum exists, note that the ray $(0,0)(1,0)(2,0)\dots$ has a tail in $G\mathring{\uparrow}f$, and hence there is $i \in \mathbb{N}$ such that $(j,0) \in V(G\mathring{\uparrow}f)$ for all $j \geqslant i$. Thus, $m \leqslant i$.

By the choice of f and m, we have $U_{m+1} \subseteq V(G{\uparrow}f)$, as U_{m+1} is complete. Hence, as V_f separates $G\mathring{\uparrow}f$ and $G\mathring{\downarrow}f$, it follows that $U_m \cap U_{m+1} \subseteq V_f \subseteq V_{t_1}$. Moreover, as $(0,2) \in V(G{\downarrow}e) \subseteq V(G{\downarrow}f)$

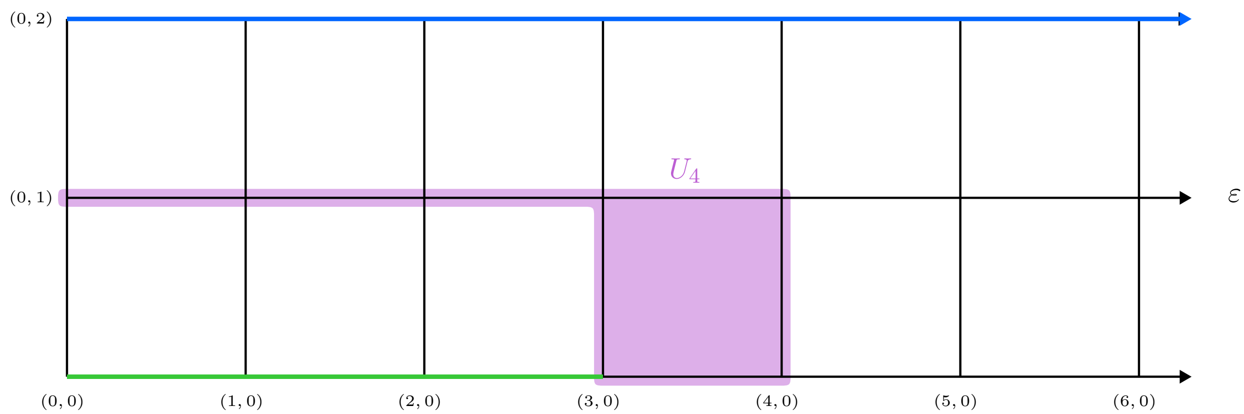

FIGURE 4.3: Depicted is the graph G' and the set U_4 in purple from Example 10. Indicated in blue is the ray $(0,2)(1,2)\dots$ that contains the vertex w and indicated in green is the path $(0,0)(1,0)(2,0)(3,0)$ that contains the vertex u in the case $m=4$.

by the choice of e, and because ε gives rise to R, the ray $(0,2)(1,2)(2,2)\dots$ meets $V_f \subseteq V_{t_1}$ in a vertex w. Similarly, the path $(0,0)(1,0)\dots(m-1,0)$ meets $V_{f'} \subseteq V_{t_1}$ in a vertex u where f' is the unique down-edge at t_1. Indeed, we have $(0,0) \in G{\downarrow}f'$ by the choice of $e < f$ and $(m-1,0) \in G{\uparrow}f'$: if $(m-1,0)$ was contained in $G\mathring{\downarrow}f'$, then $U_m \subseteq V(G{\downarrow}f')$ since U_m is complete in G, which contradicts the $\leqslant_T$-minimal choice of f.

We now let $s := t_1 =: t$ and define $Z_1 := (U_m \cap U_{m+1}) \cup \{w\}$ and $Z_2 := (U_m \cap U_{m+1}) \cup \{u\}$. By construction, we have $|Z_1| = |Z_2| = m+3$ and $Z_1, Z_2 \subseteq V_{t_1}$. Hence, $(T, \mathcal{V})$ violates the property of being lean for $s = t_1 = t$ since $U_m \cap U_{m+1}$ separates w and u and hence witnesses that G contains at most $m+2$ disjoint Z_1–Z_2 paths. □

4.4 Upwards disjointness of adhesion sets

As mentioned in the introduction, we have shown in Section 3.4.2 (Theorem 1') a more detailed version of Theorem 1 which, among others, yields that the adhesion sets of the tree-decomposition intersect 'not more than necessary':

Theorem 4.4.1. *Every graph G of finite tree-width admits a rooted tree-decomposition $(T, \mathcal{V})$ into finite parts that is linked, tight and componental. Moreover, we may assume that*

(1) *for every $e <_T e' \in E(T)$ with $|V_e| \leqslant |V_{e'}|$, each vertex of $V_e \cap V_{e'}$ either dominates some end of G that lives in $G{\uparrow}e'$, or is contained in a critical vertex*[2] *set of G that is included in $G{\uparrow}e'$.*

Halin [82, Theorem 2] showed that every locally finite, connected graph has a linked ray-

[2] A set X of vertices of G is *critical* if infinitely many components of $G - X$ have neighbourhood X in G.

decomposition[3] into finite parts with disjoint adhesion sets. He used this result in [83, Satz 10] to establish Theorem 4.4.1 for locally finite graphs with at most two ends, replacing (1) by the stronger condition of having disjoint adhesion sets. In light of this, we discuss here that (1) describes how close one may come to having 'disjoint adhesion sets' in the general case.

If G is not locally finite, we generally cannot require the tree-decomposition $(T, \mathcal{V})$ in Theorem 1 to have disjoint adhesion sets while having finite parts, as every dominating vertex of an end ε will be eventually contained in all adhesion sets along the ray of T which arises from ε. Moreover, as every critical vertex set has to lie in an adhesion set of any tree-decomposition into finite parts, and since the tree-decomposition is linked, one can easily check that the adhesion sets also intersect in critical vertex sets. Thus one might hope to obtain a tree-decomposition as in Theorem 1 that satisfies the following condition:

(1') for every $e <_T e' \in E(T)$ each vertex of $V_e \cap V_{e'}$ either dominates some end of G that lives in $G{\uparrow}e'$, or is contained in a critical vertex set of G that is included in $G{\uparrow}e'$.

But (1) allows for more than (1'): If $e <_T e' \in T$ and $|V_e| > |V_{e'}|$, then V_e and $V_{e'}$ are allowed to intersect also in vertices that do not dominate an end and that are not contained in a critical vertex set. The following example shows that allowing this is in fact necessary. It presents a locally finite graph that does not admit a tree-decomposition $(T, \mathcal{V})$ as in Theorem 1 with (1'), the stronger version of (1). As locally finite graphs do not have any dominating vertices and critical vertex sets, (1') boils down to the property that the tree-decomposition $(T, \mathcal{V})$ has *upwards disjoint adhesion sets*, that is $V_e \cap V_{e'} = \emptyset$ for every $e <_T e' \in E(T)$.

Example 4.4.2. There is a locally finite connected graph G which does not admit a linked, tight, componental rooted tree-decomposition $(T, \mathcal{V})$ into finite parts with upwards disjoint adhesion sets, i.e. one which satisfies (1').

Proof. Let $c' \subseteq c_{00}(\mathbb{N})$ be the set of all the sequences $\mathcal{S} = (s_n)_{n\in\mathbb{N}}$ for which there exists $N \in \mathbb{N}$ such that $s_i \geqslant 1$ for all $i \leqslant N$ and $s_i = 0$ for all $i > N$.[4] Let G be the graph depicted in Figure 4.4, that is the graph on the vertex set

$$V(G) := \{((s_n)_{n\in\mathbb{N}}, i, j) \in c' \times \mathbb{N} \times \{1, 2, 3\}\}$$

and with edges between $(\mathcal{S}, i, j)$ and $(\mathcal{S}, i', j')$ whenever $|i - i'| + |j - j'| = 1$ and, for $\mathcal{S} = (s_0, \dots, s_{n-1}, 0, \dots)$ and $\mathcal{S}' = (s'_0, \dots, s'_{n-1}, s'_n, 0, \dots)$, with edges between $(\mathcal{S}, i-1, 3)$ and $(\mathcal{S}', 0, j')$ and between $(\mathcal{S}, i, 3)$ and $(\mathcal{S}', 0, j')$ whenever $s'_n = i \geqslant 1$ and $s_k = s'_k$ for all $k < n$. Note that G is locally finite.

[3] A *ray-decomposition* is a tree-decomposition whose decomposition tree is a ray.
[4] We restrict to the set c' instead of $c_{00}(\mathbb{N})$ to ensure that G is connected.

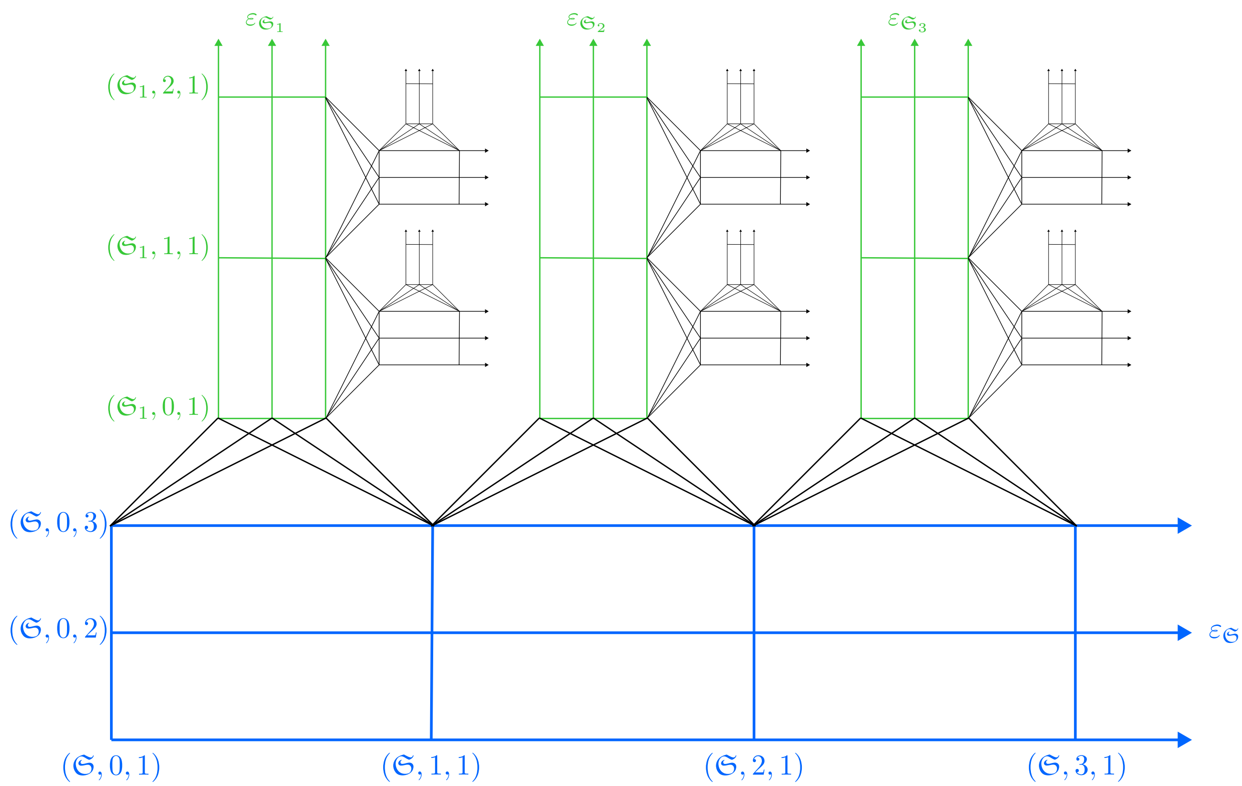

FIGURE 4.4: Depicted is the graph G from Example 4.4.2. The blue subgraph is induced by the vertices $(\mathfrak{S}, i, j)$ of G with $\mathfrak{S} = (0, 0, 0, \dots)$. The green subgraphs are induced by the vertices $(\mathfrak{S}_x, i, j)$ of G with $\mathfrak{S}_x = (x, 0, \dots)$ for $x \in \mathbb{N}_{\geqslant 1}$, respectively.

By Lemma 3.3.4 every linked, tight, componental rooted tree-decomposition of G into finite parts displays all the ends of G and their (combined) degrees. Thus, it suffices to show that G has no rooted tree-decomposition with upwards disjoint adhesion sets which displays all its ends and their (combined) degrees. Let $(T, \mathcal{V})$ be a tree-decomposition of G which displays all ends of G and their (combined) degrees. Consider the rays $R_{\mathcal{S},j} = \{(\mathcal{S}, i, j) \mid i \in \mathbb{N}\}$ for all $\mathcal{S} \in c'$ with $j \in \{1, 2, 3\}$. Then for every fixed $\mathcal{S} \in c'$ the rays $R_{\mathcal{S},1}, R_{\mathcal{S},2}, R_{\mathcal{S},3}$ all belong to the same end $\varepsilon_{\mathcal{S}}$ of G; and these ends $\varepsilon_{\mathcal{S}}$ are pairwise distinct and have (combined) degree 3. Since $(T, \mathcal{V})$ displays all ends of G and their (combined) degrees, there exist for each $\varepsilon_{\mathcal{S}}$ infinitely many edges e of T such that V_e has size three and meets every ray $R_{\mathcal{S},j}$; we fix for each end $\varepsilon_{\mathcal{S}}$ one such edge $e_{\mathcal{S}}$. Let $v_{\mathcal{S}} = (\mathcal{S}, i_{\mathcal{S}}, 3)$ be the (unique) vertex in $V_{e_{\mathcal{S}}} \cap V(R_{\mathcal{S},3})$.

Set $\mathcal{S}_0 := (0, 0, \dots)$ and $\mathcal{S}_n := (s_0, \dots, s_{n-1}, i_{\mathcal{S}_{n-1}} + 1, 0, \dots)$ where $\mathcal{S}_{n-1} = (s_0, \dots, s_{n-1}, 0, \dots)$. Then the $v_{\mathcal{S}_n}$ define a (unique) end ε of G, in that every ray that meets all the $v_{\mathcal{S}_n}$ belongs to the same end ε. This end has degree 2 as witnessed by the sets $S_n := \{(\mathcal{S}_n, i_{\mathcal{S}_n}, 3), (\mathcal{S}_n, i_{\mathcal{S}_n} + 1, 3)\}$ (see Figure 4.5).

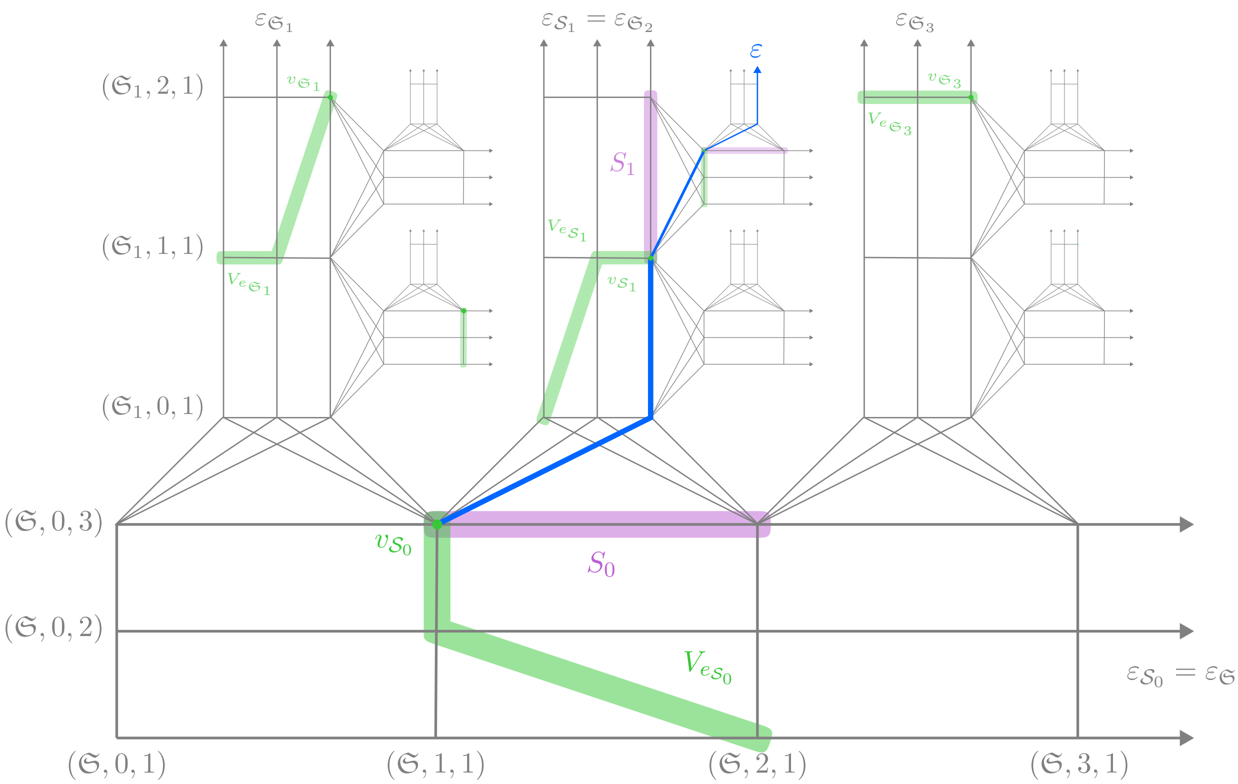

Figure 4.5: A ray in G that meets all the sets $V_{e_{\mathcal{S}_n}}$. The end to which it belongs has degree 2.

Now since $(T, \mathcal{V})$ displays the (combined) degree of ε, and because the sets S_n are the only separators witnessing that ε has degree 2, there is an edge $e \in T$ with $V_e = S_n$ for some $n \in \mathbb{N}$. But since every such S_n meets the set $V_{e_{\mathcal{S}_n}}$, the tree-decomposition $(T, \mathcal{V})$ does not have upwards disjoint separators. $\square$

This example might explain why Halin never extended his result [83, Satz 10] mentioned above to graphs with more than two ends: His precursor notion to a tree-decomposition, namely the quasi-trees and pseudo-trees discussed in [85], required upwards disjoint separators, and hence could not possibly capture the types of locally finite graphs described in Example 4.4.2.

5 Tangle-tree duality in infinite graphs

We extend Robertson and Seymour's tangle-tree duality theorem to infinite graphs.

This chapter is based on [5].

5.1 Introduction

All graphs in this chapter may be infinite, unless otherwise stated.

Tree-decompositions are a central object in structural graph theory. They were not only a crucial tool in the Graph Minor Project of Robertson and Seymour [110], but also attracted attention as several computationally hard problems can be solved efficiently on graphs of small tree-width. Because of this, the question arose which graphs have *small tree-width*, that is, admit a tree-decomposition into bags that all contain only few vertices, and, conversely, what kind of substructures prevent a graph from having small tree-width.

There are a number of substructures, e.g. large grid or clique minors, or k-blocks for large k, that are known to force a graph to have large tree-width. While these objects differ in their concrete shape, they have one thing in common: they witness high cohesion somewhere in the graph.

In their Graph Minors Project [110], Robertson and Seymour introduced *tangles* as a unified way to capture all such highly cohesive substructures in a graph. Formally, a k-tangle in a graph G is a certain orientation of all its separations of order less than k. The idea is that every highly cohesive substructure of G will lie mostly on one side of such a low-order separation, and therefore orient it towards that side. All these orientations, collectively, are then called a tangle.

One of the two major theorems in Robertson and Seymour's original work on tangles is the following duality between tangles of high order and small tree-width [114], rephrased here in the terminology of [41]:

Theorem 5.1.1. *For every finite graph G and $k \in \mathbb{N}$, exactly one of the following assertions holds:*

(i) *There exists a k-tangle in G.*

(ii) *There exists an $S_k(G)$-tree over $\mathcal{T}^*$.*

Theorem 5.1.1 is known as the *Tangle-tree duality theorem.* For a definition of $S_k(G)$-trees over $\mathcal{T}^*$ we refer the reader to Section 5.2. Theorem 5.1.1 implies an approximate duality for tangles and tree-width: every graph G with a k-tangle has tree-width at least $k - 1$, while a tree as in (ii) induces a tree-decomposition of G of width at most $3k - 4$.

The definition of a tangle extends verbatim to infinite graphs. There are several papers that extend results about tangles in finite graphs to infinite ones, or which deal with new questions that arise from tangles in infinite graphs only [31, 40, 64, 65, 90]. In this chapter we contribute to this a duality statement about tangles and small tree-width in infinite graphs, which extends Theorem 5.1.1 to infinite graphs.

The first thing to notice is that, in infinite graphs, high-order tangles no longer force the tree-width up. Indeed, every infinite graph G contains a tangle of infinite order [40], an orientation of all the finite-order separations of G. By restricting it to only those oriented separations that have order less than k, every such tangle induces a k-tangle for every $k \in \mathbb{N}$. Hence, every infinite graph has a k-tangle for every $k \in \mathbb{N}$. However, there are infinite graphs, e.g. infinite trees, that have small tree-width. Thus, in contrast to finite graphs, high-order tangles in infinite graphs are in general not an obstruction to small tree-width. Specifically, infinite locally finite trees are an example of graphs that have both a 3-tangle and an $S_3(G)$-tree over $\mathcal{T}^*$, thus witnessing that Theorem 5.1.1 fails for infinite graphs.

So what is the difference between finite and infinite graphs that causes high-order tangles to force large tree-width in finite graphs but not in infinite graphs? In finite graphs, tangles arise *only* from highly connected substructures (which may be fuzzy) as indicated earlier. In infinite graphs, however, there are also tangles that arise from infinite phenomena of the graph that do not reflect high local cohesion.

Let us first consider locally finite graphs. Every infinite locally finite, connected graph G has an *end*, an equivalence class of rays in G where two rays are equivalent if they cannot be separated by deleting finitely many vertices. Every end induces an infinite tangle by orienting every finite-order separation to the side which contains a tail of one (equivalently each) of its rays [40]. The *degree* of an end is the maximum number of disjoint rays in it.

Ends of large degree do force the tree-width up: It is not difficult to see that every graph with an end of degree at least k has tree-width at least k. But this is sharp in the sense of Theorem 5.1.1: For every $k \in \mathbb{N}$, there exists a locally finite graph (e.g. the rectangular $(k-1) \times \infty$ grid) whose single end has degree $k-1$, and that has an $S_k(G)$-tree over $\mathcal{T}^*$. In particular, ends of small degree do not force the tree-width up. So if we want to extend Theorem 5.1.1 to locally finite graphs in a way that retains its duality between tree structure on the one hand and the existence of high local cohesion on the other, we need to adjust (i) to ban tangles that are induced by ends of small degree.

We also have to adjust (ii) of Theorem 5.1.1, for a different reason. Since no infinite graph has a finite $S_k(G)$-tree over $\mathcal{T}^*$, we have to allow infinite $S_k(G)$-trees in (ii) when we extend Theorem 5.1.1 to infinite graphs. But this creates another problem. For example, consider the graph G which is obtained from a ray on vertex set $\mathbb{N}$ by gluing a large clique K on to 0. Then G has an infinite

$S_4(G)$-tree (R, α) over $\mathcal{T}^*$, where R is the natural ray on vertex set $\mathbb{N}$ and $\alpha : \vec{E}(R) \to \vec{S}_4(G)$ with $\alpha(i, i+1) = (\{0, \dots, i+1\}, \mathbb{N}_{\geqslant i} \cup V(K))$. But G has large tree-width (and a high-order tangle), as witnessed by the large clique K. The problem is that, in contrast to finite $S_k(G)$-trees, (R, α) does not induce a tree-decomposition.

Since it is our aim to extend Theorem 5.1.1 to infinite graphs in a way that retains its duality between tree structure and the existence of high local cohesion, we need to exclude such $S_k(G)$-trees from (ii). This will be formalised by 'weakly exhaustive' $S_k(G)$-trees in Section 5.2: every weakly exhaustive $S_k(G)$-tree does induces a tree-decomposition.

Our tangle-tree duality theorem for locally finite graphs then reads as follows:

Theorem 11. *For every locally finite, connected graph G and $k \in \mathbb{N}$, exactly one of the following assertions holds:*

(i) *There exists a k-tangle in G that is not induced by an end of degree $< k$.*

(ii) *There exists a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^*$.*

Let us now consider arbitrary infinite graphs. There is another type of tangle that can occur in infinite graphs which also does not reflect any highly cohesive substructure. For example, let G be the edgeless graph[1] with vertex set $\mathbb{N}$. Let β be a non-principal ultrafilter on $\mathbb{N}$, and orient every 1-separation of G (a bipartition of $\mathbb{N}$) towards its side in β. This is a 1-tangle in G, since $\mathbb{N}$ is not a union of three subsets not in β. More generally, a k-tangle is *principal* if it contains for every set X of fewer than k vertices a separation of the form $(V(G) \setminus V(C), V(C) \cup X)$ where C is a component of $G - X$. A tangle is *non-principal* if it is not principal. As in our example, an infinite graph G contains a non-principal k-tangle if there is a set X of fewer than k vertices of G whose deletion separates G into infinitely many components; and every such tangle, one for each non-principal ultrafilter on the set of components of $G - X$, contains all the separations of the form $(V(C) \cup X, V(G) \setminus V(C))$ for components C of $G - X$ [40].[2] As we have seen, such non-principal tangles do not force a graph to have large tree-width, and hence give rise to counterexamples to Theorem 5.1.1 and Theorem 11 for arbitrary infinite graphs. Hence, for graphs that are not locally finite, we shall have to adjust (i) again, to ban non-principal tangles.

We shall have to adjust (i) in another way too. A vertex v of G *dominates* an end ε of G if no finite set of vertices other than v separates v from a ray in ε. The *combined degree* of an end is the sum of its degree and the number of its dominating vertices.[3] Similarly to ends of large degree, also ends of large combined degree force the tree-width up: It is not difficult to see that every

[1] If you do not like tangles of edgeless graphs, let G be a countably infinite disjoint union of copies of K_2.

[2] Note that such tangles cannot exist in connected locally finite graphs, where deleting finitely many vertices never leaves infinitely many components.

[3] Note that in locally finite graphs a vertex cannot dominate an end, so the combined degree of an end is simply its degree.

graph with an end of combined degree k has tree-width at least k. Hence, we need to adjust (i) as follows:

Theorem 12. *For every countable graph G and $k \in \mathbb{N}$, exactly one of the following holds:*

(i) *There exists a principal k-tangle in G that is not induced by an end of combined degree $< k$.*

(ii) *There exists a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^*$.*

Note that we restricted the graphs in Theorem 12 to those that are countable. We did so for a reason: There is no uncountable graph G that has a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^*$ for any $k \in \mathbb{N}$. Indeed, let $(T, \mathcal{V})$ be the tree-decomposition induced by any weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^*$. Then the definition of $\mathcal{T}^*$ ensures that the bags V_t of $(T, \mathcal{V})$ all have size $\leqslant 3k - 3$, and that the tree T has maximum degree at most 3, and hence is countable. But then G is countable.

However, there are uncountable graphs, e.g. stars with uncountably many leaves, that have no 3-tangles as in (i) of Theorem 12, and that even have tree-width 1. We could now try to update (i) again, so that our duality theorem always outputs (i) if the graph is uncountable; but then (i) would no longer capture high local cohesion in a graph, which remains our aim. So we need to adjust (ii).

As indicated earlier, the definition of $S_k(G)$-trees over $\mathcal{T}^*$ is too restrictive to capture uncountable graphs of small tree-width, as those $S_k(G)$-trees have maximum degree at most 3. Hence, we will allow the $S_k(G)$-tree in (ii) to have infinite-degree nodes. For this, we allow $S_k(G)$-trees over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$ rather than just $\mathcal{T}^*$, where $\mathcal{U}_k^\infty$ is the set of all infinite stars of separations of order $< k$ whose interior has size $< k$.[4] Now a tree-decomposition induced by a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$ may have nodes t of infinite degree as long as their bags V_t have size less than k. This modification of (ii) will be in line with our aim that (ii) describes graphs that have no large highly cohesive substructures: such graphs may have non-principal tangles, and these can now happily live in the $S_k(G)$-tree in its nodes of infinite degree.

Our tangle-tree duality theorem for arbitrary graphs now reads as follows:

Theorem 13. *For every graph G and $k \in \mathbb{N}$, exactly one of the following assertions holds:*

(i) *There exists a principal k-tangle in G that is not induced by an end of combined degree $< k$.*

(ii) *There exists a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$.*

We remark that Theorems 11 and 12 are simple applications of Theorem 13 (see Section 5.6 for details). In particular, Theorem 13 contains Theorem 5.1.1 as a special case. Indeed, in finite graphs all tangles are principal, so a finite graph satisfies (i) of Theorem 13 if and only if it satisfies (i) of Theorem 5.1.1. Moreover, in finite graphs the set $\mathcal{U}_k^\infty$ is empty and every finite

[4] Formally, $\mathcal{U}_k^\infty := \{\sigma = \{(A_i, B_i) : i \in I\} \subseteq \vec{S}_k(G) : \sigma$ is a star, $|\bigcap_{i \in I} B_i| < k$ and $|\sigma| = \infty\}$; see Section 5.2 for details.

$S_k(G)$-tree over $\mathcal{T}^*$ is weakly exhaustive. So a finite graph satisfies (ii) of Theorem 13 if and only if it satisfies (ii) of Theorem 5.1.1.

Moreover, similarly to Theorem 5.1.1, a tree as in (ii) induces a tree-decomposition of G of width at most $3k-4$, while every graph with a tangle as in (i) has tree-width at least $k-1$.

Diestel and Oum [55] generalized Theorem 5.1.1 to so-called '$\mathcal{F}$-tangles'. The 'standard' k-tangles are orientations of the separations of a graph G of order $< k$ that *avoid* the set $\mathcal{T}^* \subseteq 2^{\vec{S}_k(G)}$: a k-tangle does not contain an element of $\mathcal{T}^*$ as a subset. This set $\mathcal{T}^*$ can be replaced by more general sets of separations, leading to the more general notion of $\mathcal{F}$-tangles.

We will in fact prove Theorem 13 more generally for $\mathcal{F}$-tangles (see Theorem 16 in Section 5.6 for the precise statement). As an application, we obtain the following exact characterization of graphs that have tree-width $k \in \mathbb{N}$, which generalizes a result of Diestel and Oum [54]:

Theorem 14. *The following assertions are equivalent for all graphs G and $k \in \mathbb{N}$:*

(i) *G has a $\mathcal{U}_k$-tangle of order k that is not induced by an end of combined degree $< k$.*

(ii) *G has a finite bramble of order at least k.*

(iii) *G has no weakly exhaustive $S_k(G)$-tree over $\mathcal{U}_k$.*

(iv) *G has tree-width at least $k-1$.*

(See Section 5.2 for definitions.) The equivalence of (ii) and (iv) yields a generalization of the 'bramble-treewidth duality theorem' of Seymour and Thomas [117] (see Theorem 5.7.2 in Section 5.7), which also includes their finite version as a corollary (without using it in the proof).

The other major theorem about tangles which Robertson and Seymour [114] proved is the *tree-of-tangles theorem*. Recall that a separation $\{A, B\}$ of a graph G *distinguishes* two tangles in G if they orient $\{A, B\}$ differently. It distinguishes them *efficiently* if they are not distinguished by any separation of smaller order.

The tree-of-tangles theorem for fixed $k \in \mathbb{N}$ asserts that every finite graph G has a tree-decomposition $(T, \mathcal{V})$ which efficiently distinguishes all its k-tangles: for every pair τ, τ' of k-tangles in G, there is an edge e of T such that the separation induced by e distinguishes τ and τ' efficiently.

Following [65], we call two k-tangles *combinatorially* distinguishable if there is a finite set $X \subseteq V(G)$ and a component C of $G - X$ such that $\{V(C) \cup X, V(G - C)\}$ distinguishes them. In particular, if two k-tangles are combinatorially indistinguishable, then they are both non-principal. For instance, in the example mentioned above, where G is the edgeless graph on vertex set $\mathbb{N}$, no two 1-tangles in G are combinatorially distinguishable.

We show that if a graph G has no k-tangle as in (i) of Theorem 13, then it has an $S_k(G)$-tree as in (ii) which additionally distinguishes all the combinatorially distinguishable k-tangles that G may have: those that are not of the form as in (i):

Theorem 15. *Let G be a graph and $k \in \mathbb{N}$. Suppose that all principal k-tangles in G are induced by ends of combined degree $< k$. Then G has a weakly exhaustive $S_k(G)$-tree (T, α) over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$ such that*

(i) *every end of G lives in an end of T, and conversely every end of T is home to an end of G,*

(ii) *every non-principal k-tangle lives at a node t of T with $\{\alpha(t', t) : t' \in N_T(t)\} \in \mathcal{U}_k^\infty$, and*

(iii) *for every pair τ, τ' of combinatorially distinguishable k-tangles in G there is an edge e of T such that $\alpha(e)$ distinguishes τ and τ' efficiently.*

In particular, (iii) ensures that no two ends of G live in the same end of T, and that no two combinatorially distinguishable non-principal k-tangles live at the same node of T. We remark that it is not possible to strengthen (iii) so that all k-tangles are efficiently distinguished by a separation of the form $\alpha(e)$ for an edge e of T [65, Corollary 3.4].

We will obtain Theorem 15 as a corollary of a more general theorem (see Theorem 5.8.1 in Section 5.8) that yields a tree-decomposition with similar properties as the $S_k(G)$-tree in Theorem 15 even if G has other k-tangles than those allowed in the premise of Theorem 15.

This chapter is organised as follows. We first recall some definitions in Section 5.2. Then we give a brief introduction to infinite graphs and their tangles in Sections 5.3 and 5.4. In Section 5.5, we first sketch the proof of Theorem 13 briefly, and then prove Lemma 5.5.4, which is one of the two main ingredients to the proof of Theorem 13. In Section 5.6, we prove Theorem 13, and then derive Theorems 11 and 12 from it. In Section 5.7 we deduce Theorem 14. In Section 5.8 we use the tools developed in Section 5.5 to show a 'refined' version of the *Tree-of-tangles theorem*, which generalizes a result of Chapter 6 to infinite graphs, and which contains Theorem 15 as a special case.

5.2 Preliminaries

An infinite increasing sequence $((A_i, B_i))_{i \in \mathbb{N}}$ of separations of a graph G is *weakly exhaustive* if the intersection of their strict big sides is empty, i.e. if $\bigcap_{i \in \mathbb{N}} B_i \setminus A_i = \emptyset$. Further, an $S_{\aleph_0}(G)$-tree (T, α) is *weakly exhaustive* if for every ray $R = r_1 r_2 \dots$ in T the sequence $(\alpha(r_i, r_{i+1}))_{i \in \mathbb{N}}$ is weakly exhaustive.

In this chapter we will need to consider a slight weakening of the profile property $(*)$ and consider $\mathcal{P}'_k$-tangles instead of k-profiles where

$$\mathcal{P}'_k := \{\sigma = \{(A_i, B_i) : i \in [3]\} \in \mathcal{P}_k : |\operatorname{int}(\sigma)| < k\}$$

is a subset of the set $\mathcal{P}_k$ defined in (2.1).

5.2.1 Principal tangles

Given some graph G and $k \in \mathbb{N} \cup \{\aleph_0\}$, an orientation of $S_k(G)$ is *principal* if it contains for every set X of fewer than k vertices a separation of the form $(V(G-K), V(K) \cup X)$ where K is a component of $G - X$. It is easy to check that the regular, principal $\mathcal{P}'_k$-tangles of $S_k(G)$ avoid

$$\mathcal{U}_k := \{\sigma = \{(A_i, B_i) : i \in I\} \subseteq \vec{S}_k(G) : \sigma \text{ is a star and } |\operatorname{int}(\sigma)| < k\}.$$

Write $\mathcal{U}_k^\infty := \{\sigma \in \mathcal{U}_k : |\sigma| = \infty\}$ for the the set of all infinite stars in $\mathcal{U}_k$.

Lemma 5.2.1. *Let τ be a regular, non-principal $\mathcal{P}'_k$-tangle of order $k \in \mathbb{N}$ in a graph G, and let $\sigma \subseteq \vec{S}_k(G)$ be a finite star with finite interior. Then $\sigma \not\subseteq \tau$.*

Proof. Suppose for a contradiction that $\sigma \subseteq \tau$. Since τ is non-principal, there is some $X \subseteq V(G)$ such that $(V(K) \cup X, V(G-K))$ for all components K of $G-X$; note that $|X| < k$ as τ has order k. By the consistency of τ, we have $\big(V\big(\bigcup \mathcal{K}_{(A,B)}\big) \cup X, V\big(G - \bigcup \mathcal{K}_{(A,B)}\big)\big) \in \tau$ for all $(A,B) \in \sigma$ where $\mathcal{K}_{(A,B)}$ is the set of all components of $G - X$ that are contained in $G[A \setminus B]$. Since $|\sigma|$ is finite, inductively applying that τ is a $\mathcal{P}'_k$-tangle yields that $\big(V\big(\bigcup \mathcal{K}\big) \cup X, V\big(G - \bigcup \mathcal{K}\big)\big) \in \tau$ where $\mathcal{K} := \bigcup_{(A,B) \in \sigma} \mathcal{K}_{(A,B)}$. As $\operatorname{int}(\sigma)$ is finite, at most finitely many components of $G - X$ are not in $\mathcal{K}$, so the same inductive argument yields that $(V(G), X) \in \tau$, which contradicts that τ is regular. □

5.2.2 Nice sets of stars

Let G be a graph, and $k \in \mathbb{N}$. Strengthening the notion 'closed under shifting' (see Section 2.3), we say that a set $\mathcal{F}$ of stars in $\vec{S}_{\aleph_0}(G)$ is *strongly closed under shifting in* $\vec{S}_k(G)$ if whenever $\vec{s} \in \vec{S}_k(G)$ emulates some $\vec{r} \in \vec{S}_{2k-1}(G)$ in $S_k(G)$, then it also emulates $\vec{r}$ in $\vec{S}_k(G)$ for $\mathcal{F}$.

Proposition 5.2.2. *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a set of stars in $\vec{S}_k(G)$ that is strongly closed under shifting. Further, let P be a k-profile in G, let $\overleftarrow{s} \in P$, and let $\{\vec{t}\} \cup \{\vec{r}_i : i \in I\} \in \mathcal{F}$ be such that $\overleftarrow{t} \in P$. Suppose further that $\vec{s} \vee \vec{x} \in \vec{S}_k(G)$ for all $\overleftarrow{x} \in P$. Then $\{\vec{s} \vee \vec{t}\} \cup \{\overleftarrow{s} \wedge \vec{r}_i : i \in I\}$ is a star in $\mathcal{F}$.*

Proof. We show that $\vec{s}$ emulates $\vec{s} \wedge \vec{t}$ in $S_k(G)$ from which the assertion follows as $\mathcal{F}$ is strongly closed under shifting and $|\vec{s} \wedge \vec{t}| \leqslant |s| + |t| < 2k-1$. For this let $\vec{x} \in \vec{S}_k(G)$ with $\vec{x} \geqslant \vec{s} \wedge \vec{t}$ be given. Then $\overleftarrow{x} \in P$ by [60, Theorem 1]. So by the assumptions on s, it follows that $\vec{s} \vee \vec{x} \in S_k(G)$. □

A set $\mathcal{F}$ of stars in $\vec{S}_k(G)$ is called *nice* if $\mathcal{F}$ is strongly closed under shifting in $\vec{S}_k(G)$, $\{(V(G), A)\} \in \mathcal{F}$ for all subsets $A \subseteq V(G)$ of fewer than k vertices, and $\mathcal{P}'_k \subseteq \mathcal{F}$.

Lemma 5.2.3. *For every $k > 0$, the sets $\mathcal{U}_k$ and $\mathcal{T}_k^*$ are nice.*

Proof. By definition, $\mathcal{P}'_k \subseteq \mathcal{U}_k \cap \mathcal{T}^*_k$, and we have $\{(V(G), A)\} \in \mathcal{U}_k \cap \mathcal{T}^*_k$ for all sets A of fewer than k vertices. The proofs that $\mathcal{U}_k$ and $\mathcal{T}^*_k$ are strongly closed under shifting in $\vec{S}_k(G)$ are analogous to the proofs that they are closed under shifting (see the proofs of [54, Lemma 6.1 & Theorem 4.1]). □

5.3 Tangles and ends

It is easy to see that every end ε of a graph G *induces* an $\aleph_0$-tangle, denoted by τ_ε, by orienting every finite-order separation of G to that side which contains a tail of some (equivalently every) ray in ε [40, §1]. We say that an orientation O of $S_k(G)$, for some $k \in \mathbb{N}$, is *induced* by an end ε of G if $\tau \subseteq \tau_\varepsilon$.

It is easy to check that the tangle τ_ε induced by an end ε of G cannot contain any stars with finite interior (cf. [40, Discussion preceding Lemma 1.6]). Hence, the following observation about ends follows easily.

Proposition 5.3.1. *Let G be any graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a set of stars in $\vec{S}_k(G)$ all of which have finite interior. Then every end of G induces an $\mathcal{F}$-tangle of $S_k(G)$.* □

We say that an end ε of a graph G *lives* in a star $\sigma \subseteq \vec{S}_{\aleph_0}(G)$ (or equivalently σ *is home to* ε) if $\sigma \subseteq \tau_\varepsilon$.

Let $(T, \mathcal{V})$ be a tree-decomposition of G whose adhesion sets are all finite, and let ε be an end of G. Let $O \subseteq \vec{E}(T)$ consist of those oriented edges (s,t) of T such that $(U_s, U_t) \in \tau_\varepsilon$. Then the orientation O of $E(T)$ points towards a node of T or to an end of T. We say that ε *lives* at that node or in that end, respectively.

The following lemma describes when an end has small combined degree.

Lemma 5.3.2 ([77, Corollary 5.8])**.** *Let ε be an end of a graph G, and let $k \in \mathbb{N}$. Then the following assertions hold:*

(i) *If τ_ε contains a weakly exhaustive increasing sequence of separations of order $\leqslant k$, then $\Delta(\varepsilon) \leqslant k$.*

(ii) *If $\Delta(\varepsilon) = k$, then τ_ε contains a weakly exhaustive increasing sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations of order k such that $G[B_i \setminus A_i]$ is connected and $(A_i \cap B_i) \cap (A_j \cap B_j) \in \mathrm{Dom}(\varepsilon)$ for all $i \neq j \in \mathbb{N}$.*

We need the following lemma which describes a condition that is sufficient to ensure that all ends of a torso have small degree.

Lemma 5.3.3. *Let σ be a star of separations of order $< k \in \mathbb{N}$ of some graph G such that every separation in σ distinguishes some pair of k-profiles in G efficiently. Then for every end ε' of* torso(σ) *there exists an end* ε *of* G *such that* $\sigma \subseteq \tau_\varepsilon$*,* $\Delta(\varepsilon') = \Delta(\varepsilon)$ *and for every* $(A,B) \in \tau_\varepsilon$ *we have* $(A \cap \mathrm{int}(\sigma), B \cap \mathrm{int}(\sigma)) \in \tau_{\varepsilon'}$ *if* $\{A \cap \mathrm{int}(\sigma), B \cap \mathrm{int}(\sigma)\}$ *is a separation of* torso(σ)*.*

Proof. Let R' be some ε'-ray. Since all separations in σ efficiently distinguish some pair of k-profiles, they are left-tight by Lemma 2.2.1, so by Proposition 3.6.1 applied to σ and $P := R'$ there is a ray R in G such that $|V(R) \cap V(R')| = \infty$. Let ε be the end of G to which R belongs. Then $\sigma \subseteq \tau_\varepsilon$, because $V(R) \cap V(R') \subseteq \mathrm{int}(\sigma)$ is infinite, and thus R has a tail in $G[B \setminus A]$ for every $(A,B) \in \sigma$. Moreover, for all $(A,B) \in \tau_\varepsilon$ that induce a separation $\{A \cap \mathrm{int}(\sigma), B \cap \mathrm{int}(\sigma)\}$ of torso(σ), we have $(A \cap \mathrm{int}(\sigma), B \cap \mathrm{int}(\sigma)) \in \tau_{\varepsilon'}$ for the same reason.

Since connected subgraphs of G induce connected subgraphs of torso(σ), every ε-ray induces an ε'-ray and $\mathrm{Dom}(\varepsilon) \subseteq \mathrm{Dom}(\varepsilon')$. Hence, $\Delta(\varepsilon) \leqslant \Delta(\varepsilon')$. We claim that also $\Delta(\varepsilon') \leqslant \Delta(\varepsilon)$, which concludes the proof. So suppose for a contradiction that $\Delta(\varepsilon') > \Delta(\varepsilon)$. Let $U \subseteq \mathrm{int}(\sigma)$ be a set of size $n := \Delta(\varepsilon) + 1 \leqslant \Delta(\varepsilon')$ and $\mathcal{P} := \{P_x : x \in U\}$ a family of n pairwise disjoint paths/rays in torso(σ) such that P_x is either an ε'-ray that starts in x or the trivial path whose single vertex x lies in $\mathrm{Dom}(\varepsilon')$. As U is finite, there exists by Lemma 5.3.2 (ii) a separation $(A,B) \in \tau_\varepsilon$ of order $\Delta(\varepsilon)$ such that $U \subseteq A_n$ and $G[B \setminus A]$ is connected.

Set $\varrho := \{(C,D) \in \sigma : \{A,B\}, \{C,D\} \text{ cross}\}$ and $\varrho' := \{(C,D) \in \varrho : C \cap D \cap (A \setminus B) \neq \emptyset\}$. Let us first show that ϱ' is finite. For this, it suffices to show that each of the pairwise disjoint strict small sides $C \setminus D$ of $(C,D) \in \varrho'$ meets the finite set $A \cap B$. Since (C,D) is left-tight, there is component $K \subseteq G[C \setminus D]$ such that $N_G(K) = C \cap D$. In particular, since $C \cap D \cap (A \setminus B) \neq \emptyset$, we have that $A \cap B$ meets K if also $C \cap D \cap (B \setminus A) \neq \emptyset$. So we may assume that $B \setminus A$ avoids $C \cap D$. Then either $B \setminus A \subseteq C \setminus D$ or $B \setminus A \subseteq D \setminus C$ as $G[B \setminus A]$ is connected, which implies that $\{A,B\}, \{C,D\}$ are nested and contradicts $(C,D) \in \varrho$.

Now set $(\bar{A}, \bar{B}) := (A,B) \wedge \bigwedge_{(C,D) \in \varrho \setminus \varrho'} (D,C)$ and $(\tilde{A}, \tilde{B}) := (\bar{A}, \bar{B}) \wedge \bigwedge_{(C,D) \in \varrho'} (D,C)$. Then we have $|\bar{A} \cap \bar{B}| = |A \cap B|$ because $(A \setminus B) \cap (C \cap D) = \emptyset$ for all $(C,D) \in \varrho \setminus \varrho'$, and $|\tilde{A} \cap \tilde{B}| \leqslant |(\bar{A} \cap \bar{B}) \cup \bigcup_{(C,D) \in \varrho'} (C \cap D)| < \infty$ because ϱ' is finite. So τ_ε contains an orientation of $\{\bar{A}, \bar{B}\}$, $\{\tilde{A}, \tilde{B}\}$; since τ_ε is consistent, we find $(\tilde{A}, \tilde{B}) \leqslant (\bar{A}, \bar{B}) \leqslant (A,B) \in \tau_\varepsilon$. Since $U \subseteq \mathrm{int}(\sigma) \subseteq D$ for all $(C,D) \in \varrho$, we also have $U \subseteq \tilde{A}$. Moreover, by Lemma 2.1.1, $(\tilde{A}, \tilde{B})$ is nested with σ, and $(\bar{A}, \bar{B})$ crosses at most those finitely many separations in σ that are contained in ϱ'. It follows by Claim 1[5] (in the proof of Lemma 6.5.3 in Chapter 6) applied to τ_ε, σ and $(\tilde{A}, \tilde{B}) \leqslant (\bar{A}, \bar{B})$ that there is a separation $(A', B') \in \tau_\varepsilon$ of order $\leqslant |\bar{A} \cap \bar{B}| = |A \cap B| = \Delta(\varepsilon)$ such that $(\tilde{A}, \tilde{B}) \leqslant (A', B')$, and thus $U \subseteq \tilde{A} \subseteq A'$. Since $(A,B) \in \tau_\varepsilon$, its big side $G[B]$ contains a tail of R, and hence $V(R'') \subseteq B$ for some tail R'' of R'. By the choice of P_x, there exists an infinite family of P_x–R'' paths in torso(σ) that meet at most in there endvertices on P_x. In particular, since (A', B') is nested with σ and thus induces a separation of torso(σ), we find $x \in B'$ if P_x is a trivial path, or

[5]In Claim 1 this is shown for finite graphs G, but the proof only uses the finiteness of G to conclude that $(\bar{A}, \bar{B})$ crosses at most finitely many separations in σ, which we have proved separately.

$V(P'_x) \subseteq B'$ for some tail P'_x of P_x if P_x is a ray. But since each P_x is connected, it follows that $V(P_x) \cap (A' \cap B') \neq \emptyset$. Since the P_x are pairwise disjoint, this implies that $|A' \cap B'| \geqslant \Delta(\varepsilon) + 1$, a contradiction. □

5.4 Tangles and critical vertex sets

Let us recall the following lemma (Lemma 3.2.5), which we proved in Chapter 3:

Lemma 5.4.1. *Let σ a star of left-tight, finite-order separations of a graph G. Then* $\mathrm{crit}(\mathrm{torso}(\sigma)) \subseteq \mathrm{crit}(G)$.

In the previous subsection we have seen that every end induces an $\aleph_0$-tangle. The following lemma asserts that also graphs with critical vertex sets have $\aleph_0$-tangles: for every critical vertex set X, every free ultrafilter on $\mathcal{C}_X$ induces an $\aleph_0$-tangle.[6]

Lemma 5.4.2 ([40, Lemmas 3.4 & 3.7])**.** *Given some finite set X of vertices of a graph G, for each free ultrafilter U on $\mathcal{C}_X$ there exists a (non-principal) $\aleph_0$-tangle τ in G such that, for all $\mathcal{K} \subseteq \mathcal{C}_X$, we have $(V(G - \bigcup \mathcal{K}), V(\bigcup \mathcal{K}) \cup X) \in \tau$ if and only if $\mathcal{K} \in U$. In particular, then*

$$\tau = \{(A, B) \in \vec{S}_{\aleph_0}(G) \mid \exists \mathcal{K} \in U : \bigcup \mathcal{K} \subseteq G[B]\}.$$

Even though the $\aleph_0$-tangles described in Lemma 5.4.2 are non-principal, critical vertex sets still induce principal k-tangles, but only for $k \in \mathbb{N}$ that are not greater than their size.

Lemma 5.4.3. *Let $k \in \mathbb{N}$, let G be any graph, and let $\mathcal{F}$ be a set of stars in $\vec{S}_k(G)$ all of which have finite interior. Further, let X be a critical vertex set of G of size $\geqslant k$. Then $\tau := \{(A, B) \in \vec{S}_k(G) : X \subseteq B\}$ is a principal $\mathcal{F}$-tangle of $S_k(G)$. In particular, if $\sigma \subseteq \vec{S}_k(G)$ is a star with $X \subseteq \mathrm{int}(\sigma)$, then $\sigma \subseteq \tau$.*

Proof. Since X is a critical vertex set and thus infinitely connected in G, and because $|X| \geqslant k$, every separation in $S_k(G)$ has a unique side which contains X. Hence, τ is an orientation of $S_k(G)$, which for the same reason is consistent. Now let σ be a star contained in τ. Since $X \cap (B \setminus A) \neq \emptyset$ for all $(A, B) \in \sigma$, every component of $G - X$ whose neighbourhood in G equals X meets B. Hence, as components are connected, each such component of $G - X$ meets $\mathrm{int}(\sigma)$. Since X is a critical vertex set of G, this implies that $|\mathrm{int}(\sigma)| = \infty$, so σ is not in $\mathcal{F}$. In particular, τ avoids $\mathcal{U}_k$ and is thus principal. □

We also need the following lemma, which describes a sufficient condition for a star to be home to a (non-principal) tangle.

[6] These $\aleph_0$-tangles are called *ultrafilter tangles* [40].

Lemma 5.4.4. *Let G be any graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a set of finite stars in $\vec{S}_k(G)$ all of which have finite interior. Further, let $\sigma \subseteq \vec{S}_k(G)$ be a star of left-tight separations, and let $X \subseteq \operatorname{int}(\sigma)$ be of size $< k$. Suppose that either X is a critical vertex set of* $\operatorname{torso}(\sigma)$ *or that infinitely many separations in σ have separator X. Then there is a (non-principal) $\mathcal{F}$-tangle of $S_k(G)$ that lives in σ.*

Proof. Set $\mathfrak{K} := \{\mathcal{K} \subseteq \mathcal{C}_X : |\mathcal{K}| = 1 \text{ or } \bigcup \mathcal{K} \subseteq G[A \smallsetminus B] \text{ for some } (A,B) \in \sigma\}$. Note that by Lemma 5.4.1 if $X \in \operatorname{crit}(\operatorname{torso}(\sigma))$ or by the left-tightness of the separations in σ if the set $\{(A,B) \in \sigma : A \cap B = X\}$ is infinite, X is critical in G. In particular, every collection $\mathcal{K} \in F := \{\mathcal{C}_X\} \cup \{\mathcal{C}_X \smallsetminus (\mathcal{K}_1 \cup \dots \cup \mathcal{K}_n) : n \in \mathbb{N}, \mathcal{K}_1, \dots, \mathcal{K}_n \in \mathfrak{K}\}$ still contains infinitely many components of $G - X$. In particular, $\emptyset \notin F$. Moreover, by definition, F is closed in $\mathcal{C}_X$ under taking supersets and under finite intersections. Thus, F is a filter on $\mathcal{C}_X$. So by Lemma 5.4.2 (and the ultrafilter lemma), there is an $\aleph_0$-tangle τ in G such that $(V(G - \mathcal{K}), V(\bigcup \mathcal{K}) \cup X) \in \tau$ for all $\mathcal{K} \in F$. By [40, Lemma 1.3 & Corollary 1.5], τ avoids $\mathcal{F}$. Moreover, by Lemma 5.4.2, there exists for every $(A,B) \in \tau$ a collection $\mathcal{K} \subseteq \mathcal{C}_X$ such that $\bigcup \mathcal{K} \subseteq G[B]$ and $\mathcal{K} \notin \mathfrak{K}$. By the definition of $\mathfrak{K}$, this implies that $(B,A) \notin \tau$ for all $(A,B) \in \sigma$, and hence $\sigma \subseteq \tau$. So $\tau \cap \vec{S}_k(G)$ is as desired. □

5.5 Refining inessential stars in infinite graphs

In this section we prove Lemma 5.5.4, which is one of the two main ingredients to the proof of Theorem 13. Let us begin by giving a brief sketch of the proof of Theorem 13. Similarly to the finite tangle-tree duality theorem, Theorem 5.1.1, it is not too difficult to show that not both, (i) and (ii) of Theorem 13, can hold at the same time. To see that at least one of the two assertions holds, we consider an arbitrary graph G without any k-tangles as in (i). For the proof that G then has an $S_k(G)$-tree as in (ii), we need two ingredients. The first one is a certain tree-decomposition of G, whose existence follows from a result in Chapter 3 (see Theorem 5.6.2 in Section 5.6 below): G admits a tree-decomposition $(T, \mathcal{V})$ of adhesion $< k$ into finite parts such that every node t of T with $\sigma_t \notin \mathcal{U}_k^\infty$ is inessential and has finite degree.

The second main ingredient to the proof of Theorem 13 is Lemma 5.5.4 below. This lemma ensures that, under mild additional assumptions on $(T, \mathcal{V})$, there exists, for every inessential node $t \in T$, a finite $S_k(G)$-tree (T^t, α^t) over $\mathcal{T}_k^* \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma_t\}$ in which each $\vec{s} \in \sigma_t$ appears as a leaf separation. In particular, all its non-leaves are associated with stars in $\mathcal{T}_k^*$. We then obtain a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}_k^* \cup \mathcal{U}_k^\infty$ by sticking the $S_k(G)$-trees (T^t, α^t) together along T. More precisely, for every edge $e = \{t, s\}$ of T we glue the trees T^t and T^s together along those leaf edges f of T^t and f' of T^s with $\alpha^t(f) = \{A, B\} = \alpha^s(f')$ where $\{A, B\}$ is the separation induced by the edge e of T (see Construction 5.6.3).

The idea of refining the inessential parts of a tree-decomposition $(T, \mathcal{V})$ with $S_k(G)$-trees as described above has its origin in [66]. There, Erde proved for finite graphs that if all edges of T induce separations that efficiently distinguish two k-tangles, then such $S_k(G)$-trees (T^t, α^t) exist for all inessential nodes t of T. The main result of this section generalizes his lemma ([66, Lemma 3.1]) not only to infinite graphs but also to certain tree-decompositions which no longer need to distinguish the k-tangles efficiently. To state this result formally, we need some further definitions.

First, we recall the definition of 'closely related' from [4]: Let G be any graph and $k \in \mathbb{N}$. A separation $(A, B) \in \vec{S}_k(G)$ is *closely related* to an orientation O of $S_k(G)$ if $(A, B) \in O$ and for every $(C, D) \in O$ we have $(A \cap C, B \cup D) \in \vec{S}_k(G)$.

Proposition 5.5.1 ([4, Proposition 3.4])**.** *Let $k \in \mathbb{N}$, and let P and P' be two k-profiles in a graph G. If a separation $(A, B) \in P$ distinguishes P and P' efficiently, then (A, B) and (B, A) are closely related to P and P', respectively.*

A finite-order separation (A, B) of a graph G is *left-ℓ-robust* for $\ell \in \mathbb{N}$ if there exist a set $U \subseteq A$ of size ℓ and a family $\{P_x : x \in A \cap B\}$ of pairwise disjoint paths in $G[A]$ such that P_x ends in x and there are ℓ U–P_x paths in $G[(A \setminus B) \cup \{x\}]$ that meet at most in their endvertices in P_x. An unoriented separation $\{A, B\}$ is ℓ-robust if both (A, B) and (B, A) are left-ℓ-robust.

Note that the property 'left-ℓ-robust' is designed to mimic the presence of a highly connected substructure of G on the small side of a separation. For example, we proved in Chapter 3 (Lemma 3.8.5) that we can obtain left-ℓ-robust separations from ends or critical vertex sets in the following way:

Lemma 5.5.2. *Let ε be an end of a graph G of finite combined degree, and let $((A_i, B_i))_{i \in \mathbb{N}}$ be a weakly exhaustive increasing sequence of separations in τ_ε such that $\liminf_{i \in \mathbb{N}} |A_i \cap B_i| = \Delta(\varepsilon)$. Then cofinitely many (A_i, B_i) with $|A_i \cap B_i| = \Delta(\varepsilon)$ are ℓ-robust.*

Proposition 5.5.3. *Let $\{A, B\}$ be a finite-order separation of a graph G, and suppose that $G[A \setminus B]$ contains infinitely many tight components of $G - X$ for some set $X \supseteq A \cap B$ of vertices of G. Then (A, B) is left-ℓ-robust for all $\ell \in \mathbb{N}$.*

Proof. This is witnessed by the trivial paths P_x in $A \cap B$ and a set U consisting of ℓ vertices that lie in pairwise distinct tight components of $G - X$ contained in $G[A \setminus B]$. $\square$

The definition of left-ℓ-robust is tailored to the proof of Theorem 13: 'left-ℓ-robust' is defined precisely so that we can prove Lemma 5.5.4 below, and, at the same time, show that there exists a tree-decomposition $(T, \mathcal{V})$ as described above whose 'relevant' edges all induce left-ℓ-robust separations (see Theorem 5.6.2), and whose nodes are thus eligible for the application of Lemma 5.5.4.

A set $\mathcal{F}$ of stars in $\vec{S}_{\aleph_0}(G)$ is *m-bounded* for some $m \in \mathbb{N}$ if $|\operatorname{int}(\varrho)| \leqslant m$ for all $\varrho \in \mathcal{F}$. It is *finitely bounded* if it is m-bounded for some $m \in \mathbb{N}$. The main result of this section then reads as follows:

Lemma 5.5.4. *Let G be a graph, $k, m \in \mathbb{N}$, and let $\mathcal{F}$ be an m-bounded, nice set of stars in $\vec{S}_k(G)$. Set $\ell := \max\{3k-2, k(k-1)m+m\}$, and let $\sigma := \{\vec{s}_1, \dots, \vec{s}_n\} \subseteq \vec{S}_k(G)$ be a finite star with finite interior. Suppose that every separation in σ is either left-ℓ-robust or has an inverse that is closely related to some k-profile in G that avoids $\mathcal{F}$. Set $\mathcal{F}' := \mathcal{F} \cup \{\{\overleftarrow{s_i}\} : i \in [n]\}$. Then either there is an $\mathcal{F}'$-tangle of $S_k(G)$ or there is a finite $S_k(G)$-tree over $\mathcal{F}'$ in which each $\vec{s}_i$ appears as a leaf separation.*

We remark that Erde [66] gave an example which shows that there need not exist an $S_k(G)$-tree over $\mathcal{F}'$ for every inessential star, even if G is finite and $\mathcal{F} = \mathcal{T}_k^*$ for some $k \in \mathbb{N}$. Thus, the additional assumptions on the separations in σ cannot be omitted. Moreover, Example 5.6.4 shows that we cannot omit the assumption that $\mathcal{F}$ is finitely bounded.

The remainder of this section is devoted to the proof of Lemma 5.5.4, which we briefly sketch here. The idea is to derive Lemma 5.5.4 from Theorem 2.3.1. In order to apply Theorem 2.3.1, we first reduce the problem to some finite separation system. For this, we define a subsystem $S_k^\sigma(G) \subseteq S_k(G)$ that consists only of those separations of G that are 'relevant' for finding either an $\mathcal{F}'$-tangle of $S_k(G)$ or an $S_k(G)$-tree over $\mathcal{F}'$. For a star $\sigma \subseteq \vec{S}_k(G)$, set

$$S_k^\sigma(G) := \{r \in S_k(G) : \vec{s} \leqslant \vec{r} \text{ or } \vec{s} \leqslant \overleftarrow{r} \text{ for every } \vec{s} \in \sigma\}.$$

As we will see in a moment, $\vec{S}_k^\sigma(G)$ is finite and $\mathcal{F}'$-separable if σ is finite and has finite interior. We can thus apply Theorem 2.3.1 to $S_k^\sigma(G)$ and $\mathcal{F}'$, which yields either an $\mathcal{F}'$-tangle of $S_k^\sigma(G)$ or an $S_k^\sigma(G)$-tree over $\mathcal{F}'$. By definition, an $S_k^\sigma(G)$-tree over $\mathcal{F}'$ is already an $S_k(G)$-tree over $\mathcal{F}'$. The main part of the proof is then concerned with showing that every $\mathcal{F}'$-tangle τ of $S_k^\sigma(G)$ *extends* to an $\mathcal{F}'$-tangle of $S_k(G)$: that there exists an $\mathcal{F}'$-tangle τ' of $S_k(G)$ such that $\tau \subseteq \tau'$.

Proposition 5.5.5. *Given a graph G and $k \in \mathbb{N}$, for every finite star $\sigma \subseteq \vec{S}_k(G)$ with finite interior the set $S_k^\sigma(G)$ is finite.*

Proof. By definition, every separation $\{A, B\} \in S_k^\sigma(G)$ is nested with σ, and thus for every $(C, D) \in \sigma$ we have $C \cap D \subseteq A$ or $C \cap D \subseteq B$. Hence, every separation $\{A, B\} \in S_k^\sigma(G)$ induces a separation $\{A \cap \operatorname{int}(\sigma), B \cap \operatorname{int}(\sigma)\}$ of $\operatorname{torso}(\sigma)$. As $\operatorname{int}(\sigma)$ is finite, there are only finitely many separations of $\operatorname{torso}(\sigma)$. It thus suffices to show that only finitely many separations in $S_k^\sigma(G)$ induce the same separation of $\operatorname{torso}(G)$.

For this, let $\{A', B'\}$ be a separation of $\operatorname{torso}(\sigma)$. Then every separation of G that induces $\{A', B'\}$ can be obtained from $\{A', B'\}$ by adding each component K of $G - \operatorname{int}(\sigma)$ to one side of

$\{A', B'\}$ that contains $N_G(K)$. It is straightforward to check that the arising separation will be in $S_k^\sigma(G)$ if and only if we added for every separation $(C, D) \in \sigma$ all components of $G - \operatorname{int}(\sigma)$ that are contained in $G[C \setminus D]$ to the same side of $\{A', B'\}$. It follows that at most $2^{|\sigma|}$ separations in $S_k^\sigma(G)$ induce the same separation of torso(G). As σ is finite, this concludes the proof. $\square$

Given an $\mathcal{F}'$-tangle τ of $S_k^\sigma(G)$, we now inductively construct an $\mathcal{F}'$-tangle τ' of $S_k(G)$ with $\tau \subseteq \tau'$. For this, recall that $\sigma := \{\vec{s}_1, \dots, \vec{s}_n\}$ is finite by assumption. So to define τ', we may proceed by extending $\tau =: \tau_n$ step-by-step to an $\mathcal{F}'$-tangle τ_i of $S_k^{\sigma_i}(G)$ where $\sigma_i := \{\vec{s}_1, \dots, \vec{s}_i\}$ for $i \in [n]$ and $\sigma_0 := \emptyset$. Then τ_0 will be the desired $\mathcal{F}'$-tangle of $S_k(G) = S_k^{\sigma_0}(G)$.

The following two lemmas describe how to obtain the tangle τ_{i-1} from τ_i. We distinguish between two cases: whether $\overleftarrow{s}_i$ is closely related to some k-profile that avoids $\mathcal{F}$ or whether $\vec{s}_i$ is left-ℓ-robust.

Lemma 5.5.6. *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a set of stars in $\vec{S}_k(G)$ that is strongly closed under shifting. Further, let $\sigma \subseteq \vec{S}_k(G)$ be a star, and suppose there is some $\vec{s} \in \sigma$ such that $\overleftarrow{s}$ is closely related to some k-profile in G that avoids $\mathcal{F}$. Set $\mathcal{F}' := \mathcal{F} \cup \{\{\overleftarrow{r}\} : \vec{r} \in \sigma\}$ and $\sigma' := \sigma \setminus \{\vec{s}\}$. Then the following assertions hold:*

(i) If τ' is an $\mathcal{F}'$-tangle of $S_k^{\sigma'}(G)$, then $\tau' \cap \vec{S}_k^\sigma(G)$ is an $\mathcal{F}'$-tangle of $S_k^\sigma(G)$.

(ii) If τ is an $\mathcal{F}'$-tangle of $S_k^\sigma(G)$, then τ extends to an $\mathcal{F}'$-tangle of $S_k^{\sigma'}(G)$.

Proof. By definition it is clear that every $\mathcal{F}'$-tangle τ of $S_k^{\sigma'}(G)$ induces an $\mathcal{F}'$-tangle $\tau \cap \vec{S}_k^\sigma(G)$ of $S_k^\sigma(G)$, so (i) holds.

For (ii), let τ be an $\mathcal{F}'$-tangle of $S_k^\sigma(G)$. We extend τ to an orientation τ' of $S_k^{\sigma'}(G)$ as follows. Let $r \in S_k^{\sigma'}(G)$ be given, and first assume that r is nested with s. Then either $r \in S_k^\sigma(G)$, and we then let $\vec{r} \in \tau'$ if and only if $\vec{r} \in \tau$, or r has an orientation that is smaller than $\vec{s}$, and we then let $\vec{r} \in \tau'$ if and only if $\vec{r} \leqslant \vec{s}$. Second, assume that r and s cross, and fix an orientation $\vec{r}$ of r. By the assumption on $\vec{s}$, its inverse $\overleftarrow{s}$ is closely related to some k-profile P in G that avoids $\mathcal{F}$. Since $r \in S_k(G)$, it is oriented by P; we set $\vec{t} := \vec{r} \vee \vec{s}$ if $\overleftarrow{r} \in P$ and $\vec{t} := \vec{r} \wedge \overleftarrow{s}$ if $\vec{r} \in P$. As $\overleftarrow{s}$ is closely related to P, it follows that t has order $< k$, and thus $t \in S_k^\sigma(G)$ by Lemma 2.1.1. Hence, τ contains an orientation of t; we let $\vec{r} \in \tau'$ if $\vec{t} \in \tau$, and $\overleftarrow{r} \in \tau'$ otherwise.

By definition, τ' is an orientation of $S_k^{\sigma'}(G)$ and $\tau \subseteq \tau'$; in particular, τ extends to τ' and $\sigma \subseteq \tau$. Hence, we are left to show that τ' is an $\mathcal{F}$-tangle of $S_k^{\sigma'}(G)$. For this, suppose for a contradiction that there is a set $\varrho \subseteq \tau'$ which has one of the following two forms: either $\varrho = \{\vec{r}_i : i \in I\}$ is a star in $\mathcal{F}$, or $\varrho = \{\vec{r}_1, \vec{r}_2\}$ with $\overleftarrow{r}_1 < \vec{r}_2$.

As P is consistent and avoids $\mathcal{F}$, we have $\varrho \not\subseteq P$, and thus P contains the inverse $\overleftarrow{r}_i$ of a separation $\vec{r}_i \in \varrho$, say $\overleftarrow{r}_1 \in P$. If ϱ is a star in $\mathcal{F}$, then $\vec{r}_i \in P$ for all $\vec{r}_i \in \varrho \setminus \{\vec{r}_1\}$ since P is consistent. As $\overleftarrow{s}$ is closely related to the k-profile P, it follows from Proposition 5.2.2 that

$\varrho' := \{\vec{r}_1 \vee \vec{s}\} \cup \{(\vec{r}_i \wedge \overleftarrow{s}) : \vec{r}_i \in \varrho \smallsetminus \{\vec{r}_1\}\}$ is a star in $\mathcal{F}$. But then $\varrho' \subseteq \tau$ by the definition of τ', which contradicts that τ is an $\mathcal{F}'$-tangle of $S_k^\sigma(G)$.

Otherwise, if $\varrho = \{\vec{r}_1, \vec{r}_2\}$ with $\overleftarrow{r}_1 < \vec{r}_2$, then, since $\overleftarrow{s}$ is closely related to P, we have $\vec{r}_1 \vee \vec{s}, \vec{r}_2 \vee \vec{s} \in \vec{S}_k(G)$ if $\overleftarrow{r}_2 \in P$, or $\vec{r}_1 \vee \vec{s}, \vec{r}_2 \wedge \overleftarrow{s} \in \vec{S}_k(G)$ if $\vec{r}_2 \in P$. By the definition of τ', it follows that $\{\vec{r}_1 \vee \vec{s}, \vec{r}_2 \vee \vec{s}\} \subseteq \tau$ or $\{\vec{r}_1 \vee \vec{s}, \vec{r}_2 \wedge \overleftarrow{s}\} \subseteq \tau$, respectively, which contradicts that τ is consistent because $(\vec{r}_1 \vee \vec{s})^* \leqslant \overleftarrow{r}_1 < \vec{r}_2 \leqslant \vec{r}_2 \vee \vec{s}$ and $(\vec{r}_1 \vee \vec{s})^* = \overleftarrow{r}_1 \wedge \overleftarrow{s} \leqslant \vec{r}_2 \wedge \overleftarrow{s}$. □

Lemma 5.5.7. *Let G be any graph, $k, m \in \mathbb{N}$, and let $\mathcal{F}$ be an m-bounded, nice set of stars in $\vec{S}_k(G)$. Set $\ell := \max\{3k-2, k(k-1)m + m\}$, let $\sigma \subseteq \vec{S}_k(G)$ be a star, and suppose that some $(C, D) \in \sigma$ is left-ℓ-robust. Set $\mathcal{F}' := \mathcal{F} \cup \{\{\overleftarrow{r}\} : \vec{r} \in \sigma\}$ and $\sigma' := \sigma \smallsetminus \{(C, D)\}$. Then the following assertions hold:*

(i) If τ' is an $\mathcal{F}'$-tangle of $S_k^{\sigma'}(G)$, then $\tau' \cap \vec{S}_k^\sigma(G)$ is an $\mathcal{F}'$-tangle of $S_k^\sigma(G)$.

(ii) If τ is an $\mathcal{F}'$-tangle of $S_k^\sigma(G)$, then τ extends to an $\mathcal{F}'$-tangle of $S_k^{\sigma'}(G)$.

Proof. By definition it is clear that every $\mathcal{F}'$-tangle τ' of $S_k^{\sigma'}(G)$ induces an $\mathcal{F}'$-tangle $\tau' \cap \vec{S}_k^\sigma(G)$ of $S_k^\sigma(G)$, so (i) holds.

For (ii), let τ be an $\mathcal{F}'$-tangle of $S_k^\sigma(G)$. Fix a set $U \subseteq C$ of size ℓ and a family $\{P_x : x \in C \cap D\}$ of disjoint paths witnessing that (C, D) is left-ℓ-robust. We first observe the following:

Claim 1. Every $\{A, B\} \in S_k(G)$ has a side that meets every P_x at least once.

Proof. Since $\ell \geqslant 3k-2$ and $|A \cap B| < k$, some strict side of $\{A, B\}$, say $A \smallsetminus B$, contains a subset U' of U of at least k vertices. By the choice of U and the P_x, there are at least k pairwise internally disjoint U'–P_x paths for every $x \in C \cap D$. Since $U' \subseteq A \smallsetminus B$ and $|A \cap B| < k$, at least one of these paths is contained in A. Thus, A meets every P_x at least once. ■

Claim 2. For every $\{A, B\} \in S_k(G)$, if A meets every P_x, then $\{A \cup C, B \cap D\}$ has order $< k$.

Proof. By assumption, there are vertices $p_x \in A \cap V(P_x)$ for every $x \in C \cap D$. Since also $V(P_x) \subseteq C$ for every $x \in C \cap D$, we have $p_x \in A \cap C$, and thus every P_x meets $A \cap C$. Additionally, as every P_x ends in $x \in C \cap D$, it also meets $B \cup D$. Since P_x is connected and $\{A \cap C, B \cup D\}$ is a separation of G, it follows that $V(P_x) \cap ((A \cap C) \cap (B \cup D)) \neq \emptyset$ for all $x \in C \cap D$. Thus $\{A \cap C, B \cup D\}$ has order $\geqslant |C \cap D|$, as the P_x are disjoint. By submodularity, this implies that $\{A \cup C, B \cap D\}$ has order $\leqslant |A \cap B| < k$. ■

Claim 3. For every $\{A, B\} \in S_k^{\sigma'}(G)$, if A meets every P_x, then τ contains an orientation of $\{A \cup C, B \cap D\}$.

Proof. By Claim 2 and Lemma 2.1.1, $\{A \cup C, B \cap D\}$ is contained in $S_k^\sigma(G)$, which clearly implies the assertion. ■

Claim 4. If both sides A and B of some $\{A,B\} \in S_k^{\sigma'}(G)$ meet every P_x at least once, then either $(A \cup C, B \cap D)$, $(A \cap D, B \cup C) \in \tau$ or $(B \cup C, A \cap D), (B \cap D, A \cup C) \in \tau$.

Proof. By Claim 3, τ orients both separations, $\{A \cup C, B \cap D\}$ and $\{B \cup C, A \cap D\}$. Since τ is consistent, it cannot contain both $(A \cup C, B \cap D)$ and $(B \cup C, A \cap D)$. We show that $D \cap (A \cup C) \cap (B \cup C)$ has size less than k. The assertion then follows as $\mathcal{P}_k' \subseteq \mathcal{F}$ because $\mathcal{F}$ is nice, and $(C,D) \in \tau$ as well as $\{(C,D), (B \cap D, A \cup C), (A \cap D, B \cup C)\} \in \mathcal{P}_k$. We have

$$D \cap (A \cup C) \cap (B \cup C) = D \cap ((A \cap B) \cup C) = (A \cap B \cap D) \cup (C \cap D) = (C \cap D) \dot{\cup} (A \cap B \cap (D \setminus C)),$$

and thus

$$|D \cap (B \cup C) \cap (A \cup C)| = |C \cap D| + (|A \cap B| - |A \cap B \cap C|) < |C \cap D| + k - |A \cap B \cap C|.$$

Since both A and B meet every P_x, and because every P_x is connected, also $A \cap B$ meets every P_x. As all P_x are pairwise disjoint and contained in $G[C]$, it follows that $|A \cap B \cap C| \geqslant |C \cap D|$. Hence,

$$|D \cap (B \cup C) \cap (A \cup C)| < |C \cap D| + k - |C \cap D| = k.$$

This completes the proof of the claim. ■

We now define an orientation τ' of $S_k^{\sigma'}(G)$ as follows:

$$\text{For every } \{A,B\} \in S_k^{\sigma'}(G), \text{ if } A \text{ meets every path } P_x \text{ at least once, then we let } (A,B) \in \tau' \text{ if } (A \cup C, B \cap D) \in \tau, \text{ and } (B,A) \in \tau' \text{ if } (B \cap D, A \cup C) \in \tau. \tag{$\triangle$}$$

By Claims 1, 3 and 4, τ' contains precisely one orientation of every separation in $S_k^{\sigma'}(G)$. We claim that τ' is an $\mathcal{F}'$-tangle of $S_k^{\sigma'}(G)$ and that τ extends to τ', i.e. $\tau \subseteq \tau'$.

We first show the latter. For this, let $(A,B) \in \tau$ be given. Then $\{A,B\} \in S_k^{\sigma}(G)$, which implies that either $(C,D) \leqslant (A,B)$ or $(C,D) \leqslant (B,A)$. In the first case, we have $V(P_x) \subseteq C \subseteq A$ for every $x \in C \cap D$. Moreover, $(A \cup C, B \cap D) = (A,B) \in \tau$, and thus $(A,B) \in \tau'$ by ($\triangle$). Analogously, we find in the second case that $(B,A) \in \tau'$.

Towards a proof that τ' is an $\mathcal{F}'$-tangle of $S_k^{\sigma'}(G)$, we first show that τ' is consistent. For this, suppose for a contradiction that there are $(A,B), (E,F) \in \tau'$ such that $(B,A) < (E,F)$. By Claim 1, $\{A,B\}$ has a side that meets every P_x at least once; we first assume that $B \cap V(P_x) \neq \emptyset$ for all $x \in C \cap D$. Then also $E \supseteq B$ meets every P_x, and thus, by ($\triangle$), we have $(A \cap D, B \cup C), (E \cup C, F \cap D) \in \tau$, which contradicts that τ is consistent as $(B \cup C, A \cap D) \leqslant (E \cup C, F \cap D)$. The case that F meets every P_x is symmetric, so we may assume that $A \cap V(P_x) \neq \emptyset \neq E \cap V(P_x)$ for every $x \in C \cap D$. Then $(A \cup C, B \cap D), (E \cup C, F \cap D) \in \tau$ by ($\triangle$), which again contradicts that τ is consistent since

$(B \cap D, A \cup C) \leqslant (B, A) < (E, F) \leqslant (E \cup C, F \cap D)$.

It remains to show that τ' avoids $\mathcal{F}'$. We have already seen that $\tau \subseteq \tau'$. Since τ avoids $\mathcal{F}'$, this implies that $\sigma \subseteq \tau \subseteq \tau'$. So suppose for a contradiction that there is a star $\varrho \subseteq \tau'$ such that $\varrho \in \mathcal{F}$.

Claim 5. There is a separation $(A, B) \in \varrho$ such that A meets every path P_x at least once.

Proof. Let us first assume that there is a separation $(A, B) \in \varrho$ whose strict small side $A \setminus B$ contains a set U' of k vertices from U. By the choice of U and the P_x, it follows that there are k internally disjoint U'–P_x paths, for every $x \in C \cap D$. Since $U' \subseteq A \setminus B$ and $|A \cap B| < k$, at least one of these paths is contained in A. Thus, A meets every P_x at least once.

Now suppose that no separation in ϱ contains more than $k-1$ vertices from U in its strict small side. Since $|\operatorname{int}(\varrho)| \leqslant m$, it follows that $U \cap (A \setminus B) \neq \emptyset$ for at least $(\ell - m)/(k-1) = km$ separations $(A, B) \in \varrho$. Let $\varrho' \subseteq \varrho$ be the set of these separations, and pick for each $(A, B) \in \varrho'$ some $v_A \in U \cap (A \setminus B)$. Further, fix some $y \in C \cap D$. By the choice of U, there is a family $\{Q_A : (A, B) \in \varrho'\}$ of internally disjoint paths such that Q_A starts in v_A and ends in a vertex of P_y. As each Q_A meets $A \cap B \subseteq \operatorname{int}(\varrho)$ in an internal vertex if $V(P_y) \cap A = \emptyset$, and since $|\operatorname{int}(\varrho)| \leqslant m$, P_y meets the small sides A of at least $(k-1)m$ separations $(A, B) \in \varrho'$. As $|C \cap D| < k$, iterating this argument for all $x \in C \cap D$ yields that the small side of some $(A, B) \in \varrho' \subseteq \varrho$ meets every P_x. ■

By Claim 5 there is a separation $(A, B) \in \varrho$ whose small side A meets every P_x at least once in some vertex p_x. Since ϱ is a star, we then have $p_x \in A \subseteq F$ for every $(E, F) \in \varrho \setminus \{(A, B)\}$. Therefore, by $(\triangle)$,

$$\varrho' := \{(A \cup C, B \cap D)\} \cup \{(E \cap D, F \cup C) : (E, F) \in \varrho \setminus \{(A, B)\}\}$$

is contained in τ.

We claim that $\varrho' \in \mathcal{F}$, contradicting that τ is an $\mathcal{F}$-tangle of $S_k^{\sigma}(G)$. For this, we show that (C, D) emulates $(A \cap C, B \cup D)$ in $S_k(G)$, from which the claim follows as $\mathcal{F}$ is strongly closed under shifting and $|(A \cap C) \cap (B \cup D)| \leqslant |A \cap B| + |C \cap D| \leqslant 2k - 2$. Indeed, let $\{E, F\} \in S_k(G)$ such that $(A \cap C, B \cup D) \leqslant (E, F)$. Then $p_x \in A \cap C \subseteq E$ for every $x \in C \cap D$, so E meets every P_x. Then by Claim 2, $\{E \cup C, F \cap D\}$ has order $< k$, which concludes the proof. □

For the proof of Lemma 5.5.4 it remains to show that $\vec{S}_k^{\sigma}(G)$ is $\mathcal{F}'$-separable.

Lemma 5.5.8. *Let G be a graph, $k \in \mathbb{N}$, and let $\sigma \subseteq \vec{S}_k(G)$ be a finite star. Suppose that every separation is either left-$(3k-2)$-robust or has an inverse that is closely related to some k-profile in G. If $\mathcal{F}$ is a set of stars in $\vec{S}_k(G)$ and closed under shifting, then $\vec{S}_k^{\sigma}(G)$ is $\mathcal{F}'$-separable where $\mathcal{F}' := \mathcal{F} \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma\}$.*

Proof. Since every $\overleftarrow{s}$ with $\vec{s} \in \sigma$ is maximal in $\vec{S}^{\sigma}_k(G)$ with respect to the partial order on $\vec{S}_k(G)$, we have $\{\overleftarrow{s} \vee \vec{t}\} = \{\overleftarrow{s}\} \in \mathcal{F}'$ for all $\vec{t}$ that emulate some $\vec{r} \leqslant \overleftarrow{s}$ in $\vec{S}^{\sigma}_k(G)$. Hence, it suffices to show that $\vec{S}^{\sigma}_k(G)$ is $\mathcal{F}$-separable. So let $\vec{r}, \vec{r'} \in \vec{S}^{\sigma}_k(G)$ be given such that $\vec{r} < \overleftarrow{r'}$, and pick a separation $s \in \vec{S}^{\sigma}_k(G)$ of minimal order such that $\vec{r} \leqslant \vec{s} \leqslant \overleftarrow{r'}$. We claim that $\vec{s}$ and $\overleftarrow{s}$ emulate $\vec{r}$ and $\vec{r'}$, respectively, in $\vec{S}^{\sigma}_k(G)$ for $\mathcal{F}$. This clearly implies the assertion.

We show that $\vec{s}$ emulates $\vec{r}$ in $\vec{S}^{\sigma}_k(G)$ for $\mathcal{F}$; the other case is symmetric. By the choice of s and Lemma 2.1.1, $\vec{x} \wedge \vec{s}$ has order at least $|s|$ for all $x \in S^{\sigma}_k(G)$ with $\vec{x} \geqslant \vec{r}$. So $\vec{s} \vee \vec{x} \in \vec{S}^{\sigma}_k$ by submodularity and Lemma 2.1.1, which implies that $\vec{s}$ emulates $\vec{r}$ in $\vec{S}^{\sigma}_k(G)$. Since $\mathcal{F}$ is closed under shifting in $\vec{S}_k(G)$, it thus suffices to show that $\vec{s}$ also emulates $\vec{r}$ in $\vec{S}_k(G)$. So suppose for a contradiction that there is some $\vec{x} \in \vec{S}_k(G)$ with $\vec{x} \geqslant \vec{r}$ such that $\vec{x} \vee \vec{s}$ has order $\geqslant k$; since σ is finite, we may choose $\vec{x}$ so that x and y are nested for as many $\vec{y} \in \sigma$ as possible. Let $\sigma' \subseteq \sigma$ consist of all those $\vec{y} \in \sigma$ such that y, x cross. As $\vec{s} \geqslant \vec{r}$ and $\vec{r} \in \vec{S}^{\sigma}_k(G)$, we have $\vec{r} \leqslant \overleftarrow{y}$ for all $\vec{y} \in \sigma'$. If also $\vec{s} \leqslant \overleftarrow{y}$ for all $\vec{y} \in \sigma'$, then $\vec{s} \wedge \vec{x} \in \vec{S}^{\sigma}_{\aleph_0}(G)$. Since $\vec{r} \leqslant \vec{s} \wedge \vec{x} \leqslant \vec{s} \leqslant \overleftarrow{r'}$, it follows by the choice of s that $|\vec{s} \wedge \vec{x}| \geqslant |s|$. By submodularity, this implies $|\vec{s} \vee \vec{x}| \leqslant |x| < k$ as desired.

Hence, we may assume that $\vec{y} \leqslant \vec{s}$ for some $\vec{y} \in \sigma'$. Then by Claims 1 and 2 in the proof of Lemma 5.5.4 if $\vec{y}$ is left-$(3k-2)$-robust or by definition if $\overleftarrow{y}$ is closely related to a k-profile in G, one of $\vec{x} \vee \vec{y}$ and $\vec{x} \wedge \overleftarrow{y}$ has order $\leqslant |x|$; let $\vec{t}$ be that corner. By Lemma 2.1.1, t is nested with more separations in σ than x, so we have that $|\vec{s} \vee \vec{t}| < k$ by the choice of x. Since $|\vec{s} \vee \vec{t}| \geqslant |\vec{s} \vee \vec{x}|$, which it is straightforward to check, this concludes the proof. □

We are now ready to prove Lemma 5.5.4:

Proof of Lemma 5.5.4. Since σ is finite and has finite interior, $S^{\sigma}_k(G)$ is finite by Proposition 5.5.5. Moreover, by Lemma 5.5.8, $\vec{S}^{\sigma}_k$ is $\mathcal{F}'$-separable. Hence, we can apply Theorem 2.3.1 to $S^{\sigma}_k(G)$ and $\mathcal{F}'$, which yields either an $\mathcal{F}'$-tangle of $S^{\sigma}_k(G)$ or an $S^{\sigma}_k(G)$-tree over $\mathcal{F}'$. Let us first assume that there is an $\mathcal{F}'$-tangle τ' of $S^{\sigma}_k(G)$. Then inductively applying Lemma 5.5.6 (ii) and 5.5.7 (ii) yields that τ' extends to an $\mathcal{F}'$-tangle τ of $S_k(G)$ as desired. So we may assume that there is an $S^{\sigma}_k(G)$-tree (T, α) over $\mathcal{F}'$. By definition, (T, α) is also an $S_k(G)$-tree over $\mathcal{F}'$, so we are left to show that every $\vec{s} \in \sigma$ appears as a leaf separation of (T, α). For this, let $\vec{s} \in \sigma$ be given, and set $O := \{\vec{r} \in \vec{S}^{\sigma}_k(G) : \vec{r} \leqslant \overleftarrow{s}\}$. As $\vec{s} \in \sigma$, the set O is a consistent orientation of $S^{\sigma}_k(G)$. Moreover, since either $\vec{s}$ is left-ℓ-robust and $\mathcal{F}$ is $(\ell - 1)$-bounded, or $\overleftarrow{s}$ is contained in an $\mathcal{F}$-tangle, O avoids $\mathcal{F}$. Thus, O has to live at a leaf of T, which then has to be associated with $\{\overleftarrow{s}\}$; so $\vec{s}$ is a leaf separation of (T, α). □

5.6 Tangle-tree duality in infinite graphs

In this section we prove Theorems 11 to 13. In fact, we prove the following more general duality theorem for $\mathcal{F}$-tangles.

Theorem 16. *Let G be a graph, and let $k \in \mathbb{N}$. Further, let $\mathcal{F}$ be a finitely bounded, nice set of stars in $\vec{S}_k(G)$. Then exactly one of the following assertions holds:*

(i) *There exists a principal $\mathcal{F}$-tangle of $S_k(G)$ which is not induced by an end of combined degree $< k$.*

(ii) *There exists a weakly exhaustive $S_k(G)$-tree over $\mathcal{F} \cup \mathcal{U}_k^\infty$.*

For the proof of Theorem 13, we first show two auxiliary statements. We will use the first of them to prove that in Theorem 16 not both, (i) and (ii), can hold.

Lemma 5.6.1. *Let G be a graph, $k \in \mathbb{N}$, and let $((A_i, B_i))_{i \in \mathbb{N}}$ be a weakly exhaustive increasing sequence in $\vec{S}_k(G)$. Then every principal, consistent orientation τ of $S_k(G)$ with $(A_i, B_i) \in \tau$, for every $i \in \mathbb{N}$, is induced by an end of G of combined degree $< k$.*

Proof. Let $\mathcal{S}$ be the set of all finite subsets of $V(G)$. A *direction* of G is a map f with domain $\mathcal{S}$ that maps every $S \in \mathcal{S}$ to a component of $G - S$ such that $f(S) \subseteq f(S')$ whenever $S' \subseteq S$. Now τ defines a direction f of G as follows. Let $S \in \mathcal{S}$ be given. Since $((A_i, B_i))_{i \in \mathbb{N}}$ is weakly exhaustive, there exists some $i \in \mathbb{N}$ such that $S \subseteq A_i$. As τ is principal, there further exists a component C_i of $G - (A_i \cap B_i)$ such that $(V(G - C_i), V(C_i) \cup (A_i \cap B_i)) \in \tau$; and since τ is consistent, we have $C_i \subseteq G[B_i \smallsetminus A_i]$. Now let $f(S) := C$ where C is the unique component of $G - S$ that contains C_i. It is straightforward to check that f is a direction because τ is consistent.

By [50, Theorem 2.2], there exists an end ε of G such that for all $\{A, B\} \in S_{\aleph_0}(G)$ we have $(A, B) \in \tau_\varepsilon$ if and only if $f(A \cap B) \subseteq G[B \smallsetminus A]$. We claim that τ is induced by ε. Indeed, let $(A, B) \in \tau$ be given. Since $A \cap B$ is finite and $((A_i, B_i))_{i \in \mathbb{N}}$ is weakly exhaustive, there exists some $i \in \mathbb{N}$ such that $A \cap B \subseteq A_i$. As C_i is connected and avoids $A \cap B \subseteq A_i$, either $C_i \subseteq G[B \smallsetminus A]$ or $C_i \subseteq G[A \smallsetminus B]$. In fact, because τ is consistent, we find $C_i \subseteq G[B \smallsetminus A]$, and hence $(A, B) \leqslant (V(G - C_i), V(C_i) \cup (A_i \cap B_i))$. It follows that $(A, B) \in \tau_\varepsilon$ since τ_ε is consistent and $(V(G - C_i), V(C_i) \cup (A_i \cap B_i)) \in \tau_\varepsilon$ by the choice of ε.

The assertion now follows since $\Delta(\varepsilon) < k$ by Lemma 5.3.2 (i). □

In order to show that in Theorem 16 at least one of (i) and (ii) holds, we will construct for every graph with no tangles as in (i) an $S_k(G)$-tree as in (ii). For this, we need the following theorem which follows easily from Theorem 4' in Chapter 3:

Theorem 5.6.2. *Let G be a graph, $\ell \in \mathbb{N}$, and suppose there exists some $k \in \mathbb{N}$ such that every end of G has combined degree $< k$ and every critical vertex set of G has size $< k$. Then G has a tree-decomposition $(T, \mathcal{V})$ of adhesion $< k$ into finite parts such that*

(i) *for every node t of T of infinite degree its associated star σ_t is in $\mathcal{U}_k^\infty$, and infinitely many $(A, B) \in \sigma_t$ are left-tight and satisfy $A \cap B = V_t$,*

(ii) *every end η of T is home to a unique end ε of G, and we have $\liminf_{i\in\mathbb{N}} |V_{r_i} \cap V_{r_{i+1}}| = \Delta(\varepsilon)$ for every η-ray $R = r_1r_2\ldots$ in T,*

(iii) *for every edge $\vec{e} = (t, s)$ in T, if $\deg_T(s) < \infty$, then the separation induced by $\vec{e}$ is left-ℓ-robust, and*

(iv) *$(T, \mathcal{V})$ is tight and displays the infinities.*[7]

Proof. Since all ends of G have finite degree, G contains no half-grid minor, and in particular no subdivision of $K_{\aleph_0}$. So by [84, Theorem 10.1] and [12, Theorem 2.2], G has finite tree-width, that is, admits a tree-decomposition into finite parts. Thus, by Theorem 4' and Lemma 3.8.3, G admits a tree-decomposition $(T, \mathcal{V})$ into finite parts which satisfies (i) – (iv). Indeed, (iii) and (iv) follow immediately by Lemma 3.8.3. (i) and (ii) hold because all bags of $(T, \mathcal{V})$ are finite and $(T, \mathcal{V})$ *displays the infinities of G*.[8] Moreover, by Theorem 4' (I1), all adhesion sets of $(T, \mathcal{V})$ are either *linked*[8] to an end or a critical vertex set of G. Since all ends of G have combined degree $< k$ and all critical vertex sets of G have size $< k$, this implies that all adhesion sets of $(T, \mathcal{V})$ have size $< k$. □

In the proof of Theorem 16, if a graph has no tangles as in (i), we will apply Lemma 5.5.4 to the stars associated with the nodes of the tree-decomposition $(T, \mathcal{V})$ from Theorem 5.6.2. We then obtain an $S_k(G)$-tree as in (ii) by sticking the $S_k(G)$-trees obtained from Lemma 5.5.4 together along T as follows:

Construction 5.6.3. Let $(T, \mathcal{V})$ be a tree-decomposition of a graph G of adhesion $< k \in \mathbb{N}$, and let $\mathcal{F}$ be a set of stars in $\vec{S}_k(G)$. Further, let $U \subseteq V(T')$ be a set of nodes of T. Assume that for every node $t \in U$, we are given a weakly exhaustive $S_k(G)$-tree (T^t, α^t) over $\mathcal{F} \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma_t\}$ in which each $\vec{s} \in \sigma_t$ appears as a leaf separation. Assume further that the stars σ_t associated in T with nodes $t \in V(T) \setminus U$ are contained in $\mathcal{F}$.

Set $T^t := T[N_T(t) \cup \{t\}]$ and $\alpha^t((s, t)) := (U_s, U_t)$ for all $t \in V(T) \setminus U$. We obtain the tree T' from the disjoint union of the trees T^t by identifying for every edge $e = \{t_1, t_2\}$ of T the nodes $s_1 \in T^{t_1}$ and $u_2 \in T^{t_2}$ as well as $s_2 \in T^{t_2}$ and $u_1 \in T^{t_1}$ where (u_i, s_i) is the unique (leaf) edge of T^{t_i} such that $\alpha^{t_i}(u_i, s_i) = (U_{t_i}, U_{t_{3-i}})$. For each edge e of T', we then set $\alpha'(\vec{e})$ to be $\alpha^t(\vec{e})$,

[7] See Section 3.2.3 for definitions. We remark that it is not important what these properties actually mean, we only need them once in Section 5.8 to apply Lemma 3.8.6 to this tree-decomposition.

[8] See Section 3.2.3 for a definition.

where t is a node of T such that $e \in E(T^t)$. It is straightforward to check that (T', α') is an $S_k(G)$-tree over $\mathcal{F}$. Moreover, since $(T, \mathcal{V})$ is a tree-decomposition of G and because each (T^t, α^t) is weakly exhaustive, (T', α') is weakly exhaustive.

We are now ready to prove Theorem 16.

Proof of Theorem 16. We first show that not both, (i) and (ii), can hold for G. For this, suppose that there is a weakly exhaustive $S_k(G)$-tree (T, α) as in (ii), and let τ be any consistent orientation of $S_k(G)$. We claim that τ is not a tangle as in (i). Indeed, since τ is consistent, it induces via α^{-1} a consistent orientation O of $E(T)$. It follows that O either contains a sink or a directed ray. If O contains a sink, that is, if there is a node t of T all whose incident edges are oriented inwards by O, then $\sigma_t \subseteq \tau$. But T is over $\mathcal{F} \cup \mathcal{U}_k^\infty$, and thus τ is either not principal or not an $\mathcal{F}$-tangle. Otherwise, if O contains a directed ray $R = r_1 r_2 ...$, then, since (T, α) is weakly exhaustive, τ contains an infinite weakly exhaustive increasing sequence $(\alpha(r_i, r_{i+1}))_{i \in \mathbb{N}}$ of separations of order $< k$. It follows, by Lemma 5.6.1, that τ is either non-principal or induced by an end of G of combined degree $< k$.

We now show that at least one of (i) and (ii) holds. For this, suppose that (i) does not hold: that all principal $\mathcal{F}$-tangles of $S_k(G)$ are induced by ends of combined degree $< k$. We show that then (ii) must hold: that there exists a weakly exhaustive $S_k(G)$-tree over $\mathcal{F} \cup \mathcal{U}_k^\infty$. By Proposition 5.3.1 and Lemma 5.4.3, all ends of G have combined degree $< k$, and all critical vertex sets of G have size $< k$. We may thus apply Theorem 5.6.2, which yields a tree-decomposition $(T, \mathcal{V})$ of G of adhesion $< k$ into finite parts. Let t be a node of T finite degree. Since $\operatorname{int}(\sigma_t) = V_t$ is finite, and because of Theorem 5.6.2 (iii), σ_t satisfies the premise of Lemma 5.5.4. Since σ_t is not home to any ends as $\operatorname{int}(\sigma_t) = V_t$ is finite, and since G does not contain any principal $\mathcal{F}$-tangles in G of order k that are not induced by an end, the star σ_t cannot be home to any principal $\mathcal{F}$-tangles of order k. Moreover, since $\mathcal{F}$ is nice, and hence $\mathcal{P}'_k \subseteq \mathcal{F}$ as well as $\{(V(G), A)\} \in \mathcal{F}$ for all sets A of fewer than k vertices, σ_t can also not be home to any non-principal $\mathcal{F}$-tangle by Lemma 5.2.1. Hence, applying Lemma 5.5.4 to σ_t yields a finite $S_k(G)$-tree (T^t, α^t) over $\mathcal{F}' := \mathcal{F} \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma_t\}$ in which each $\vec{s} \in \sigma_t$ appears as a leaf separation.

Since $\sigma_t \in \mathcal{U}_k^\infty$ for all infinite-degree nodes t of T by Theorem 5.6.2 (i), applying Construction 5.6.3 to $(T, \mathcal{V})$ and the (T^t, α^t) yields a weakly exhaustive $S_k(G)$-tree over $\mathcal{F} \cup \mathcal{U}_k^\infty$, as desired. □

Proof of Theorem 13. By Lemma 5.2.3, $\mathcal{T}_k^*$ is nice. As the $\mathcal{T}_k^*$-tangles of $S_k(G)$ are precisely the k-tangles in G if $|G| \geqslant k$ [54, Lemma 4.2], the assertion follows immediately by applying Theorem 16 to the $(3k-3)$-bounded, nice set $\mathcal{T}_k^*$. □

Proof of Theorem 11. By definition, $\mathcal{U}_k^\infty$ is empty if G is locally finite. Moreover, since every

k-tangle is a k-profile, inductively applying the profile property yields that every k-tangle in a locally finite graph is principal (see also [41, Exercise 43 in Ch. 12]). The assertion thus follows immediately from Theorem 13. □

Proof of Theorem 12. By Theorem 13, it is enough to show that if (ii) of Theorem 13 holds, then also (ii) of Theorem 12 holds. For this, assume that (T,α) is a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$. If (T,α) is even over $\mathcal{T}^*$, then we are done. Otherwise we define an $S_k(G)$-tree (T',α') as follows.

By pruning the tree T if necessary, we may assume that (T,α) is *irredundant*: for every node t of T and neighbours t',t'' of t we have $\alpha(t',t)=\alpha(t'',t)$ if and only if $t'=t''$.[9] Then all nodes in T have countable degree. Indeed, let t be a node of of T of infinite degree and consider $\sigma_t := \{\alpha(s,t) : \{s,t\} \in E(T)\}$. Then σ_t is a star in $\vec{S}_k(G)$ because (T,α) is an $S_k(G)$-tree over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$. Since G is countable, there are only countably many small separations of G of the form $(A, V(G))$ for some set A of fewer than k vertices, and also σ_t can contain at most countably many separations of the form (A,B) with $B \neq V(G)$, as any such separation contains a vertex in its strict small side $A \setminus B$ that is not contained in the strict small side of any other separation in σ_t. Hence, σ_t is countable, and thus $N_T(t)$ is countable since (T,α) is irredundant.

Let r be an arbitrary node of T, and for every infinite-degree node t of T, let $\{s_i^t : i \in \mathbb{N}_0\}$ be an enumeration of its neighbourhood such that s_0^t is the unique vertex of rTt that is incident with t. Let F be the forest obtained from T by deleting all edges e of the form $e=\{t,s_i^t\}$ where $\deg(t)=\infty$ and $i \geqslant 2$. Now the tree T' is obtained from F by simultaneously adding, for every infinite-degree node t of T, a ray $R_t := r_2^t r_3^t \dots$, the edge $\{t, r_2^t\}$, and all edges of the form $\{r_i^t, s_i^t\}$ for $i \geqslant 2$. Further, let $\alpha' : \vec{E}(T') \to \vec{S}_k(G)$ be defined via $\alpha'(\vec{e}) := \alpha(\vec{e})$ for all edges $e \in E(T') \cap E(T)$, and $\alpha'(s_i^t, r_i^t) := \alpha(s_i^t, t)$ and $\alpha'(r_i^t, r_{i+1}^t) := \bigvee_{j \geqslant i} \alpha(s_j^t, t)$ for all $i \geqslant 2$ as well as $\alpha'(t, r_2^t) := \alpha(s_0^t, t) \vee \alpha(s_1^t, t)$.

Then (T',α') is again an $S_k(G)$-tree, and it is straightforward to check that (T',α') is weakly exhaustive and over $\mathcal{T}^* \cup \{\sigma \in \mathcal{U}_k : |\sigma| = 3\}$. To turn (T',α') into an $S_k(G)$-tree over $\mathcal{T}^*$, we add a subdivision vertex v_e to those edges $e=\{s,t\}$ of T' whose endvertices s and t are both associated in T' with stars σ'_s, σ'_t in $\mathcal{U}_k \setminus \mathcal{T}^*$. We denote the arising tree with T''. To define α'', let e be an edge of T'', and first assume that $e=\{s,t\}$ for nodes s,t of T'. If $\sigma'_s, \sigma'_t \in \mathcal{T}^*$, then let $\alpha''(\vec{e}) := \alpha'(\vec{e})$. Otherwise, if $\sigma'_t \notin \mathcal{T}^*$, then let $\alpha''((s,t)) := \alpha'((s,t)) \wedge (V(G), \mathrm{int}(\sigma'_t))$ (and $\alpha''((t,s))$ accordingly). Second, if $e=\{v_f,t\}$ where $f=\{s,t\} \in E(T')$, then let $\alpha''((v_f,t)) := \alpha'((s,t)) \wedge (V(G), \mathrm{int}(\sigma'_t))$ (and $\alpha''((t,v_f))$ accordingly). Since $|\mathrm{int}(\sigma'_t)| < k$ if $\sigma'_t \notin \mathcal{T}^*$, the image of α'' is contained in $\vec{S}_k(G)$;

[9] Pick any node r of T and for every separation $(A,B) \in \sigma_r$ a neighbour $t_{(A,B)}$ of r such that $\alpha(t_{(A,B)}, r) = (A,B)$. Deleting from T all components of $T-r$ that do not contain $t_{(A,B)}$ for any $(A,B) \in \sigma_r$ turns (T,α) into an $S_k(G)$-tree $(T', \alpha\restriction_{T'})$ in which the neighbourhood of r has changed but σ_r has not, and neither has σ_t for any other node t of T'. So $(T', \alpha\restriction_{T'})$ is still a weakly exhaustive $S_k(G)$-tree over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$. Now think of T' as rooted in r and proceed along its levels.

so (T'', α'') is again an $S_k(G)$-tree, which by definition is over $\mathcal{T}^*$. □

We conclude this section with an example that shows that Theorem 13 fails for sets $\mathcal{F}$ of stars that are nice but not finitely bounded.

Example 5.6.4. Let $G = (V, E)$ be the graph with vertex set $V := \{v_{ij} : (i,j) \in [4] \times \mathbb{N}\}$ and edge set $E := \{\{v_{ij}, v_{i'j'}\} \in V(G) : j' \in \{j, j+1\}\}$ (see Figure 5.1). Set $\mathcal{F}' := \{\{(A_k, B_k)\} : k \in \mathbb{N}\}$ where $(A_k, B_k) := (\{v_{ij} : i \in [4], j \geqslant k\}, \{v_{ij} : i \in [4], j \leqslant k\})$ (see Figure 5.1), and let $\mathcal{F} := \mathcal{F}' \cup \mathcal{P}'_5 \cup \{(V(G), A) : A \subseteq V(G), |A| < 5\}$. Clearly, $\mathcal{F}$ is strongly closed under shifting, and thus a nice set of stars in $\vec{S}_5(G)$.

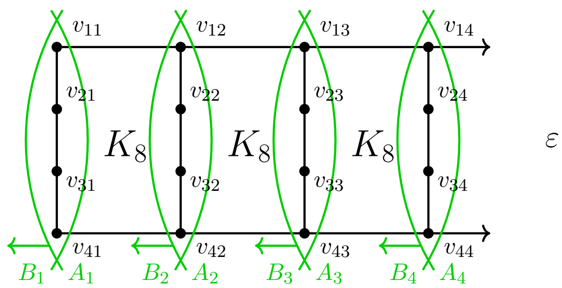

Figure 5.1: Example 5.6.4

It is easy to check that G has precisely one $\mathcal{F}$-tangle of order 5, the one induced by its end ε. Indeed, any consistent orientation of $S_5(G)$ that is not induced by ε has to contain some (B_i, A_i), and is hence not an $\mathcal{F}$-tangle. Thus, G has no $\mathcal{F}$-tangle of order 5 that is not induced by an end of combined degree < 5.

But G has no $S_5(G)$-tree over $\mathcal{F} \cup \mathcal{U}_k^\infty$ either. Indeed, any such tree would have to contain a ray whose edges are associated with separations that form an increasing sequence in τ_ε. By the definition of $\mathcal{P}'_k$, and because $\mathcal{U}_k^\infty$ is empty since G is locally finite, the nodes of that ray would eventually be associated with stars in $\mathcal{F}'$, a contradiction because $\mathcal{F}'$ contains only singleton stars.

5.7 Bramble-treewidth duality: an application of the tangle-tree duality theorem

A set U of vertices of a graph G is *connected* if $G[U]$ is connected. A *bramble* in G is a set $\mathcal{B}$ of mutually touching connected sets of vertices of G where two sets of vertices are said to *touch* if they have a vertex in common or if G contains an edge between them. The *order* of a bramble is the least number of vertices that *cover* the bramble, in that they meet every element of it.

Seymour and Thomas proved the following duality between high-order brambles and small tree-width (see also [41, Theorem 12.4.3]):

Theorem 5.7.1 ([117])**.** *Let $k \in \mathbb{N}$. A finite graph has tree-width $< k$ if and only if it contains no bramble of order $> k$.*

Theorem 5.7.1 extends to infinite graphs with one adaptation. For this, let us first note that every graph G with a ray contains a bramble of infinite order. Indeed, if $R = r_0 r_1 \dots$ is a ray in G, then $\mathcal{B} := \{\{r_i : i \geqslant n\} : n \in \mathbb{N}\}$ is a bramble, and it clearly cannot be covered by finitely many vertices. However, the graph that consists of just a single ray has clearly tree-width 1.

Thus, in order to ensure that brambles of high order force the tree-width of a graph up, we have to restrict the class of brambles we consider to those that are finite. Here, a bramble is *finite* if all its elements are finite. Note that, clearly, every bramble in a finite graph is finite.

With this definition, Theorem 5.7.1 extends to infinite graphs:

Theorem 5.7.2. *Let $k \in \mathbb{N}$. A graph has tree-width $< k$ if and only if it contains no finite bramble of order $> k$.*

The probably shortest way to prove Theorem 5.7.2 is to use a result discovered by Thomas (see [121, Theorem 14] for a proof) which says that an infinite graph G has tree-width $\leqslant k$ if and only if every finite subgraph of G has tree-width $\leqslant k$. Hence, if a graph G has tree-width $\geqslant k$ for some $k \in \mathbb{N}$, then some finite subgraph $H \subseteq G$ has tree-width $\geqslant k$. Theorem 5.7.1 then yields a bramble $\mathcal{B}$ in H of order $> k$. Since H is finite, $\mathcal{B}$ will be finite as well, and it is easy to see that $\mathcal{B}$ is also a bramble of order $> k$ in G. The other implication can be proved similarly to the finite case.

In this section we present an alternative proof of Theorem 5.7.2, which derives Theorem 5.7.2 from Theorem 16. Though this proof is not as short as the one indicated above, it provides an example of how Theorem 16 can be employed to obtain other duality theorems for infinite graphs. Moreover, the proof we present in this section is direct in that it does not make use of the finite result (Theorem 5.7.1); but, of course, it easily implies it. Further, we prove Theorem 5.7.2 by showing a more general duality that includes brambles and tree-width as well as $\mathcal{U}_k$-tangles and $S_k(G)$-trees over $\mathcal{U}_k$.

Recall that $\mathcal{U}_k := \{\sigma \subseteq \vec{S}_{\aleph_0}(G) : \sigma \text{ is a star with } |\operatorname{int}(\sigma)| < k\}$ for $k \in \mathbb{N}$. The main result of this section is Theorem 14, which we restate here for convenience:

Theorem 14. *The following assertions are equivalent for all graphs G and $k \in \mathbb{N}$:*

(i) *G has a $\mathcal{U}_k$-tangle of order k that is not induced by an end of combined degree $< k$.*

(ii) *G has a finite bramble of order at least k.*

(iii) *G has no weakly exhaustive $S_k(G)$-tree over $\mathcal{U}_k$.*

(iv) *G has tree-width at least $k-1$.*

Theorem 14 generalizes a result of Diestel and Oum [54, Theorem 6.5] to infinite graphs; its proof is inspired by theirs.

Proof of Theorem 14. (i) $\Leftrightarrow$ (iii) is Lemma 5.2.3 and Theorem 16. (iii) $\Leftrightarrow$ (iv) is analogous to [54, Lemma 6.3].

It remains to show (i) $\Leftrightarrow$ (ii). For (ii) $\Rightarrow$ (i) we closely follow the proof of [54, Lemma 6.4]. Let $\mathcal{B}$ be a finite bramble of order at least k. For every $\{A, B\} \in S_k(G)$, since $A \cap B$ is too small to cover $\mathcal{B}$ but every two sets in $\mathcal{B}$ touch and are connected, exactly one of the sets $A \setminus B$ and $B \setminus A$ contains an element of $\mathcal{B}$. Thus, $O := \{(A, B) \in \vec{S}_k(G) : B \setminus A \text{ contains an element of } \mathcal{B}\}$ is an orientation of $S_k(G)$, and it is consistent for the same reason.

To show that O avoids $\mathcal{U}_k$, let $\sigma = \{(A_i, B_i) : i \in I\} \in \mathcal{U}_k$ be given. Then $|\operatorname{int}(\sigma)| < k$, so some $C \in \mathcal{B}$ avoids $\operatorname{int}(\sigma)$, and hence lies in the union of the sets $A_i \setminus B_i$. But these sets are disjoint, since σ is a star, and they have no edges between them. Hence, C lies in one of them, $A_j \setminus B_j$ say, which implies that $(B_j, A_j) \in O$. But then $(A_j, B_j) \notin O$, so $\sigma \not\subseteq O$ as claimed.

To see that O is not induced by an end of G of combined degree $< k$, let ε be an end of G. If $\Delta(\varepsilon) < k$, then in particular $\operatorname{dom}(\varepsilon) < k$. Since $\mathcal{B}$ has order at least k, there exists some $C \in \mathcal{B}$ that avoids $\operatorname{Dom}(\varepsilon)$. Moreover, again because $\Delta(\varepsilon) < k$, there exists, by Lemma 5.3.2 (ii), a weakly exhaustive increasing sequence $((A_i, B_i))_{i \in \mathbb{N}}$ of separations in τ_ε such that $(A_i \cap B_i) \cap (A_j \cap B_j) \subseteq \operatorname{Dom}(\varepsilon)$ for $i \neq j \in \mathbb{N}$. Since $C \cap \operatorname{Dom}(\varepsilon) = \emptyset$ and C is finite, this implies that there is some $j \in \mathbb{N}$ such that $C \subseteq A_j \setminus B_j$. But then $(B_j, A_j) \in O$, and thus O is not induced by ε.

For (i) $\Rightarrow$ (ii) assume that G has an $\mathcal{U}_k$-tangle τ of order k that is not induced by an end of combined degree $< k$.

Claim 1. For every separation $(A, B) \in \tau$ the set $\{(C, D) \in \tau : (A, B) \leqslant (C, D)\}$ has a maximal element.

Proof. Suppose towards a contradiction that the set $S := \{(C, D) \in \tau : (A, B) \leqslant (C, D)\}$ has no maximal element, and let $S' \subseteq S$ consist of all separations (C, D) in S that are right-tight and have a connected strict right side $D \setminus C$. Since τ avoids $\mathcal{U}_k$ and is hence principal, there exists for every $(C, D) \in \tau$ a component K of $G - (A \cap B)$ such that $(V(G - K), V(K) \cup N_G(K)) \in \tau$. In particular, for every $(C, D) \in S$ there exists $(C', D') \in S'$ such that $(C, D) \leqslant (C', D')$. It follows that S' is non-empty and that S' has no maximal element either. Thus, S' contains a strictly increasing sequence $((C_i, D_i))_{i \in \mathbb{N}}$ of separations.

Since all (C_i, D_i) have order $< k$, we may assume that $|C_i \cap D_i| = \ell$ for some $\ell < k$ and all $i \in \mathbb{N}$, by passing to a subsequence of $((C_i, D_i))_{i \in \mathbb{N}}$ if necessary. We claim that $((C_i, D_i))_{i \in \mathbb{N}}$ is weakly exhaustive. Indeed, since all $D_i \setminus C_i$ are connected, all separators $C_i \cap D_i$ are distinct. Hence, $X := \bigcap_{i\in\mathbb{N}} \bigcup_{j\geqslant i}(C_j \cap D_j)$ has size $< \ell$; let $j \in \mathbb{N}$ such that $X \subseteq C_i \cap D_i$ for all $i \geqslant j$. Pick some $u \in (C_j \cap D_j) \setminus X$, and let $j' \in \mathbb{N}$ such that $u \in C_{j'} \setminus D_{j'}$. Since (C_j, D_j) is right-tight and $D_j \setminus C_j$ is connected, there is for every $v \in D_j \setminus C_j$ a v–u path in G that avoids X. So if $D := \bigcap_{i\in\mathbb{N}} D_i \setminus C_i$ is non-empty, then there exists a finite D–$(C_{j'} \setminus D_{j'})$ path that avoids X. But this path has to meet all separators $C_i \cap D_i$ with $i \geqslant j'$ in vertices outside of X, a contradiction.

Thus, $((C_i, D_i))_{i\in\mathbb{N}}$ is weakly exhaustive, which by Lemma 5.6.1 contradicts the assumption that τ is not induced by an end of combined degree $< k$. ■

We now define a finite bramble $\mathcal{B}$ as follows. Let $V(G)$ be equipped with a fixed well-ordering. Then for every non-empty set $U \subseteq V(G)$ there exists a unique element in U which is least in the well-ordering; we denote this vertex with v_U.

Now for every separation $(A, B) \in \tau$ that is maximal in τ (with respect to the partial order on τ induced by $\vec{S}_k(G)$), we pick a finite, connected set $U_{(A,B)} \subseteq B \setminus A$ which contains $v_{B\setminus A}$ and for every vertex in $A \cap B$ at least one of its neighbours. For this note that such a set exists since (A, B) is maximal in τ, and thus $G[B \setminus A]$ is connected and $N_G(B \setminus A) = A \cap B$. We then put in $\mathcal{B}$ precisely all sets $U_{A,B}$. By definition, all elements of $\mathcal{B}$ are finite and connected. Moreover, $\mathcal{B}$ has order at least k. Indeed, since τ avoids $\mathcal{U}_k$, there exists for every set U of at most $k-1$ vertices of G a component C of $G - U$ such that $(V(G-C), V(C) \cup N_G(C)) \in \tau$. Then for every maximal separation (A, B) in τ with $(V(G-C), V(C) \cup N_G(C)) \leqslant (A, B)$ the set $U_{(A,B)}$ avoids U; and such an (A, B) exists by Claim 1.

To conclude the proof, it remains to show that the sets in $\mathcal{B}$ mutually touch. For this, let $U := U_{(A,B)}$, $U' := U_{(A',B')} \in \mathcal{B}$ be given. If $v := v_{B\setminus A} = v_{B'\setminus A'} =: v'$, then U and U' intersect by construction. So one of v and v', say v, is strictly smaller in the well-ordering, which by the choice of v' implies that $v \in A'$. We claim that $(A \cap B) \cap (B' \setminus A') \neq \emptyset$. Then the assertion follows. Indeed, let u be any vertex in that set. Since $A \cap B \subseteq N_G(U)$ by the choice of U, there is a v–u path in $G[U \cup \{u\}]$. But as $v \in A'$ and $u \in B' \setminus A'$, it follows that U meets $A' \cap B'$. Hence U and U' touch since $A' \cap B' \subseteq N_G(U')$ by the choice of U'.

To prove the claim suppose towards a contradiction that $(A \cap B) \cap (B' \setminus A') = \emptyset$. If also $(B \setminus A) \cap (B' \setminus A') = \emptyset$, then it follows that $(B', A') \leqslant (A, B)$, which contradicts the consistency of τ as $(A', B'), (A, B) \in \tau$. Hence, there is a vertex $u' \in (B \setminus A) \cap (B' \setminus A')$. Since both (A', B') and (A, B) are maximal separations in τ, we have $B' \setminus A' \nsubseteq B \setminus A$, and hence the exists a vertex $w \in (B' \setminus A') \cap A \neq \emptyset$. As $B' \setminus A'$ is connected, there exists a u'–w path in $G[B' \setminus A']$. But this path has to meet $A \cap B$ since $u' \in B$ and $w \in A$, which concludes the proof. □

Proof of Theorem 5.7.2. This is (ii) ⇔ (iv) of Theorem 14. □

We conclude this section with the following corollary, which we used in Chapter 3:

Corollary 5.7.3. *Let G be a graph of tree-width $\leqslant w \in \mathbb{N}$, and let σ be a finite star of separations of G of order $\leqslant w+1$ whose interior is finite. Suppose that all separations in σ are left-ℓ-robust for $\ell := (w+1)^2(w+2)+w+1$. Then* $\operatorname{torso}(\sigma)$ *has tree-width $\leqslant w$.*

Proof. By definition, $\mathcal{U}_{w+2}$ is $(w+1)$-bounded, and by Lemma 5.2.3, $\mathcal{U}_{w+2}$ is nice, so we may apply Lemma 5.5.4, which yields that there is either an $\mathcal{U}_{w+2}$-tangle τ of G with $\sigma \subseteq \tau$ or a finite $S_{w+2}(G)$-tree (T, α) over $\mathcal{U}_{w+2} \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma\}$ in which each $\vec{s} \in \sigma$ appears as a leaf separation. Suppose first that the former holds. Since the interior of σ is finite and $\sigma \subseteq \tau$, the $\mathcal{U}_{w+2}$-tangle τ cannot be induced by an end of G. But since G has tree-width $\leqslant w$, it has no $\mathcal{U}_{w+2}$-tangles of order $w+2$ that are not induced by an end by Theorem 14, a contradiction. So we may assume the latter. It is easy to check that (T, α) induces a tree-decomposition $(T, \mathcal{V})$ of G (cf. [54, Lemma 6.3]) whose bags have size $\leqslant w+1$, unless they are associated with leaves of T whose incident edge induces a separation in σ. By restricting the bags in $\mathcal{V}$ to $\operatorname{int}(\sigma)$, we obtain a tree-decomposition $(T, \mathcal{V}')$ of $G[\operatorname{int}(\sigma)]$ of width $\leqslant w$. In fact, since (T, α) contains all $\vec{s} \in \sigma$ as leaf separations, $(T, \mathcal{V}')$ is even tree-decomposition of $\operatorname{torso}(\sigma)$. Thus, $\operatorname{torso}(\sigma)$ has tree-width $\leqslant w$. □

5.8 Refining trees of tangles in infinite graphs

Besides the tangle-tree duality theorem, Robertson and Seymour [114] proved the *tree-of-tangles theorem*, which asserts that for every $k \in \mathbb{N}$ every finite graph has a tree-decomposition such that its k-tangles live at different nodes of the tree. Erde [66] combined this theorem and the tangle-tree duality theorem into one, by constructing a single tree-decomposition such that every node either accommodates a single k-tangle or is too small to accommodate one, in that it is associated with a star in $\mathcal{T}_k$. In fact, he showed that such a tree-decomposition can be obtained from any given one that efficiently distinguishes all the k-tangles, by refining its inessential parts.

The author [3] improved Erde's result by constructing further refinements of the essential parts of that tree-decomposition, yielding a tree-decomposition that has the additional property that all its essential bags are as small as possible. In this section, we extend this result to infinite graphs. We then obtain Theorem 15 as a simple corollary.

To state the main result of this section, we first need some further definitions. Following [65], we call two regular k-profiles τ, τ' in a graph G *combinatorially distinguishable* if at least one of them is principal or they are both non-principal but such that there exists a set $X \subseteq V(G)$ such that $(V(K) \cup X, V(G-K)) \in \tau$ for all $K \in \mathcal{C}_X$ and such that $(V(G-K), V(K) \cup X) \in \tau'$ for a component $K \in \mathcal{C}_X$.

A k-profile in G is *bounded* if it does not extend to an $\aleph_0$-profile.[10]

[10]Equivalently, a principal k-profile τ in G is bounded if and only if it is neither induced by an end nor of the form $\{(A,B) \in \vec{S}_k(G) : X \subseteq B\}$ for a set $X \in \operatorname{crit}(G)$ of size $\geqslant k$ (cf. [40, Theorem 3]). Moreover, every non-principal

Theorem 5.8.1. *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a finitely bounded, profile-respecting, nice set of finite stars in $S_k(G)$. Further, let $(\tilde{T}, \tilde{\mathcal{V}})$ be a tree-decomposition of G which distinguishes all combinatorially distinguishable $\mathcal{F}$-tangles of order k such that every separation induced by an edge of $\tilde{T}$ distinguishes a pair of $\mathcal{F}$-tangles of order k efficiently. Then there exists a tree-decomposition $(T, \mathcal{V})$ of G which refines $(\tilde{T}, \tilde{\mathcal{V}})$ and which is such that*

(i) *every end of G of combined degree $< k$ lives in an end of T;*

(ii) *if every end of $\tilde{T}$ is home to an end of G, then also every end of T is home to an end of G;*

(iii) *every non-principal $\mathcal{F}$-tangle of order k which does not live in an end of $\tilde{T}$ lives at a node t of T with $\sigma_t \in \mathcal{U}_k^\infty$;*

(iv) *for every inessential node t of T we have $\sigma_t \in \mathcal{F}$; and*

(v) *every bag V_t of $(T, \mathcal{V})$ that is home to a bounded $\mathcal{F}$-tangle of order k is of smallest size among all the exclusive bags of tree-decompositions of G that are home to the $\mathcal{F}$-tangle living in V_t.*

We remark that if G is locally finite, then there exists by [31, Theorem 7.3] for every $k \in \mathbb{N}$ a tree-decomposition $(\tilde{T}, \tilde{\mathcal{V}})$ of G which efficiently distinguishes all k-tangles in G, and which thus satisfies the premise of Theorem 5.8.1.

In the remainder of this section we prove Theorem 5.8.1. For this, we need two more refining lemmas. The first one lets us refine stars which are home to a bounded tangle, and the second one generalizes our refining lemma for inessential stars, Lemma 5.5.4, to stars whose interior is infinite.

To show the first lemma, we need the following result of [3]:

Lemma 5.8.2 ([3, Proof of Lemma 4.3])**.** *Let $k \in \mathbb{N}$, let $\mathcal{Q}$ be some set of k-profiles in a graph*[11] *G, and let $\tau \in \mathcal{Q}$. Further, let $\sigma \subseteq \tau$ be a finite star with finite interior, and suppose that every separation in σ efficiently distinguishes some pair of k-profiles in $\mathcal{Q}$. Then there exists a star $\varrho \subseteq \tau$ with $\sigma \leqslant \varrho$ whose interior is of smallest size among all stars in τ that are exclusive for $\mathcal{Q}$ and which has the further property that all the separations in ϱ are closely related to τ.*

The following lemma refines essential stars that are home to a bounded tangle in a similar way as Lemma 5.5.4 refines inessential stars. It generalizes [3, Lemma 4.3] to infinite graphs.

Lemma 5.8.3. *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a finitely bounded, profile-respecting, nice set of finite stars in $\vec{S}_k(G)$. Let $\sigma \subseteq \vec{S}_k(G)$ be a star, and suppose that every separation in σ efficiently distinguishes some pair of $\mathcal{F}$-tangles of $S_k(G)$. Further, suppose there is a unique $\mathcal{F}$-tangle τ of $S_k(G)$ that satisfies $\sigma \subseteq \tau$. If τ is bounded, then there exists a star $\sigma' \subseteq \tau$ whose interior is of smallest size among all exclusive stars in τ, and a finite $S_k(G)$-tree over $\mathcal{F} \cup \{\sigma'\} \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma\}$ in which each $\vec{s} \in \sigma$ appears as a leaf separation.*

k-profile is unbounded.

[11]In [3] the assertion of Lemma 5.8.2 is shown only for finite graphs. However, the same proof works for infinite graphs as long as σ is finite and has finite interior.

Note that we show in the proof that the interior of every exclusive star in τ is finite. Thus, 'of smallest size' is well-defined.

Proof. Since τ is bounded and thus not induced by an end, and because σ is τ-exclusive, no end of G lives in σ by Proposition 5.3.1. Hence, by Proposition 3.6.1 and because all separations in σ are left-tight by Lemma 2.2.1, torso(σ) is rayless. Moreover, by Lemmas 5.4.3 and 5.4.4, torso(σ) is tough, and so by Theorem 2.5.2, int(σ) is finite. In particular, again by Lemma 5.4.4, σ is finite.

We can thus apply Lemma 5.8.2 to σ and τ, which yields a star $\sigma' \subseteq \tau$ with $\sigma \leqslant \sigma'$ which is closely related to τ and whose interior is of smallest size among all exclusive stars in τ. As all separations in σ distinguish some pair of $\mathcal{F}$-tangles in G efficiently, and are hence closely related to some $\mathcal{F}$-tangle by Proposition 5.5.1, it follows that all the stars $\varrho_s := \{\overleftarrow{s}\} \cup \{\vec{r} : \vec{r} \leqslant \vec{s}, \vec{r} \in \sigma\}$ for $\vec{s} \in \sigma'$ satisfy the premise of Lemma 5.5.4. We thus obtain, for every ϱ_s, a finite $S_k(G)$-tree (T^s, α^s) over $\mathcal{F} \cup \{\{\overleftarrow{r}\} : \vec{r} \in \varrho_s\}$ in which each $\vec{r} \in \varrho_s$ appears as a leaf separation.

Let T be the tree obtained from the disjoint union of the trees T^s by identifying their leaves v_s where v_s is the unique leaf of T^s whose incident edge induces $\vec{s}$. Then (T, α) with $\alpha(\vec{e}) = \alpha^s(\vec{e})$ where T^s is the unique tree containing e is an $S_k(G)$-tree over $\mathcal{F} \cup \{\sigma'\} \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma\}$. □

The next lemma generalizes Lemma 5.5.4 to stars with infinite interior:

Lemma 5.8.4. *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a finitely bounded, nice set of stars in $\vec{S}_k(G)$. Let $\sigma := \{\vec{s}_i : i \in I\} \subseteq \vec{S}_k(G)$ be a star, and suppose that every separation in σ efficiently distinguishes some pair of k-profiles in G that avoid $\mathcal{F}$. Set $\mathcal{F}' := \mathcal{F} \cup \{\{\overleftarrow{s}_i\} : i \in I\}$. Then either there is a principal $\mathcal{F}'$-tangle of $S_k(G)$ that is not induced by an end of G of combined degree $< k$, or there is a weakly exhaustive $S_k(G)$-tree (T, α) over $\mathcal{F}' \cup \mathcal{U}_k^\infty$ in which each $\vec{s}_i$ appears as a leaf separation.*

In particular, (T, α) can be chosen so that every end of T is home to an end of G.

If $(T, \mathcal{V})$ is a tree-decomposition of a graph G, then a tree T' obtained from T by edge contractions induces the tree-decomposition $(T', \mathcal{V}')$ of G whose bags are $V'_t = \bigcup_{s \in t} V_s$ for every $t \in T'$, where we denote the vertex set of T' as the set of branch sets.

Before we prove Lemma 5.8.4, we recall Lemma 3.8.6 from Chapter 3:

Lemma 5.8.5. *Let G be a graph, and let $(T, \mathcal{V})$ be the tree-decomposition of G that satisfies* (ii) *and* (iv) *of Theorem 5.6.2. Let F be some set of edges of T such that no edge in F is incident with a node of infinite degree and such that for every end η of T, the set F avoids cofinitely many edges e of the η-ray R in T with $|V_e| = \Delta(\varepsilon)$ where ε is the end of G that lives in η. Then the tree-decomposition obtained from $(T, \mathcal{V})$ by contacting all edges in F has still finite parts.*

Proof of Lemma 5.8.4. Let us assume that there is no $\mathcal{F}'$-tangle that is as desired and show that there then exists a weakly exhaustive $S_k(G)$-tree over $\mathcal{F}' \cup \mathcal{U}_k^\infty$. Set $\ell := \max\{3k-2, k(k-1)m+m\}$ where $m \in \mathbb{N}$ is such that $\mathcal{F}$ is m-bounded.

Assume first that σ has finite interior. Deviating from the assumptions of the lemma, we allow in this case that separations in σ do not efficiently distinguish two $\mathcal{F}$-tangles if they are left-ℓ-robust instead. We will use this in second case below. We show that there then exists an $S_k(G)$-tree (T, α) as desired such that T is rayless. Set $\mathcal{X} := \{X \subseteq V(G) : X = A \cap B$ for infinitely many $(A, B) \in \sigma\}$ and $\varrho := \{(A, B) \in \sigma : A \cap B \notin \mathcal{X}\} \cup \{(A_X, B_X) : X \in \mathcal{X}\}$ where (A_X, B_X) is the supremum over all $(A, B) \in \sigma$ with $A \cap B = X$. Clearly, we have $A_X \cap B_X = X$, and hence ϱ is included in $\vec{S}_k(G)$. We claim that ϱ satisfies the premise of Lemma 5.5.4. Indeed, ϱ is still a star, and it is finite since $\operatorname{int}(\sigma)$ is finite. Moreover, every separation in $\varrho \cap \sigma$ is either left-ℓ-robust or has an inverse that is closely related to some $\mathcal{F}$-tangle of order k by Proposition 5.5.1. Further, every separation $(A, B) \in \varrho \setminus \sigma$ is left-ℓ-robust by Proposition 5.5.3, as $G[A \setminus B]$ contains infinitely many tight components of $G - (A \cap B)$ by Lemma 2.2.1.

Hence, we may apply Lemma 5.5.4 to ϱ. As ϱ is finite and has finite interior, it cannot be home to any $\mathcal{F}$-tangles that are non-principal or induced by an end by Lemma 5.2.1 and Proposition 5.3.1. Since ϱ is also not home to any other principal $\mathcal{F}$-tangles by assumption, we thus obtain a finite $S_k(G)$-tree (T, α) over $\mathcal{F} \cup \{\{\overleftarrow{s}\} : \vec{s} \in \varrho\}$ in which each $(A, B) \in \varrho$ appears as a leaf separation; let us denote the respective leaf of T with $t_{(A,B)}$. We now obtain the desired rayless $S_k(G)$-tree over $\mathcal{F} \cup \mathcal{U}_k^\infty$ by adding, for every $(A, B) \in \sigma \setminus \varrho$ a leaf $u_{(A,B)}$ and the edge $\{u_{(A,B)}, t_{(A_X,B_X)}\}$ where $X = A \cap B$.

We now turn to the general case that σ may have infinite interior. By Proposition 5.3.1, no end of G of combined degree $\geqslant k$ lives in σ, and hence, by Lemma 5.3.3, all ends of $\operatorname{torso}(\sigma)$ have combined degree $< k$. Further, by Lemmas 5.4.1 and 5.4.3, all critical vertex sets of $\operatorname{torso}(\sigma)$ have size $< k$. Hence, we may apply Theorem 5.6.2 to $\operatorname{torso}(\sigma)$, which yields a tree-decomposition $(T', \mathcal{V}')$ of $\operatorname{torso}(\sigma)$ of adhesion $< k$. Since every separator $A \cap B$ of a separation $(A, B) \in \sigma$ is complete in $\operatorname{torso}(\sigma)$, there exists for every such $A \cap B$ a node $t_{A \cap B}$ of T' such that $A \cap B \subseteq V'_{t_{A \cap B}}$. We then obtain a tree-decomposition $(T'', \mathcal{V}'')$ of G by adding for every $(A, B) \in \sigma$ a node $u_{(A,B)}$ and the edge $\{u_{(A,B)}, t_{A \cap B}\}$ to T' and assigning the bag A to $u_{(A,B)}$.

Let F be the set of edges $e = \{s, t\}$ of T'' with $\sigma''_s, \sigma''_t \notin \mathcal{U}_k^\infty$ whose induced separation is neither ℓ-robust nor in σ. Note that if $\sigma''_s \in \mathcal{U}_k^\infty$, then $\vec{e} = (s, t)$ is left-ℓ-robust by Proposition 5.5.3: infinitely many separations in σ'_s are left-tight by Theorem 5.6.2 (i), and hence, as all separations in σ are left-tight, also infinitely many separations in σ''_s are left-tight; so (U_s, U_t) contains infinitely many tight components of $G - V_e$ in its small side.

Let $(\tilde{T}'', \tilde{\mathcal{V}}'')$ and $(\tilde{T}', \tilde{\mathcal{V}}')$ be the tree-decompositions obtained from $(T'', \mathcal{V}'')$ and $(T', \mathcal{V}')$, respectively, by contracting all edges in F. Then all bags $\tilde{V}''_t$ for nodes t of $\tilde{T}''$ not of the form $u_{(A,B)}$ are finite: By Theorem 5.6.2 (ii) every end η of T' is home to an end ε' of torso(σ) with $\liminf_{e\in R}|V_e| = \Delta(\varepsilon')$ where R is any η-ray. Then by Lemma 5.3.3 the same holds true for the ray R in $T'' \supseteq T'$ and some end ε of G with $\Delta(\varepsilon) = \Delta(\varepsilon')$. Thus, by Lemma 5.5.2 applied to ε and the separations of G induced by the ray R in T'', the set F and $(T', \mathcal{V}')$ satisfy the premise of Lemma 5.8.5, and thus $(\tilde{T}', \tilde{\mathcal{V}}')$ has still finite parts. As $\tilde{V}''_t = \tilde{V}'_t$ for all nodes t of $\tilde{T}''$ not of the form $u_{(A,B)}$, these bags are finite.

Now let t be a node of $\tilde{T}''$ which is neither a leaf of the form $u_{(A,B)}$ for some $(A, B) \in \sigma$ nor associated with a star $\tilde{\sigma}''_t$ in $\mathcal{U}^\infty_k$. Then all separations in $\tilde{\sigma}''_t$ are either left-ℓ-robust or efficiently distinguish two $\mathcal{F}$-tangles by construction. Since $\mathrm{int}(\tilde{\sigma}''_t)$ is finite, we may apply the first case to $\tilde{\sigma}''_t$. As σ is not home to any principal $\mathcal{F}$-tangles that are not induced by an end of combined degree $< k$, $\tilde{\sigma}''_t$ is not home to any such $\mathcal{F}$-tangles either. Hence, we obtain a rayless $S_k(G)$-tree (T^t, α^t) over $\mathcal{F} \cup \mathcal{U}^\infty_k \cup \{\{\overleftarrow{s}\} : \vec{s} \in \tilde{\sigma}''_t\}$ in which each $\vec{s} \in \tilde{\sigma}''_t$ appears as a leaf separation. Applying Construction 5.6.3 to $(\tilde{T}'', \tilde{\mathcal{V}}'')$ and the (T^t, α^t) yields a weakly exhaustive $S_k(G)$-tree (T, α) over $\mathcal{F} \cup \mathcal{U}^\infty_k \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma\}$. Moreover, by construction, each $\vec{s}_i \in \sigma$ appears as a leaf separation of (T, α). For the 'in particular'-part, note that every end of T' is home to an end of torso(σ) by Theorem 5.6.2 (ii), and hence every end of $\tilde{T}''$ is home to an end of G by Lemma 5.3.3. Since all T^t are rayless, the assertion follows. □

With Lemmas 5.8.3 and 5.8.4 at hand, we are ready to prove the main result of this section.

Proof of Theorem 5.8.1. We define for every node t of $\tilde{T}$ an $S_k(G)$-tree (T^t, α^t) as follows. If t is home to a bounded $\mathcal{F}$-tangle, then let (T^t, α^t) be the finite $S_k(G)$-tree obtained from applying Lemma 5.8.3 to $\tilde{\sigma}_t$. If t is inessential or only home to $\mathcal{F}$-tangles of order k that are either non-principal or induced by ends of combined degree $< k$, then let (T^t, α^t) be the $S_k(G)$-tree obtained from applying Lemma 5.8.4 to $\tilde{\sigma}_t$.

Then applying Construction 5.6.3 to $(\tilde{T}, \tilde{\mathcal{V}})$ and the (T^t, α^t) yields a weakly exhaustive $S_k(G)$-tree (T, α). It is now straightforward to check that $(T, \mathcal{V})$ with $V_t := \mathrm{int}(\sigma_t)$ is a tree-decomposition of G (cf. [29, §4]); in particular, every edge $\vec{e}$ of T induces the separation $\alpha(\vec{e})$. Then $(T, \mathcal{V})$ satisfies (i) to (v): By construction, no end of combined degree $< k$ can live at a node of T, and hence they have to live in ends of T; so (i) holds. (ii) follows by the 'in particular' part of Lemma 5.8.4. Property (iii) holds because non-principal $\mathcal{F}$-tangles that live at nodes of $\tilde{T}$ have to live at nodes of T by the 'in particular' part of Lemma 5.8.4, but they cannot live at nodes that are associated with stars in $\mathcal{F}$. (iv) and (v) follow by Lemmas 5.8.3 and 5.8.4. □

We conclude this section with the proof of Theorem 15. For this, we need the following theorem, which is immediate from the proof of Corollary 5' in Chapter 3:

Theorem 5.8.6. *Every graph G without half-grid minor admits a tree-decomposition $(T, \mathcal{V})$ which efficiently distinguishes all combinatorially distinguishable $\aleph_0$-tangles in G.*

Moreover, $(T, \mathcal{V})$ can be chosen so that every end of T is home to an end of G, and every non-principal $\aleph_0$-tangle lives at a node t of T.

Proof of Theorem 15. By assumption, G has no end of combined degree $\geqslant k$, and hence no half-grid minor. Let $(T', \mathcal{V}')$ be the tree-decomposition of G from Theorem 5.8.6. Since G has no bounded k-tangle by assumption, $(T', \mathcal{V}')$ in particular distinguishes all combinatorially distinguishable k-tangles. By contracting all edges of T whose induced separations do not efficiently distinguish some pair of k-tangles, we obtain a tree-decomposition $(\tilde{T}, \tilde{\mathcal{V}})$ that satisfies the premise of Theorem 5.8.1. Let $(T, \mathcal{V})$ be the tree-decomposition obtained from applying Theorem 5.8.1 to $(\tilde{T}, \tilde{\mathcal{V}})$. Then (T, α) is an $S_k(G)$-tree, where $\alpha(t_0, t_1) := (U_{t_0}, U_{t_1})$ for all edges $(t_0, t_1) \in \vec{E}(T)$. In particular, (T, α) is over $\mathcal{T}^* \cup \mathcal{U}_k^\infty$ by (i), (iii) and (iv) of Theorem 5.8.1 and because all principal k-tangles in G are induced by ends of combined degree $< k$.

We claim that (T, α) is as desired. Indeed, it satisfies (i) by Theorem 5.8.1 (i) and (ii) and because every end of T', an hence every end of $\tilde{T}$, is home to an end of G. Moreover, (T, α) satisfies (ii) by Theorem 5.8.1 (iii) and because every non-principal k-tangle in G lives at a node of T' and hence of $\tilde{T}$. Finally, (T, α) satisfies (iii) because $(T, \mathcal{V})$ refines $(\tilde{T}, \tilde{\mathcal{V}})$, and $(\tilde{T}, \tilde{\mathcal{V}})$ distinguishes all combinatorially distinguishable k-tangles in G. □

Part II

Tangles in finite graphs

6 Optimal trees of tangles: refining the essential parts

We combine the two fundamental fixed-order tangle theorems of Robertson and Seymour into a single theorem that implies both, in a best possible way. We show that, for every $k \in \mathbb{N}$, every tree-decomposition of a graph G which efficiently distinguishes all its k-tangles can be refined to a tree-decomposition whose bags are either too small to be home to a k-tangle, or as small as possible while being home to a k-tangle.

This chapter is based on [3, §1-4].

6.1 Introduction

All graphs in this chapter are finite.

Tangles were introduced by Robertson and Seymour [114] as a way to capture 'highly cohesive' structures of a graph. Formally, a tangle of a graph G is an orientation of all its separations up to some order. The idea is that every highly cohesive substructure of G will lie mostly on one side of such a low-order separation, and therefore orients it towards that side. All these orientations, collectively, are then called a tangle.

For a given graph G and an integer $k > 0$ the *tree-of-tangles theorem* asserts the existence of a tree-decomposition $(T, \mathcal{V})$ of G, with $\mathcal{V} = (V_t)_{t \in T}$, say, in which each k-tangle lives in a different part. Such a tree-decomposition provides information about the overall structure of the graph and the location of the tangles inside it. A bag $V_t \in \mathcal{V}$ is called *essential* if there is a tangle living in it, and otherwise *inessential*.

In general, the inessential parts might contain a large portion of G. The tree-decomposition then tells us nothing about the structure of this portion. However, if there are no k-tangles in G at all, there is another theorem which does tell us something about its structure: the *tangle-tree duality theorem*. This asserts the existence of a tree-decomposition in which each part is too small to be home to a tangle, and which thus witnesses that G has no k-tangles (since any k-tangle would have to live in some part).

As inessential parts are not home to any tangles, the tangle-tree duality theorem applies locally to these parts. We thus obtain tree-decompositions of all the inessential parts in $\mathcal{V}$ into smaller parts, too small to accommodate a k-tangle. Can all these tree-decompositions be combined into a single tree-decomposition of G that refines $(T, \mathcal{V})$?

In general, this will not be possible. However, Erde [66] showed that it can be done if those local tree-decompositions are constructed carefully for this purpose. The following theorem follows directly from [66, Lemma 8]:

Theorem 6.1.1. (Erde 2017) *Let G be a graph, $k \geqslant 3$, and let $\mathcal{F} \subseteq 2^{\vec{S}_k(G)}$ be a friendly set of stars. Let $(\tilde{T}, \tilde{\mathcal{V}})$ be a tree-decomposition of G which distinguishes all the $\mathcal{F}$-tangles of $S_k(G)$ and is such that every separation of G induced by an edge of $\tilde{T}$ efficiently distinguishes some pair of $\mathcal{F}$-tangles of $S_k(G)$. Then there exists a tree-decomposition $(T, \mathcal{V})$ of G which refines $(\tilde{T}, \tilde{\mathcal{V}})$ and is such that every star associated with an inessential bag of $(T, \mathcal{V})$ is a star in $\mathcal{F}$.*

See Section 6.3 for the definition of 'friendly'.

The procedure of refining $(T, \mathcal{V})$ by further decomposing its inessential parts can be seen as a way to describe the structure of those parts in a more precise way. However, just as the inessential parts of $(T, \mathcal{V})$ may contain a large portion of G, another large portion may be hidden inside its essential parts. The tree-decomposition $(T, \mathcal{V})$ does not tell us much about this portion. We know from the fact that these parts are essential that there is a tangle living there, but an essential part of $(T, \mathcal{V})$ can also contain many vertices and edges that have nothing to do with that tangle.

To get a better insight into the structure of the essential parts of $(T, \mathcal{V})$ too, it would be desirable to decompose them by further tree-decompositions as well. Overall, it would be useful to have a single tree-decomposition of G whose inessential and essential parts are both as small as possible. As every essential part is home to a tangle, by definition, and every tangle corresponds to some highly cohesive (but maybe fuzzy) substructure of G which cannot be divided into small parts, 'small' can only be achieved relatively to the tangle living in that part. In other words, the essential parts should be as small as possible so that they can still be home to a tangle.

For example, consider the graph G in Figure 6.1a. It contains a central K_6, which induces a 3-tangle τ of G: the set of all separations of G of order at most 2 oriented towards the central K_6. The star of separations $\sigma \subseteq \tau$ indicated in the figure, which distinguishes τ from the other 3-tangles of G (those induced by the other three K_6 subgraphs), lies far away from the central K_6, the 'essence' of τ. This results in an essential part which contains a lot more vertices than those contained in the central K_6. But although τ lives in that part, those other vertices do not really 'belong to τ'.

In Figure 6.1b the star σ has been moved closer to the core of τ, the K_6 that induces it. In this decomposition of G, every essential part contains only the vertices of its K_6, while every inessential part is small in that it contains only few vertices. Clearly, the tree-decomposition in Figure 6.1b captures the structure of G in a more precise way than the tree-decomposition in Figure 6.1a does, and thus provides more information about the graph.

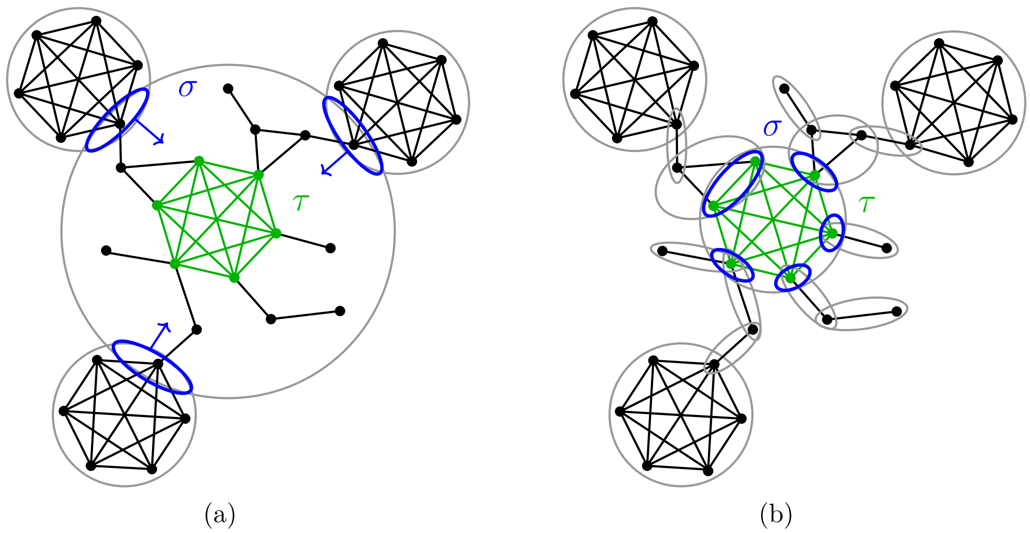

FIGURE 6.1: A tree-decomposition of G that distinguishes all its tangles, and a refinement of its central essential part.

Erde [66] already introduced a method to refine the essential parts of a given tree-decomposition so that the newly arising essential parts do not contain vertices that are 'inessentially separated' from the tangle living in that part. This result comes somewhat close to our goal of decreasing the size of the essential parts by refining them. However, it is not strong enough to decrease the size of the essential parts as much as possible, so that even in the example described above, where the tangle τ is induced by a clique, there are several vertices not in the clique which are not 'inessentially separated' from τ. Thus, if we apply the method from [66] to the essential part in Figure 6.1a that accommodates τ, then the central essential part of the refined tree-decomposition will still contain more vertices than just those in the clique.

In this chapter we show that the essential parts of tangle-distinguishing tree-decompositions can be refined in the best possible way: so that the newly arising essential parts contain as few vertices as possible while being home to a tangle:

Theorem 17. *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}_k(G)$. Let $(\tilde{T}, \tilde{\mathcal{V}})$ be a tree-decomposition of G which distinguishes all the $\mathcal{F}$-tangles of $S_k(G)$ so that every separation induced by an edge of $\tilde{T}$ efficiently distinguishes some pair of $\mathcal{F}$-tangles of $S_k(G)$. Then there exists a tree-decomposition $(T, \mathcal{V})$ of G which refines $(\tilde{T}, \tilde{\mathcal{V}})$ and is such that*

(i) *every star associated with an inessential bag of $(T, \mathcal{V})$ is a star in $\mathcal{F}$;*

(ii) *every essential bag V_t of $(T, \mathcal{V})$ is of smallest size among all the exclusive bags of tree-decompositions of G that are home to the $\mathcal{F}$-tangle living in V_t.*

If we apply Theorem 17 to the set of all k-tangles in G (for some $k \in \mathbb{N}$), then this tree-decomposition $(T, \mathcal{V})$ has the following three properties. First, each of its inessential bags has size $\leqslant 3k-3$. Second, the unique k-tangle τ that lives in an essential bag V_t of size $> 3k-3$ is induced by V_t, in the sense that a separation $\{A, B\}$ of G of order less than k is oriented by τ as (A, B) if and only if $|A \cap V_t| < |B \cap V_t|$. Third, the restriction of τ to $G[V_t]$ is a k-tangle of $G[V_t]$; thus, $G[V_t]$ alone already witnesses that τ is indeed a k-tangle in G.

If a tangle τ in G is induced by a separable k-block in G (see Section 6.5), then the essential bag of our tree-decomposition that is home to τ will be equal to that k-block. Thus, the tree-decomposition $(T, \mathcal{V})$ of G isolates all the separable k-blocks that induce k-tangles, in that they appear as bags of the tree-decomposition. See Theorem 6.5.7 for details.

This chapter is organised as follows. In Section 6.2 we recall some connections between nested sets and S-trees, and in Section 6.3 we give the definition of 'friendly'. We then recall the definitions and lemmas from [4] about refining inessential stars in Section 6.4. In Section 6.5 we prove Theorem 17.

6.2 Interplay between nested sets and S-trees

Let N be a nested set of separations. A star $\sigma \subseteq \vec{N}$ is called a *node of* N if there is a consistent orientation O of N such that σ is the set of maximal elements in O. (In [43] the nodes of N are called 'splitting stars'.)

It is easy to see that the set $\vec{N} := \operatorname{im}(\alpha)$ *induced* by an S-tree (T, α) over a set of stars is nested. Conversely, if N is a *regular* nested set in S, i.e. all separations in N are proper, then one can obtain an S-tree (T, α) as follows. We take the set of all nodes of $\vec{N}$ as the vertex set of T and N as the edges of T where we let a separation $s \in N$ be incident to the two nodes of $\vec{N}$ that contain $\vec{s}$ and $\overleftarrow{s}$, respectively.

Theorem 6.2.1. [43, Theorem 6.9] *Let G be a graph, $k \in \mathbb{N}$ and let $\vec{S} \subseteq \vec{S}_k(G)$. Let further $N \subseteq S$ be a regular nested set. Then there exists an S-tree (T, α) with $\operatorname{im}(\alpha) = \vec{N}$ such that the stars associated with the nodes of T are precisely the nodes of $\vec{N}$.*

This motivates the name 'nodes' for the splitting stars of N. It is shown in [43] that the S-tree from Theorem 6.2.1 is unique up to isomorphisms. Hence, we may say that N *induces* the S-tree (T, α).

We say that a consistent orientation O *lives* at a node σ of N (or equivalently σ *is home* to O) if $\sigma \subseteq O$. Similarly, we say that O *lives* at a node t of an S-tree (T, α) if $\sigma_t \subseteq O$. It is easy to see that every consistent orientation of S has to live at a (unique) node of any regular nested set N.

Given a set $\mathcal{O}$ of consistent orientations of S and an S-tree (T, α), we call a node $t \in T$ *essential (for $\mathcal{O}$)* if there is an orientation in $\mathcal{O}$ which lives at t and otherwise *inessential (for $\mathcal{O}$)*.

Similar as to S-trees, it is easy to check that the set of separations induced by the edges of a tree-decomposition is nested. Conversely, it follows from Theorem 6.2.1 that every regular nested set N of separations of a graph G is induced by the edges of some tree-decomposition $(T, \mathcal{V})$ of G. In fact, if we require the tree T to be minimal, then there is a unique such tree-decomposition (up to isomorphisms). We say that N *induces* this tree-decomposition. In particular, if $(T, \mathcal{V})$ is induced by N, then the stars σ_t associated with the bags V_t of $(T, \mathcal{V})$ are precisely the nodes of N.

6.3 Friendly sets of stars

Given a graph G and some $k \in \mathbb{N}$, a set $\mathcal{F}$ of stars in $\vec{S}$ is *friendly* if $\mathcal{F}$ is standard, profile-respecting, closed under shifting and contains $\{(V(G), A)\}$ for every separation $\{A, V(G)\} \in S$ where $A \subseteq V(G)$.

The condition that a set of stars is friendly might seem to be rather strong at first, but in practice it can typically be satisfied. Diestel, Eberenz and Erde [45] showed that for any graph G and $k \in \mathbb{N}$ any set $\mathcal{F}$ of subsets of $\vec{S}_k(G)$ can be transformed into a standard set $\hat{\mathcal{F}}$ of stars which is closed under shifting so that an orientation of $S_k(G)$ is a regular $\hat{\mathcal{F}}$-tangle if and only if it is a regular $\mathcal{F}$-tangle.

Lemma 6.3.1. [45, Lemma 11 & 14] *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F} \subseteq 2^{\vec{S}_k(G)}$ be a standard set. Then there exists a standard set $\hat{\mathcal{F}}$ of stars in $\vec{S}_k(G)$ that is closed under shifting such that an orientation $\mathcal{O}$ of $S_k(G)$ is a regular $\mathcal{F}$-tangle if and only if it is a regular $\hat{\mathcal{F}}$-tangle.*

In particular, any set $\mathcal{F}$ of subsets of $\vec{S}_k(G)$ can be transformed into a friendly set $\bar{\mathcal{F}} := \hat{\mathcal{F}} \cup \mathcal{P}_k \cup \{\{V(G), A\} : A \subseteq V(G) \wedge |A| < k\}$ of stars in $\vec{S}_k(G)$ so that an orientation of $S_k(G)$ is an $\bar{\mathcal{F}}$-tangle if and only if it is a regular profile and an $\mathcal{F}$-tangle of $S_k(G)$.

Moreover, it is easy to check that the set $\mathcal{T}_k^*$ that consists of all stars in $\mathcal{T}_k$ is a friendly set of stars. Further, Diestel and Oum have shown that the $\mathcal{T}_k^*$-tangles of $S_k(G)$ are precisely the k-tangles in G if $|G| \geqslant k$ [54, Lemma 4.2].

6.4 Refining inessential stars in finite graphs

The idea of refining tangle-distinguishing tree-decompositions has its origin in [66]. There, Erde proved that one can, under certain circumstances, refine the inessential nodes of tangle-distinguishing nested sets of separations of a graph so that all the inessential nodes of the refinement are 'small'.

Since then, this theorem has been generalized to 'abstract separation systems' [4]. Much of the theory that was introduced in [4], including the main result which generalizes Theorem 6.1.1, will play an important role in the proof of Theorem 17. Therefore, in this section we recap those definitions and lemmas.

Let $k \in \mathbb{N}$, let P be a k-profile in a graph G, and let $\vec{s} \in \vec{S}_k(G)$ be a separation. We say that $\vec{s}$ is *closely related to* P if $\vec{s} \in P$ and $\vec{r} \wedge \vec{s} \in \vec{S}_k(G)$ for every $\vec{r} \in P$. Further, we call a subset of $\vec{S}_k(G)$ *closely related to* P if all its elements are closely related to P.

A separation $\vec{r} \in P$ is *maximal* in P if it is a maximal element in P with respect to the partial order on P induced by $\vec{S}_k(G)$. An example of separations that are closely related to P are those which are maximal in P. In fact, even more is true:

Proposition 6.4.1. [4, Lemma 4.4] *Let $k \in \mathbb{N}$, and let P be a k-profile in some graph G. Further, let $Y \subseteq P$ be a nested set such that the inverse $\overleftarrow{y}$ of every $\vec{y} \in Y$ is closely related to some k-profile in G, and let $P_Y \subseteq P$ be the set of all separations in P that are nested with Y. Then every maximal separation in P_Y is closely related to P.*

The next observation, which we will use throughout this chapter, describes a sufficient condition for a separation to be closely related to some profile:

Proposition 6.4.2. [4, Proposition 3.3] *Let $k \in \mathbb{N}$, let G be a graph, and let $\vec{s} \in \vec{S}_k(G)$ be closely related to some k-profile P in G. Further, let $\vec{r} \in \vec{S}_k(G)$ with $\vec{r} \leqslant \vec{s}$, and suppose that $\vec{r} \wedge \vec{u} \in \vec{S}_k(G)$ for every $\vec{u} \leqslant \vec{s}$. Then $\vec{r}$ is closely related to P.*

The key tool in the proof of Theorem 6.1.1 is the following lemma which makes it possible to refine inessential stars that consist of separations whose inverses are each closely related to some $\mathcal{F}$-tangle. This lemma is taken from [4]; a slightly less general version was first proved in [66].

Lemma 6.4.3. [4, Lemma 3.5] *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}_k(G)$. Further, let $\sigma = \{\vec{s}_1, ..., \vec{s}_k\} \subseteq \vec{S}_k(G)$ be a star which is inessential for the set of all $\mathcal{F}$-tangles of $S_k(G)$ such that each $\overleftarrow{s}_i$ is closely related to some $\mathcal{F}$-tangle of $S_k(G)$. Then there is an $S_k(G)$-tree over $\mathcal{F} \cup \{\{\overleftarrow{s}_1\}, ..., \{\overleftarrow{s}_k\}\}$ in which each $\vec{s}_i$ appears as a leaf separation.*

Note that Erde [66] gave an example which shows that it is not possible to refine every inessential star, even if $\mathcal{F} = \mathcal{T}_k^*$.

6.5 Refining trees of tangles in finite graphs

Let G be a graph, $k \in \mathbb{N}$, and let $(\tilde{T}, \tilde{\mathcal{V}})$ be a tree-decomposition of G which distinguishes all its k-tangles so that every separation induced by an edge of $\tilde{T}$ efficiently distinguishes some pair of k-tangles. Then Theorem 6.1.1 asserts that $(\tilde{T}, \tilde{\mathcal{V}})$ extends to a tree-decomposition all whose inessential parts are small. This is done by refining the inessential parts of $(\tilde{T}, \tilde{\mathcal{V}})$ with further tree-decompositions that decompose those parts into smaller parts. In this section we show that we can refine the essential parts with further tree-decompositions as well so that we then obtain a single tree-decomposition of G whose parts are all relatively small. The main result of this section is then Theorem 17, which extends Theorem 6.1.1 and also strengthens a result of Erde ([66, Theorem 18]).

In order to make the proof of Theorem 17 more straightforward, and to reduce clutter caused by translating back and forth between the tree-decompositions and the nested sets of separations they induce, we first prove the following 'nested set version' of Theorem 17. Since every tree-decomposition of G induces a nested set of separations of G and vice versa, Theorem 17 then follows immediately.

Theorem 17'. *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}_k(G)$. Further, let $\tilde{N} \subseteq S_k(G)$ be a nested set that distinguishes all the $\mathcal{F}$-tangles of $S_k(G)$ so that every separation in N efficiently distinguishes a pair of $\mathcal{F}$-tangles of $S_k(G)$. Then there exist a nested set $N \subseteq S_k(G)$ with $\tilde{N} \subseteq N$ such that*

(i) *every inessential node of N is a star in $\mathcal{F}$;*

(ii) *the interior of every essential node of N is of smallest size among all exclusive stars contained in the $\mathcal{F}$-tangle living at that node.*

Note that the set $\mathcal{T}_k^*$ of all stars in $\mathcal{T}_k$ satisfies the premise of Theorem 17' (see Section 6.3), and hence we can apply Theorem 17' to the set of all k-tangles in G for some $k \in \mathbb{N}$. In particular, we can use Theorem 17' to refine the canonical nested set from [27], which efficiently distinguishes all the k-tangles in G:

Corollary 6.5.1. *Let G be a graph, and $k \in \mathbb{N}$. Then there exist nested sets $\tilde{N} \subseteq N \subseteq S_k(G)$ such that:*

(i) *$\tilde{N}$ is fixed under all automorphisms of G and efficiently distinguishes all k-tangles in G;*

(ii) *every inessential node of N is a star in $\mathcal{T}_k$;*

(iii) *the interior of every essential node of N is of smallest size among all exclusive stars contained in the k-tangle living at that node.* □

Before we prove Theorem 17', let us first emphasize the optimality of its condition (ii). Since $\tilde{N}$ distinguishes all $\mathcal{F}$-tangles of $S_k(G)$, the refinement N does so as well. Hence, each of its essential nodes is necessarily exclusive, and thus their interiors can only be of smallest size among all exclusive stars in the tangle they are home to. The following example shows that in general there may be non-exclusive stars in a tangle whose interiors contain fewer vertices than the interiors of the exclusive stars in that tangle. Thus, condition (ii) cannot be strengthened.

Example 6.5.2. Let G be the graph depicted in Figure 6.2a, which consists of five complete graphs K_{20} and one K_8 that are glued together as depicted in Figure 6.2a. Whenever a K_{20} is denoted inside a face of the drawing, we assume that the K_{20} contains the boundary vertices and edges of this face, but no other depicted vertices or edges. Moreover, the K_8 contains all vertices in the grey cycle.

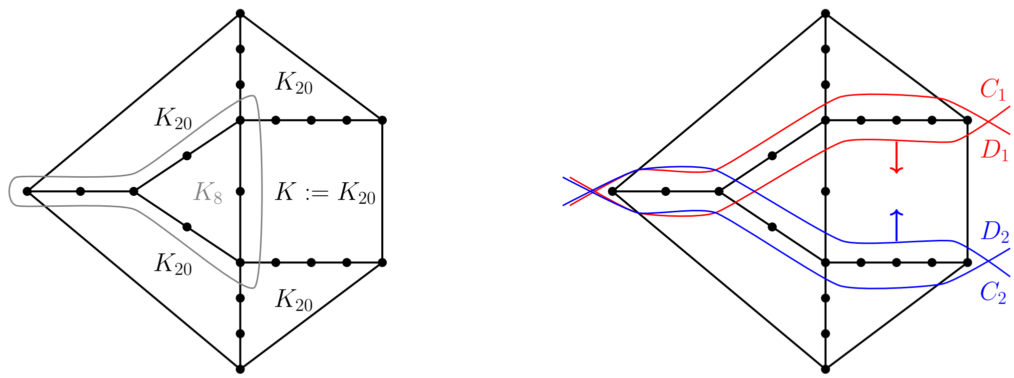

(a) (b)

Figure 6.2: Example 6.5.2

Let τ be the 10-tangle defined by orienting every separation towards that side which contains the clique K. Further, observe that every separation $\{A, B\} \in S_{10}(G)$ with $V(K_8) \subseteq A \cap B$ is of the form $\{A, V(G)\}$. Thus, we obtain a consistent orientation τ' of $S_{10}(G)$ by letting $(A, V(G)) \in \tau'$ for every set $A \subseteq V(G)$ of at most nine vertices, and $(A, B) \in \tau'$ if $V(K_8) \subseteq B$ for all other separations $\{A, B\} \in S_{10}(G)$. It is straightforward to check that τ' is a 10-tangle in G.

Moreover, it is easy to see that every star ϱ in τ whose interior is of smallest size among all stars in τ consists of $(C_1, D_1), (C_2, D_2)$ (see Figure 6.2b) and possibly some separations of the form $(A, V(G))$. But these separations are also contained in τ', and thus ϱ is not exclusive for the set of all 10-tangles in G.

Let us now turn to the proof of Theorem 17'. By Theorem 6.1.1, there exists a nested set $N' \subseteq S_k(G)$ which refines the inessential nodes of $\tilde{N}$ and satisfies (i). Therefore, we are left to refine the essential nodes of N' so that this refinement satisfies (i) and (ii). For this, we show the following lemma about refining essential stars in graphs, which will imply Theorem 17':

Lemma 6.5.3. *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}_k(G)$. Further, let $\sigma = \{\vec{s}_1, \dots, \vec{s}_n\} \subseteq \vec{S}_k(G)$ be a star which is home to a unique $\mathcal{F}$-tangle τ of $S_k(G)$, and suppose that each s_i efficiently distinguishes some pair of $\mathcal{F}$-tangles of $S_k(G)$.*

Then there exists a star $\sigma' \subseteq \tau$ whose interior is of smallest size among all exclusive stars *in τ, and an $S_k(G)$-tree over $\mathcal{F}' := \mathcal{F} \cup \{\sigma'\} \cup \{\{\overleftarrow{s}_1\}, \ldots, \{\overleftarrow{s}_n\}\}$ in which each $\vec{s}_i$ appears as a leaf separation.*

As τ has to live at a node of any $S_k(G)$-tree over $\mathcal{F}'$, and since σ' is the unique star in $\mathcal{F}'$ that is contained in τ, every $S_k(G)$-tree over $\mathcal{F}'$ has to contain a node which is associated with σ'. Thus, there can only exist an $S_k(G)$-tree over $\mathcal{F}'$ in which each $\vec{s}_i \in \sigma$ appears as a leaf separation if σ' is nested with σ. Our first step towards the proof of Lemma 6.5.3 will thus be to show that under all exclusive stars in τ whose interiors are of smallest size there is at least one star which is nested with σ.

Lemma 6.5.4. *Let $k \in \mathbb{N}$, let $\mathcal{P}$ be some set of k-profiles in a graph G, and let $P \in \mathcal{P}$. Further, let $\sigma \subseteq P$ be a star, and suppose that every separation in σ efficiently distinguishes some pair of k-profiles in $\mathcal{P}$. Then there exists a star $\varrho \subseteq P$ with $\sigma \leqslant \varrho$ whose interior is of smallest size among all stars in P that are exclusive for $\mathcal{P}$.*

Proof. Let ϱ be a star in P which is exclusive for $\mathcal{P}$ and which is nested with as many separations in σ as possible such that its interior is of smallest size among all stars in P that are exclusive for $\mathcal{P}$. We claim that ϱ is in fact nested with σ, which clearly implies the assertion. For this, suppose for a contradiction that there is a separation $(C,D) \in \sigma$ which is not nested with ϱ. By assumption, $\{C,D\}$ efficiently distinguishes some pair $Q \ni (D,C)$ and Q' of k-profiles in $\mathcal{P}$. Since Q is consistent and because $\varrho \not\subseteq Q$ as ϱ is P-exclusive for $\mathcal{P}$, there is a unique separation $(A,B) \in \varrho$ with $(B,A) \in Q$. Then $(C,D) \wedge (A,B)$ and $(C,D) \wedge (B',A')$, for every $(A',B') \in \varrho \smallsetminus \{(A,B)\}$, have order at least $|C \cap D|$ because $\{C,D\}$ distinguishes Q and Q' efficiently. By submodularity, it follows that $(C,D) \vee (B',A')$ has order $\leqslant |A' \cap B'| < k$, and further

$$|A \cap B| \geqslant |(C \cup A) \cap (D \cap B)| = |(A \cap B \cap D) \cup (B \cap C \cap D)| = |A \cap B \cap (D \smallsetminus C)| + |B \cap C \cap D|. \quad (6.1)$$

Thus, the star $\varrho' := \{(A \cup C, B \cap D)\} \cup \{(A' \cap D, B' \cup C) : (A',B') \in \varrho \smallsetminus \{(A,B)\}\}$ is a subset of $\vec{S}_k(G)$. In particular, since P is a profile, we have $\varrho' \subseteq P$. Moreover, ϱ' is P-exclusive for $\mathcal{P}$ since for every k-profile $P' \neq P$ in $\mathcal{P}$ either $(B,A) \in P'$ or $(D,C) \in P'$ and thus $(B \cap D, A \cup C) \in P'$, or there exists some $(A',B') \in \varrho$ such that $(B',A') \in P'$, and we then have $(B' \cup C, A' \cap D) \in P'$ because P' is a profile.

We claim that the interior X' of ϱ' has at most the size of the interior X of ϱ. This then contradicts the choice of ϱ since ϱ' is nested with $(C,D) \in \sigma$ by construction, and thus, by Lemma 2.1.1, ϱ' is nested with at least one separation in σ more than ϱ. Indeed, we have

$$X' = (B \cap D) \cap \bigcap_{(A',B') \in \varrho \smallsetminus \{(A,B)\}} (B' \cup C) = (X \cap D) \cup (B \cap C \cap D) = (X \cap (D \smallsetminus C)) \mathbin{\dot\cup} (B \cap C \cap D).$$

Since $\{C, D\}$ is a separation of G, we have $X = (X \cap C) \,\dot\cup\, (X \cap (D \setminus C))$, and thus

$$|X'| = |X| - |X \cap C| + |B \cap C \cap D| \leqslant |X| - |(A \cap B) \cap C| + |B \cap C \cap D|,$$

where we have used that $A \cap B \subseteq X$. Combining this with (6.1) yields that

$$|X'| \leqslant |X| - |(A \cap B) \cap C| + (|A \cap B| - |A \cap B \cap (D \setminus C)|) = |X|$$

since $|A \cap B| = |(A \cap B) \cap C| + |A \cap B \cap (D \setminus C)|$, which concludes the proof. □

For the proof of Lemma 6.5.3 we also need the next proposition, which follows from [66, Lemma 14 & 15].

Proposition 6.5.5. *Let G be a graph, $k \in \mathbb{N}$, and let $\vec{s} \in \vec{S}_k(G)$ be closely related to some k-profile P in G. Further, let $\vec{s'} \in P$ be any separation of minimal order such that $\vec{s} \leqslant \vec{s'}$. Then $\vec{s'}$ is closely related to P, and for every $\vec{r} \in P$ with $\vec{r} \leqslant \overleftarrow{s}$ we have $\vec{r} \wedge \overleftarrow{s'} \in \vec{S}_k(G)$. Moreover, if $\vec{r}$ is closely related to P, then $\vec{r} \wedge \overleftarrow{s'}$ is closely related to P.*

Proof. Let $\vec{x} \in P$ be a maximal separation in P such that $\vec{s'} \leqslant \vec{x}$. Applying [66, Lemma 14] to $\vec{s}$ and $\vec{x}$ yields that $\vec{s'} \wedge \vec{t} \in \vec{S}_k(G)$ for every $\vec{t} \leqslant \vec{x}$, which by Proposition 6.4.1 and 6.4.2 implies that $\vec{s'}$ is closely related to P. Further, applying [66, Lemma 15] to $\vec{s'}$ and $\vec{r}$ yields that $\vec{r} \wedge \overleftarrow{s'}$ is an element of $\vec{S}_k(G)$ and that $(\vec{r} \wedge \overleftarrow{s'}) \wedge \vec{t} \in \vec{S}_k(G)$ for every $\vec{t} \leqslant \vec{r}$, which by Proposition 6.4.2 implies that $\vec{r} \wedge \overleftarrow{s'}$ is closely related to P if $\vec{r}$ is closely related to P. □

We can now prove Lemma 6.5.3:

Proof of Lemma 6.5.3. Let $\mathcal{P}$ be the set of all $\mathcal{F}$-tangles of $S_k(G)$, and first assume that there is a star $\sigma' \subseteq \tau$ with $\sigma \leqslant \sigma'$ whose interior is of smallest size among all stars in τ that are exclusive for $\mathcal{P}$, and which has the further property that every separation in σ' is closely related to τ. Then the assertions follows. Indeed, by applying Lemma 6.4.3 to the inessential nodes of $N' := \{s : \vec{s} \in \sigma \cup \sigma'\}$ we obtain a refinement N of N' all whose inessential nodes are stars in $\mathcal{F}$. By Theorem 6.2.1, this set N induces an $S_k(G)$-tree over $\mathcal{F} \cup \{\sigma'\} \cup \{\{\overleftarrow{s}_1\}, \dots, \{\overleftarrow{s}_n\}\}$ in which each $\vec{s}_i$ appears as a leaf separation. By the choice of σ', this clearly implies the assertion.

Thus, it suffices to find a star σ' as above. For this, let $\sigma' \subseteq \tau$ be a star such that

(1) $\sigma \leqslant \sigma'$ and $|\operatorname{int}(\sigma')| = \min\{|\operatorname{int}(\varrho)| : \varrho \subseteq \tau \text{ and } \varrho \text{ is exclusive for } \mathcal{P}\}$, and

(2) the number of separations in σ' that are not closely related to τ is minimal among all stars in τ that satisfy (1).

Note that, by Lemma 6.5.4, there exists a star which satisfies (1). We claim that σ' is closely related to τ, which implies that σ' is the star for the argument above, and thus concludes the proof. For this, suppose for a contradiction that there is a separation $(C', D') \in \sigma'$ which is not closely related to τ, and let $(A, B) \in \tau$ be of minimal order such that $(C', D') \leqslant (A, B)$. We first show that there is such a separation $\{A, B\}$ which is nested with σ.

Claim 1. For every $\vec{r} \in \tau$ there exists a separation $\vec{r'} \in \tau$ of order $\leqslant |r|$ such that r' is nested with σ. Moreover, if $\vec{s} \in \tau$ is nested with σ such that $\vec{s} \leqslant \vec{r}$, then r' can be chosen so that $\vec{s} \leqslant \vec{r'}$.

Proof. Let $\vec{r'}(\geqslant \vec{s})$ be a separation of order $\leqslant |r|$ which is nested with as many separations in σ as possible. We claim that $\vec{r'}$ is nested with σ, which then clearly implies the assertion. For this, suppose for a contradiction that $\vec{r'}$ crosses some $\vec{t} \in \sigma$. By the assumption on σ, there exists a pair of k-profiles $Q \ni \overleftarrow{t}$ and $Q' \ni \vec{t}$ in G such that t distinguishes them efficiently. Since Q is a k-profile, it contains an orientation of r'; by symmetry we may assume that $\vec{r'} \in Q$. If $\vec{r'} \vee \overleftarrow{t} \in \vec{S}_k(G)$, then $\vec{r'} \vee \overleftarrow{t} \in Q$ and $\overleftarrow{r'} \wedge \vec{t} \in Q'$ as Q and Q' are k-profiles, and thus $\vec{r} \vee \overleftarrow{t}$ distinguishes Q and Q'. Since t efficiently distinguishes Q and Q', this implies that $\vec{r'} \vee \overleftarrow{t}$ has order $\geqslant |t|$, and hence, by submodularity, $\vec{r'} \wedge \overleftarrow{t}$ has order $\leqslant |r'|$. Then $\vec{r'} \wedge \overleftarrow{t} \in \tau$ because τ is a profile, which contradicts the choice of $\vec{r'}$ since $\vec{r'} \wedge \overleftarrow{t}$ is nested with one separation in σ more than $\vec{r'}$ by Lemma 2.1.1 (and still $\vec{s} \leqslant \vec{r'} \wedge \overleftarrow{t}$ as $\vec{s} \leqslant \vec{r}$ and s is nested with t). ■

By Claim 1, we may assume that $\{A, B\}$ is nested with σ. Then, by Proposition 6.5.5, the star

$$\sigma'' := \{(A, B)\} \cup \{(C \cap B, D \cup A) : (C, D) \in \sigma' \setminus \{(C', D')\}\}$$

is a subset of $\vec{S}_k(G)$, which by the consistency of τ implies that $\sigma'' \subseteq \tau$. Moreover, we have $\sigma \leqslant \sigma''$ as $\sigma \leqslant \sigma'$ and (A, B) is nested with σ. We claim that σ'' witnesses that σ' does not satisfy (2), which then contradicts the choice of σ' and thus concludes the proof. By Proposition 6.5.5, (A, B) is closely related to τ, and if $(C, D) \in \sigma'$ is closely related to τ, then $(C \cap B, D \cup A)$ is closely related to τ, too. Thus, there are fewer separations in σ'' that are not closely related to τ than in σ'. Therefore, we are left to show that σ'' satisfies (1). By the argument above, we have $\sigma \leqslant \sigma''$, so we only need to show that the interior X'' of σ'' has at most the size of the interior X' of σ'. By definition, we have

$$X'' = B \cap \bigcap_{(C,D) \in \sigma'} (D \cup A) = B \cap (X' \cup A) = (B \cap X') \cup (B \cap A) = (B \cap X') \,\dot{\cup}\, ((B \cap A) \setminus X').$$

Since $\{A, B\}$ is a separation of G, we have $X' = (X' \cap B) \,\dot{\cup}\, (X' \cap (A \setminus B))$ and thus

$$|X''| = |X'| - |(A \setminus B) \cap X'| + |(A \cap B) \setminus X'|.$$

So we are done if $|(A \cap B) \setminus X'| \leqslant |(A \setminus B) \cap X'|$. Set $\varrho := \sigma' \setminus \{(C', D')\}$. By the choice of (A, B), the separation $(A,B) \wedge \bigwedge_{(C,D)\in\varrho}(D,C)$ has at least order $|A\cap B|$, and thus

$$|A\cap B| \leqslant \Big|A\cap \bigcap_{(C,D)\in\varrho} D\cap \Big(B\cup \bigcup_{(C,D)\in\varrho} C\Big)\Big| = \Big|\Big(A\cap B\cap \bigcap_{(C,D)\in\varrho} D\Big)\cup\Big(A\cap \bigcap_{(C,D)\in\varrho} D\cap \bigcup_{(C,D)\in\varrho} C\Big)\Big|.$$

Since σ' is a star, we have $C\subseteq D'$ for every $(C,D)\in\varrho$, and hence $\bigcup_{(C,D)\in\varrho} C\subseteq D'$. Moreover, by the choice of (A,B), we have $B\subseteq D'$, and thus

$$|A\cap B| \leqslant |(A\cap B\cap X')\cup(A\cap X')| = |A\cap B\cap X'| + |(A\smallsetminus B)\cap X'|.$$

Combining this inequality with $|A\cap B| = |A\cap B\cap X'| + |(A\cap B)\smallsetminus X'|$ yields that

$$|(A\smallsetminus B)\cap X'| \geqslant |A\cap B| - |A\cap B\cap X'| = |A\cap B| - (|A\cap B| - |(A\cap B)\smallsetminus X'|) = |(A\cap B)\smallsetminus X'|,$$

which concludes the proof. □

With Lemma 6.5.3 at hand Theorem 17' follows immediately:

Proof of Theorem 17'. Apply Theorem 6.1.1 to $\tilde{N}$ to obtain a nested set $N'\subseteq S_k(G)$ with $\tilde{N}\subseteq N'$ which satisfies (i). Then, apply Lemma 6.5.3 to the essential nodes of N'. □

Proof of Theorem 17. Apply Theorem 17' to the set of separations induced by $(\tilde{T},\tilde{\mathcal{V}})$. □

Recall that the essential parts of the tree-decomposition from Theorem 17 are as small as possible so that they are still home to their tangle. Thus, this refinement is optimal in that it decreases the size of the essential parts as much as possible.

Now, one may ask what else could be said about the essential parts of that refinement besides their size being as small as possible. Our main reason for decreasing the size of the parts was to more precisely exhibit where the tangles are located in the graph, so it seems natural to ask whether the vertices in the essential parts in some sense 'belong' to the tangle living in that part, and additionally whether all vertices that 'belong' to a tangle are contained in the part which is home to that tangle.

In general, this question cannot be answered because tangles are, by nature, a somewhat fuzzy object, which is, at least in general, not related to a specific set of vertices. So usually it cannot be determined at all whether a certain vertex 'belongs' to a tangle or not. Thus, in general, we also cannot evaluate whether the essential parts contain precisely those vertices which 'belong' to that tangle.

However, in what follows, we present two properties of the refined tree-decomposition from Theorem 17 that we believe validate that its essential parts, at least in many cases, 'closely correspond' to the tangle they are home to. For this, let G be some graph, and let $(T,\mathcal{V})$ be the tree-decomposition which one obtains by applying Theorem 17 to any tree-decomposition of G which efficiently distinguishes all the regular k-profiles in G. (Such tree-decompositions exist by [47, Theorem 3.6].)

Then each part of $(T, \mathcal{V})$ of size $> 3k-3$ is home to a k-tangle, which it then 'witnesses' and 'induces'. More precisely, let us say that a subgraph H of G *witnesses* a k-tangle τ if for every subset $\{(A_i, B_i) : i \in [3]\}$ of τ of at most three elements we have $\bigcup_{i \in [3]} G[A_i \cap V(H)] \neq H$. Further, a set U of vertices of G *induces* a k-profile τ if for every separation $(A, B) \in \tau$ we have $|A \cap U| < |B \cap U|$.

If a bag $V_t \in \mathcal{V}$ has size $> 3k-3$, then σ_t cannot be a star in $\mathcal{P}_{S_k(G)} \subseteq \mathcal{T}_k$, and thus by (i), V_t is home to some regular k-profile τ in G. By (ii) and the following lemma, it then follows that $|A \cap V_t| < k$ for every $(A, B) \in \tau$.

Lemma 6.5.6. [46, Proof of Theorem 12] *Let $k \in \mathbb{N}$, let $\mathcal{P}$ be some set of k-profiles in a graph G, and let $P \in \mathcal{P}$. Further, let $\sigma \subseteq P$ be a star whose interior is of smallest size among all stars in P that are exclusive for $\mathcal{P}$. Then $|A \cap \operatorname{int}(\sigma)| < k$ for all $(A, B) \in P$.*

Note that in [46] the statement of Lemma 6.5.6 is shown for stars in P whose interior is of smallest size among all stars in P. However, it is easy to see that the same proof also yields the statement of Lemma 6.5.6, where the interior of σ is only of smallest size among all *exclusive* stars in P.

By Lemma 6.5.6 we have $|A \cap V_t| < k \leqslant (3k-2) - (k-1) \leqslant |V_t| - |A \cap V_t| = |(B \setminus A) \cap V_t| \leqslant |B \cap V_t|$, and hence V_t induces τ. Moreover, for every subset $\{(A_i, B_i) : i \in [3]\} \subseteq \tau$, we have $\bigcup_{i \in [3]} G[A_i \cap V_t] \neq G[V_t]$ since $|V_t \cap \bigcap_{i \in [3]} (B_i \setminus A_i)| \geqslant |V_t| - |A_1| - |A_2| - |A_3| \geqslant (3k-2) - 3(k-1) = 1$. Hence, τ is in fact a k-tangle in G as witnessed by $G[V_t]$.

Therefore, every large enough part of $(T, \mathcal{V})$ induces a k-tangle. Conversely, if a k-tangle is induced by some large and highly connected substructure of G, does that substructure appear as a part of $(T, \mathcal{V})$? In general, this will not be the case even if a k-tangle is induced by a large clique, since even a clique must not be equal to a part of any tree-decomposition of adhesion $< k$ [28]. However, if a k-tangle is induced by a 'k-block', then the k-block will be equal to a bag of $(T, \mathcal{V})$ if there exists any tree-decomposition of adhesion $< k$ at all that contains the k-block as a bag.

For some $k \in \mathbb{N}$, a *k-block* in a graph G is a maximal set b of at least k vertices such that no two vertices $v, w \in b$ can be separated in G by fewer than k vertices. It is straight forward to check that every k-block *induces* a regular k-profile by orienting $\{A, B\} \in S_k(G)$ as (A, B) if and only if $b \subseteq B$. A k-block b in G is *separable* if it is the interior of some star in $S_k(G)$, i.e. if there exists a star $\sigma \subseteq \vec{S}_k(G)$ such that $\operatorname{int}(\sigma) = b$.

Now suppose that b is a separable k-block in G, and let τ be the k-profile induced by b. Then the star $\varrho := \{(V(C) \cup N(C), V(G) \setminus V(C)) : C \in \mathcal{C}(G-b)\}$ is in $\vec{S}_k(G)$ [30, Lemma 4.1], and it is easy to see that $\operatorname{int}(\varrho) = b$ and that τ is the unique k-profile which contains ϱ. Moreover, we have

$b \subseteq B$ for every $(A, B) \in \tau$, and thus $b \subseteq \text{int}(\sigma)$ for every star $\sigma \subseteq \tau$. Hence, if b is separable, then every star in τ whose interior is of smallest size among all exclusive stars in τ will be equal to b. It follows by condition (ii) that the essential part of the tree-decomposition $(T, \mathcal{V})$ from Theorem 17 which is home to τ is equal to b. Thus, every separable k-block in G appears as a bag of $(T, \mathcal{V})$. All in all, we have the following theorem:

Theorem 6.5.7. *Let G be a graph and $k \in \mathbb{N}$. Then there exists a tree-decomposition $(T, \mathcal{V})$ of G of adhesion $< k$ which has the following properties:*

(i) *$(T, \mathcal{V})$ efficiently distinguishes all regular k-profiles in G;*

(ii) *each bag of $(T, \mathcal{V})$ of size $> 3k - 3$ is home to a k-tangle, which it then witnesses and induces;*

(iii) *every separable k-block in G appears as a bag of $(T, \mathcal{V})$.* □

We remark that the existence of tree-decompositions satisfying (i) and (iii) is already known due to Carmesin and Gollin [30]. What is new is property (ii): that all the inessential bags of the tree-decomposition $(T, \mathcal{V})$ are small, while all of its bags that are not small induce k-tangles.

7 Refining tree-decompositions so that they display the k-blocks

Carmesin and Gollin proved that every finite graph has a canonical tree-decomposition $(T, \mathcal{V})$ of adhesion less than k that efficiently distinguishes every two distinct k-profiles, and which has the further property that every separable k-block is equal to the unique part of $(T, \mathcal{V})$ in which it is contained.

We give a shorter proof of this result by showing that such a tree-decomposition can in fact be obtained from any canonical tight tree-decomposition of adhesion less than k. For this, we decompose the parts of such a tree-decomposition by further tree-decompositions. As an application, we also obtain a generalization of Carmesin and Gollin's result to locally finite graphs.

This chapter is based on [6].

7.1 Introduction

All graphs in this chapter may be infinite, unless otherwise stated.

A *k-block* in a graph G, for some $k \in \mathbb{N}$, is a maximal set of at least k vertices no two of which can be separated in G by removing fewer than k other vertices. For large enough k, the k-blocks of a graph are examples of highly connected substructures.

Another example of such a substructure, one which also indicates high local connectivity but is of a more fuzzy kind than blocks, is that of a tangle. Tangles were introduced by Robertson and Seymour in [114]. Formally, a *k-tangle* in a graph G is a consistent orientation of all the separations $\{A, B\}$ of G of order less than k, as (A, B) say, such that no three such oriented separations together cover the whole graph by the subgraphs induced on their 'small sides' A.

Since k-blocks cannot be separated by deleting fewer than k vertices, they induce an orientation of every separation of order less than k: towards that side which contains the k-block. Although these orientations are consistent in that they all point towards the same k-block, they need not be tangles if the k-block is too small. But they are *k-profiles*: a common generalization of tangles and blocks, in that every k-tangle is a k-profile, and every k-block induces a k-profile in the way described above.

Robertson and Seymour [114] proved that every finite graph G has a tree-decomposition $(T, \mathcal{V})$ of adhesion less than k, with $\mathcal{V} = (V_t)_{t \in T}$ say, that *distinguishes* all its k-tangles: for every pair of k-tangles there exists an edge of T which induces a separation that is oriented differently by these two k-tangles. For this, recall that G reflects the separation properties of T: similar as the deletion of an edge $e = \{t_1, t_2\} \in E(T)$ separates T into two subtrees $T_1 \ni t_1$ and $T_2 \ni t_2$, the sets $U_i := \bigcup_{t \in V(T_i)} V_t$, for $i = 1, 2$, form a separation $\{U_1, U_2\}$ of G. In fact, the tree-decomposition of Robertson and Seymour even distinguishes the k-tangles in G *efficiently*: If two k-tangles are distinguished by a separation $\{A, B\}$ of G of *order* $|A \cap B| \leqslant \ell \in \mathbb{N}$, then they are also distinguished by a separation of order $\leqslant \ell$ that is induced by an edge of T.

Carmesin, Diestel, Hamann and Hundertmark [27] generalized this result by showing that every finite graph has a tree-decomposition of adhesion less than k that efficiently distinguishes all its regular k-profiles. In addition, the tree-decomposition they constructed has the additional property that it is *canonical*: its construction commutes with all isomorphisms $G \to G'$.

Carmesin and Gollin [30] improved this result even further and showed that every finite graph G admits a tree-decomposition $(T, \mathcal{V})$ as above which additionally displays the structure of the k-blocks in G, in that every k-block in G which can be isolated by any tree-decomposition at all[1] appears as a bag of $(T, \mathcal{V})$:

Theorem 7.1.1 ([30, Theorem 1])**.** *Every finite graph G has a canonical tree-decomposition $(T, \mathcal{V})$ of adhesion less than k that efficiently distinguishes every two distinct regular k-profiles, and which has the further property that every separable k-block is equal to the unique bag of $(T, \mathcal{V})$ that contains it.*

They also proved the following related result:

Theorem 7.1.2 ([30, Theorem 2])**.** *Every finite graph G has a canonical tree-decomposition $(T, \mathcal{V})$ that efficiently distinguishes every two distinct maximal robust profiles, and which has the further property that every separable block inducing a maximal robust[2] profile is equal to the unique bag of $(T, \mathcal{V})$ that contains it.*

In this chapter we give a short proof of Theorems 7.1.1 and 7.1.2 by showing the following more general result, which allows us to decompose the parts of a *given* tree-decomposition further, so that the resulting tree-decomposition displays the structure of the blocks:

Theorem 18. *Let G be any graph, and let $\mathcal{B}$ be a set of separable blocks in G. Suppose that G has a tight tree-decomposition $(\tilde{T}, \tilde{\mathcal{V}})$ that distinguishes all the blocks in $\mathcal{B}$. Then there exists a tree-decomposition $(T, \mathcal{V})$ that refines $(\tilde{T}, \tilde{\mathcal{V}})$ and is such that every block in $\mathcal{B}$ is equal to the unique bag of $(T, \mathcal{V})$ that contains it. Moreover, $(T, \mathcal{V})$ is canonical if $\mathcal{B}$ and $(\tilde{T}, \tilde{\mathcal{V}})$ are canonical.*

[1] We call such k-blocks *separable*.

[2] For some $n \in \mathbb{N}$, a profile in G is *n-robust* if for every $(A, B) \in P$ and every $\{C, D\} \in S_n(G)$ the following holds: if both $(A \cup C, B \cap D)$ and $(A \cup D, B \cap C)$ have order less than $|A \cap B|$, then one of them is contained in P. Clearly, every k-profile is k-robust. A profile is *robust* if it is n-robust for every $n \in \mathbb{N}$.

Here, an (unrooted) tree-decomposition $(T, \mathcal{V})$ is *tight* if every separation induced by an edge of T is tight.

For the proof of Theorem 7.1.1 Carmesin and Gollin gave one particular algorithm to construct a canonical tree-decomposition which distinguishes all k-profiles efficiently and which displays all separable k-blocks. However, there are a number of different algorithms to construct canonical tree-decompositions that distinguish all the k-profiles in a graph [27, 34, 63]. By Theorem 18, we can now choose whichever algorithm we like to construct an initial tree-decomposition, perhaps in order to have some control over the structure of those parts that do not contain any blocks, and we can still conclude that the tree-decomposition extends to one which additionally displays all separable k-blocks.

Moreover, Theorem 18 also applies to infinite graphs. Carmesin, Hamann and Miraftab [31] and Elbracht, Kneip and Teegen [64] showed that every locally finite graph has a canonical tree-decomposition that distinguishes all its k-profiles. Moreover, Jacobs and Knappe [90] showed that every locally finite graph without half-grid minor has a canonical tree-decomposition that distinguishes all its maximal robust profiles. Applying Theorem 18 to these tree-decompositions yields the following generalisations of Theorems 7.1.1 and 7.1.2:

Theorem 19. *Every locally finite graph G has a canonical tree-decomposition $(T, \mathcal{V})$ of adhesion less than k that efficiently distinguishes every two distinct k-profiles, and which has the further property that every separable k-block is equal to the unique bag of $(T, \mathcal{V})$ that contains it.*

Theorem 20. *Every locally finite graph G without half-grid minor has a canonical tree-decomposition $(T, \mathcal{V})$ that efficiently distinguishes every two distinct maximal robust profiles, and which has the further property that every separable block inducing a maximal robust profile is equal to the unique bag of $(T, \mathcal{V})$ that contains it.*

7.2 Blocks

For some $k \in \mathbb{N}$, a *k-block* in a graph G is a maximal set b of at least k vertices such that no two vertices $v, w \in b$ can be separated in G by removing fewer than k vertices other than v, w. A set $b \subseteq V(G)$ is a *block* if it is a k-block for some $k \in \mathbb{N}$.

It is straightforward to check that every k-block in G *induces* a regular k-profile in G by orienting $\{A, B\} \in S_k(G)$ as (A, B) if and only if $b \subseteq B$. Moreover, distinct k-blocks induce distinct k-profiles. We say that a tree-decomposition of G *(efficiently) distinguishes* two blocks in G if it (efficiently) distinguishes their induced profiles.

A k-block b in G is *separable* if it is the interior of some star in $S_k(G)$, i.e. if there exists a star $\sigma \subseteq \vec{S}_k(G)$ such that $\operatorname{int}(\sigma) = b$. We need the following equivalent characterization of separable

k-blocks:

Lemma 7.2.1 ([30, Lemma 4.1])**.** *Let b be a k-block in a graph G. Then b is separable if and only if $|N_G(C)| < k$ for all components C of $G - b$.*

7.3 Refining stars whose interior contains a block

In this section we prove Theorem 18 and then derive Theorems 7.1.1 and 7.1.2 and Theorems 19 and 20 from it. For this, we first show the following lemma. It asserts that given a part of a tight tree-decomposition which contains a separable k-block, then we can further decompose that part in a star-like way so that the central bag of that decomposition is equal to the k-block:

Lemma 7.3.1. *Let $k \in \mathbb{N}$, and let b be a separable k-block in a graph G. Further, let $\sigma \subseteq \vec{S}_{\aleph_0}(G)$ be a star of tight separations such that $b \subseteq \operatorname{int}(\sigma)$. Then there exists a star $\varrho_b^\sigma \subseteq \vec{S}_k(G)$ such that $\sigma \leqslant \varrho_b^\sigma$ and $\operatorname{int}(\varrho_b^\sigma) = b$. Moreover, ϱ_b^σ can be chosen so that if $\varphi : G \to G'$ is an isomorphism, then $\varphi(\varrho_b^\sigma) = \varrho_{\varphi(b)}^{\varphi(\sigma)}$.*

Proof. Let $\mathcal{C}(G-b)$ be the set of all components of $G-b$ and let

$$\mathcal{C}' := \{C \in \mathcal{C}(G-b) : V(C) \cap B \neq \emptyset \text{ for all } (A,B) \in \sigma\}$$

be the set of all components of $G-b$ that are not completely contained in the strict small side $G[A \smallsetminus B]$ of some $(A,B) \in \sigma$. Further, for every component $C \in \mathcal{C}'$, we define a star $\sigma_C := \{(A,B) \in \sigma : A \cap V(C) \neq \emptyset\}$ and separation

$$(X_C, Y_C) := \Big(V(C) \cup N_G(C) \cup \bigcup_{(A,B)\in\sigma_C} A \smallsetminus B,\ V(G) \smallsetminus \Big(V(C) \cup \bigcup_{(A,B)\in\sigma_C} A \smallsetminus B\Big)\Big)$$

(see Figure 7.1). Note that $N_G(C) \subseteq b \subseteq B$ for all $(A,B) \in \sigma$, and thus $X_C \cap Y_C = N_G(C)$.

Let us first show that $\{X_C, Y_C\}$ is a separation of G. Clearly, its sides cover $V(G)$, so it remains to prove that $N_G(X_C \smallsetminus Y_C) \subseteq X_C$. By the definition of $\{X_C, Y_C\}$, this is the case if $N_G(A \smallsetminus B) \subseteq V(C) \cup N_G(C)$ for all $(A,B) \in \sigma_C$. So let $(A,B) \in \sigma_C$ be given. Then $V(C) \cap A \neq \emptyset$, and moreover $V(C) \cap B \neq \emptyset$ because $C \in \mathcal{C}'$, which implies that $V(C) \cap (A \cap B) \neq \emptyset$ as C is connected. Since $\{A,B\}$ is tight, there is a component $\tilde{C}_A \subseteq G[A \smallsetminus B]$ of $G - (A \cap B)$ such that $N_G(\tilde{C}_A) = A \cap B$. In particular, there is an edge between $\tilde{C}_A$ and C. As C is a component of $G - b$ and $\tilde{C}_A$ is connected and disjoint from b by assumption, this implies that $\tilde{C}_A \subseteq C$. Hence, $N_G(A \smallsetminus B) = A \cap B = N_G(\tilde{C}_A) \subseteq N_G(C) \cup V(C)$, and thus $\{X_C, Y_C\}$ is a separation of G. In particular, $\{X_C, Y_C\}$ has order $|N_G(C)|$.

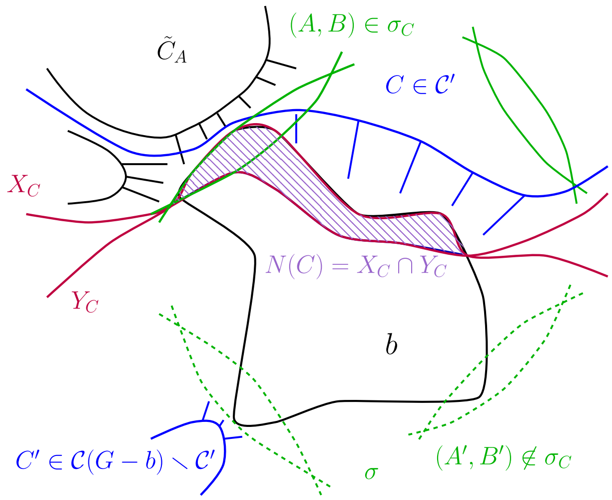

FIGURE 7.1: A component $C \in \mathcal{C}'$ and the separation (X_C, Y_C). The separations $(A, B) \in \sigma_C$ are indicated with solid lines, the separations $(A', B') \in \sigma \smallsetminus \sigma_C$ are indicated with dashed lines. The component $\tilde{C}_A$ of $G - (A \cap B)$ is contained in C.

Since $|N_G(C)| < k$ by Lemma 7.2.1, this implies that $\{X_C, Y_C\} \in S_k(G)$. Now set

$$\varrho_b^\sigma := \{(X_C, Y_C) : C \in \mathcal{C}'\} \cup \{(A, B) \in \sigma : A \cap B \subseteq b\}.$$

Since every $(A, B) \in \sigma$ is tight, Lemma 7.2.1 implies that every $(A, B) \in \sigma$ with $A \cap B \subseteq b$ has order less than k. As also every $\{X_C, Y_C\}$ for $C \in \mathcal{C}'$ is of order less than k by the argument above, it follows that $\varrho_b^\sigma \subseteq \vec{S}_k(G)$. We claim that ϱ_b^σ is as desired.

First, we prove that ϱ_b^σ is a star. For this, we show that $(X_C, Y_C) \leqslant (Y_{C'}, X_{C'})$ for distinct $C \neq C' \in \mathcal{C}'$ and that $(X_C, Y_C) \leqslant (B, A)$ for all $C \in \mathcal{C}'$ and $(A, B) \in \sigma$ with $A \cap B \subseteq b$. Since σ is a star itself, this then concludes the proof that ϱ_b^σ is a star.

We first show the former. To this end, let two distinct components $C \neq C' \in \mathcal{C}'$ be given. Then they cannot both meet the small side A of the same separation $(A, B) \in \sigma$, as otherwise $\tilde{C}_A \subseteq C \cap C'$ by the argument above, and then $C = C'$. Therefore, $\sigma_C \cap \sigma_{C'} = \emptyset$, and thus $X_C \smallsetminus Y_C = V(C) \cup \bigcup_{(A,B) \in \sigma_C} A \smallsetminus B$ and $X_{C'} \smallsetminus Y_{C'} = V(C') \cup \bigcup_{(A,B) \in \sigma_{C'}} A \smallsetminus B$ are disjoint. Hence $(X_C, Y_C) \leqslant (Y_{C'}, X_{C'})$.

Now let $C \in \mathcal{C}'$ and $(A, B) \in \sigma$ with $A \cap B \subseteq b$ be given. Then $V(C) \cap B \neq \emptyset$ by the definition of $\mathcal{C}'$, which implies that $V(C) \subseteq B \smallsetminus A$ as C is connected and avoids $b \supseteq A \cap B$. Thus, $(X_C, Y_C) \leqslant (B, A)$.

Second, we show that $\sigma \leqslant \varrho_b^\sigma$. For this, let $(A, B) \in \sigma$ be given. We need to find a separation

$(A', B') \in \varrho_b^\sigma$ with $(A,B) \leqslant (A',B')$. If $A \cap B \subseteq b$, then $(A,B) \in \varrho_b^\sigma$ is as desired. Otherwise, $A \cap B$ meets a component C of $G - b$, which then has to lie in $\mathcal{C}'$ since σ is a star. In particular, $(A,B) \in \sigma_C$ and $(A,B) \leqslant (X_C, Y_C)$ by the definition of $\{X_C, Y_C\}$. Since $(X_C, Y_C) \in \varrho_b^\sigma$ as $C \in \mathcal{C}'$, this completes the proof that $\sigma \leqslant \varrho_b^\sigma$.

Next, we show that $\mathrm{int}(\varrho_b^\sigma) = b$. For this, we first observe that b is disjoint from every component C of $G - b$ and from every strict small side $A \smallsetminus B$ of every $(A,B) \in \sigma$ by assumption. By the definition of ϱ_b^σ, this implies that $b \subseteq \mathrm{int}(\varrho_b^\sigma)$. Moreover, every vertex $v \in V(G) \smallsetminus b$ is contained in a component C of $G - b$. If $C \in \mathcal{C}'$, then v lies in the strict small side $X_C \smallsetminus Y_C$ of (X_C, Y_C) by definition, and hence $v \notin \mathrm{int}(\varrho_b^\sigma)$. Otherwise, there is a separation $(A,B) \in \sigma$ such that $v \in V(C) \subseteq A \smallsetminus B$. Since $\sigma \leqslant \varrho_b^\sigma$ as show earlier, this implies that $v \notin \mathrm{int}(\varrho_b^\sigma)$. Therefore, $\mathrm{int}(\varrho_b^\sigma) \subseteq b$.

We are left to show the 'moreover'-part. For this, let $\varphi : G \to G'$ be an isomorphism. We show that $(\varphi(X_C), \varphi(Y_C)) = (X_{\varphi(C)}, Y_{\varphi(C)})$, which clearly implies the assertion. For this, note that $\varphi(b)$ is a k-block in G', that $\varphi(C)$ is a component of $G' - \varphi(b)$, and that $\varphi(\sigma)$ is a star of tight separations in $\vec{S}_k(G')$ with $\varphi(b) \subseteq \mathrm{int}(\varphi(\sigma))$. Thus, $\{X_{\varphi(C)}, Y_{\varphi(C)}\}$ is defined. Moreover, if $V(C) \cap A \neq \emptyset$ for some $(A,B) \in \sigma$, then $\varphi(V(C)) \cap \varphi(A) \neq \emptyset$. Hence, $\varphi(\sigma)_{\varphi(C)} = \varphi(\sigma_C)$, which, by the definition of $\{X_C, Y_C\}$ and because φ is an isomorphism, implies that $(\varphi(X_C), \varphi(Y_C)) = (X_{\varphi(C)}, Y_{\varphi(C)})$. $\square$

We can now prove Theorem 18.

Proof of Theorem 18. Applying Lemma 7.3.1 to every star σ_t that is associated with a node $t \in \tilde{T}$ such that $\tilde{V}_t$ contains a block b in $\mathcal{B}$ yields stars $\varrho_b^{\sigma_t}$ with $\sigma_t \leqslant \varrho_b^{\sigma_t}$ and $\mathrm{int}(\varrho_b^{\sigma_t}) = b$.

We now construct the desired tree-decomposition $(T, \mathcal{V})$. For this, we first define tree-decompositions $(T^t, \mathcal{V}^t)$ of the parts $G[\tilde{V}_t]$ as follows. If $\tilde{V}_t$ does not contain a block from $\mathcal{B}$, then we set $T^t := (\{t\}, \emptyset)$ and $V_t^t := \tilde{V}_t$. Otherwise, if b is the (unique) block from $\mathcal{B}$ that is contained in $\tilde{V}_t$, then we let T^t be the star with centre t and with $|\varrho_b^{\sigma_t}|$ many leaves $u_{(A,B)}$, one for each $(A,B) \in \varrho_b^{\sigma_t}$. Further, we set $V_t^t := b = \mathrm{int}(\varrho_b^{\sigma_t})$ and $V_{u_{(A,B)}}^t := A \cap \tilde{V}_t$ for all $(A,B) \in \varrho_b^{\sigma_t}$. It is straightforward to check that $(T^t, \mathcal{V}^t)$ is a tree-decomposition of $G[V_t]$.

We then let T be the tree obtained from the disjoint union over the trees T^t by adding for every edge $\{t_1, t_2\} \in \tilde{T}$ the edge $\{v_1, v_2\}$ where $v_i = t_i$ if $T^{t_i} = (\{t_i\}, \emptyset)$ and $v_i := u_{(A,B)}$ where $(A,B) \in \varrho_b^{\sigma_{t_i}}$ is the unique separation with $\vec{s}_{(t_{3-i}, t_i)} = (U_{3-i}, U_i) \leqslant (A,B)$ otherwise. Note that such a separation exists because $(U_{3-i}, U_i) \in \sigma_{t_i} \leqslant \varrho_b^{\sigma_{t_i}}$. Further, we set $V_s := V_s^t$ for all $s \in T$ where t is the unique node of $\tilde{T}$ such that $s \in T^t$. It is straightforward to check that $(T, \mathcal{V})$ is a tree-decomposition of G. Moreover, by construction, for every edge $\{t_1, t_2\} \in \tilde{T}$, the edge $\{v_1, v_2\}$ of T induces the same separation of G, i.e. $s_{(t_1,t_2)} = s_{(v_1,v_2)}$, so $(T, \mathcal{V})$ refines $(\tilde{T}, \tilde{\mathcal{V}})$. Finally, by the 'moreover'-part of Lemma 7.3.1, $(T, \mathcal{V})$ is canonical if $(\tilde{T}, \tilde{\mathcal{V}})$ is canonical. Hence, $(T, \mathcal{V})$ is as desired. $\square$

Proof of Theorem 19. By [31, 64], G admits a canonical tree-decomposition $(\tilde{T}, \tilde{\mathcal{V}})$ that efficiently distinguishes all the k-profiles in G.[3] In particular, $(\tilde{T}, \tilde{\mathcal{V}})$ is tight by Lemma 2.2.1. Moreover, since every k-block induces a k-profile, and since distinct k-blocks induce distinct k-profiles, $(\tilde{T}, \tilde{\mathcal{V}})$ distinguishes all k-blocks in G. Apply Theorem 18 to $(\tilde{T}, \tilde{\mathcal{V}})$ and the set $\mathcal{B}$ of all separable k-blocks in G. □

Proof of Theorem 20. By [64, Theorem 6.6] and [90, Theorem 5.4] (see also [90, Theorem 1 and the comment after the proof of Theorem 1]), G admits a canonical tree-decomposition $(\tilde{T}, \tilde{\mathcal{V}})$ that efficiently distinguishes all its maximal robust profiles. In particular, $(\tilde{T}, \tilde{\mathcal{V}})$ is tight by Lemma 2.2.1. Apply Theorem 18 to $(\tilde{T}, \tilde{\mathcal{V}})$ and the set $\mathcal{B}$ of all separable blocks in G that induce a maximal robust profile. □

Proof of Theorem 7.1.1. Apply Theorem 19. □

Proof of Theorem 7.1.2. Apply Theorem 20. □

[3] In [31, Theorem 7.3] Carmesin, Hamann and Miraftab prove that G admits a canonical tree-decomposition that distinguishes efficiently all its *robust* k-profiles. In [64, Theorem 6.6] Elbracht, Kneip and Teegen give an independent proof of this result. Here, it can be seen from the proof that one may omit 'robust', as the authors remark in the preliminary section.

8 On vertex sets inducing tangles

Diestel, Hundertmark and Lemanczyk asked whether every k-tangle in a graph is induced by a set of vertices by majority vote. We reduce their question to graphs whose size is bounded by a function in k. Additionally, we show that if for any fixed k this problem has a positive answer, then every k-tangle is induced by a vertex set whose size is bounded in k. More generally, we prove for all k that every k-tangle in a graph G is induced by a weight function $V(G) \to \mathbb{N}$ whose total weight is bounded in k. As the key step of our proofs, we show that any given k-tangle in a graph G is the lift of a k-tangle in some topological minor of G whose size is bounded in k.

This chapter is based on [7] and joint work with Hanno von Bergen, Raphael W. Jacobs, Paul Knappe and Paul Wollan.

8.1 Introduction

All graphs in this chapter are finite.

8.1.1 Vertex sets inducing tangles

Tangles are an abstract notion of 'clusters' in graphs that originates in the theory of graph minors developed by Robertson and Seymour [110]. They allow for a unified treatment of various concrete highly cohesive substructures in graphs, such as large clique or grid minors. Tangles describe these clusters in a graph indirectly. Instead of describing what the cluster is composed of, they describe its position, in that they orient the low-order separations of the graph towards it. Intuitively, a concrete cluster orients all low-order separations by majority vote, that is the cluster orients such a separation $\{A, B\}$ towards its side, A or B, which contains most of the cluster. Such a side exists; otherwise, the cluster would be separated by few vertices, which contradicts its high cohesion.

The orientations of the low-order separations induced by concrete clusters are 'consistent', in that they all point to the cluster. Robertson and Seymour's key innovation was to distil from this an abstract notion of 'consistency' which leads to the notion of tangle [114]: Formally, a *k-tangle* τ in a graph G is an orientation of the separations of G of order $< k$ such that there do not exist three separations $(A_i, B_i) \in \tau$ such that the union of the small sides $G[A_i]$ covers the whole graph. We refer to k as the *order* of τ, and we will denote a tangle of unspecified order as simply a *tangle in G*.

While every concrete cluster induces a tangle by majority vote, is the converse also true in that all tangles stem from concrete clusters in this way? Without a precise definition of concrete cluster, it seems difficult to answer this question. However, the question remains interesting if we consider arbitrary vertex sets instead of concrete clusters: Is every tangle at least induced by the majority vote of some set of vertices? Diestel, Hundertmark and Lemanczyk [47, Section 7] formalised this problem as follows. A set X of vertices of a graph G *induces* a tangle τ in G if for every separation $(A, B) \in \tau$ we have $|X \cap A| < |X \cap B|$. For example, the vertex set of a complete subgraph induces a tangle in this way.

Problem 8.1.1. [47] *Is every tangle in a graph G induced by some set $X \subseteq V(G)$?*

In what follows we will often consider Problem 8.1.1 for all tangles of some fixed order $k \in \mathbb{N}$ and then say for short: Problem 8.1.1 for k.

We remark that Problem 8.1.1 is already answered in the affirmative for $k \leqslant 3$: Such sets X inducing k-tangles exist for $k \leqslant 2$ due to the well-known correspondence of these tangles to components and blocks, respectively (cf. [114, (2.6)]). For $k = 3$, Elbracht [61, Theorem 1.2] proved Problem 8.1.1 directly. Independently, Grohe [78, Theorem 4.8] proved a direct correspondence between the 3-tangles and the 'proper triconnected components' of a graph, which yields another proof of Problem 8.1.1 for $k = 3$. Similarly, it should be possible to derive Problem 8.1.1 for $k = 4$ from the recent characterisation of 4-tangles in terms of 'internal 4-connectedness' by Carmesin and Kurkofka [33, Theorem 1]. Besides these results for small k, Diestel, Elbracht and Jacobs [46, Theorem 12] showed that Problem 8.1.1 is true for every k-tangle τ in a graph G which *extends* to a $2k$-tangle τ' in G, that is, $\tau \subseteq \tau'$.

For general k, Elbracht, Kneip and Teegen [62, Theorem 2] made substantial progress towards Problem 8.1.1 by proving a relaxed weighted version. A *weight function* $w\colon V(G) \to \mathbb{N}$ on the vertex set $V(G)$ of a graph G *induces* a tangle τ in G if $w(A) < w(B)$ for every $(A, B) \in \tau$. Note that a set $X \subseteq V(G)$ induces a tangle τ if and only if its indicator function $\mathbb{1}_X$ on $V(G)$ induces τ.

Theorem 8.1.2. [62] *Every tangle in a graph G is induced by some weight function on $V(G)$.*

Contrary to their positive result, Elbracht, Kneip and Teegen [62, Theorem 10] explicitly construct an example which shows that not only Problem 8.1.1, but also Theorem 8.1.2 fails for tangles in general discrete contexts, such as matroids or data sets (see e.g [42, 47, 54, 63]). However, no such example is known for tangles in graphs. Thus, Problem 8.1.1 is open for all $k \geqslant 4$.

8.1.2 Our contributions to Problem 8.1.1

In this chapter we reduce Problem 8.1.1 for every k to graphs whose size is bounded by a function in k:

Theorem 21. *For every integer $k \geqslant 1$, there exists $M = M(k) \in O(3^{k^{k^5}})$ such that for every k-tangle τ in a graph G, there exists a k-tangle τ' in a connected topological minor G' of G with fewer than M edges such that if a weight function w' on $V(G')$ induces the tangle τ', then the weight function w on $V(G)$ which extends w' by zero*[1] *induces the tangle τ. In particular, a set of vertices which induces τ' also induces τ.*

As an immediate corollary of Theorem 21, we obtain that one may verify the validity of Problem 8.1.1 for any fixed k computationally in explicitly bounded, though impractically long, time by checking Problem 8.1.1 for every k-tangle in every connected graph with fewer than M edges.

Corollary 22. *For $k \geqslant 1$, there exists $M = M(k) \in O(3^{k^{k^5}})$ such that Problem 8.1.1 holds for k if it holds for all k-tangles in connected graphs G with fewer than M edges.*

Our second corollary of Theorem 21 asserts that whenever Problem 8.1.1 is true for some fixed k, then every k-tangle is induced by a set of vertices of size bounded in k. This extends the known fact that one may choose the set X in Problem 8.1.1 that induces a given k-tangle with $k \leqslant 2$ to be of size at most k; this follows from the characterisation of 1- and 2-tangles by components and blocks, respectively. More generally, we prove that for every k-tangle τ in a graph G one may choose a weight function w that induces τ, which exists by Theorem 8.1.2, in such way that its *total weight* $w(V(G)) = \sum_{v \in V(G)} w(v)$, which it distributes on $V(G)$, is bounded in k.

Corollary 23. *For every integer $k \geqslant 1$, there exists $K = K(k)$ such that for every k-tangle τ in a graph G there exists a weight function $w : V(G) \to \mathbb{N}$ which induces τ and whose total weight $w(V(G))$ is bounded by K. In particular, the support of w has size $\leqslant K$.*

Moreover, if Problem 8.1.1 holds for k, then every k-tangle in a graph is induced by a set of at most $M(k)$ vertices, where $M(k)$ is given by Theorem 21.

We derive Corollary 23 from Theorem 21 by first fixing a weight function for every k-tangle in a connected graph with fewer than M edges (such weight functions exist by Theorem 8.1.2) and then taking K as the maximum of the total weights of these finitely many fixed weight functions. The moreover-part follows immediately from Theorem 21 by fixing each such weight function to be an indicator function of some inducing set given by the assumed positive answer to Problem 8.1.1 (see Section 8.8 for details).

8.1.3 An inductive proof method for tangles

Our proof of Theorem 21 is based on the following theorem, which allows inductive proofs for statements about tangles in graphs. We expect that this inductive proof method will be of

[1] Given two sets $X' \subseteq X$, a function $w \colon X \to \mathbb{N}$ *extends* a function $w' \colon X' \to \mathbb{N}$ *by zero* if w restricted to X' is w' and w restricted to $X \smallsetminus X'$ is 0.

independent interest.

Theorem 24. *For every integer $k \geqslant 1$ there is some $M(k) \in O(3^{k^{k^5}})$ such that the following holds: Let τ be a k-tangle in a graph G. Then there exists a sequence $G_0, \ldots, G_m$ of graphs and k-tangles τ_i in G_i for every $i \in \{0, \ldots, m\}$ such that*

- *$G_0 = G$, $\tau_0 = \tau$;*
- *G_i is obtained from G_{i-1} by deleting an edge, suppressing a vertex, or taking a proper component;*
- *the k-tangle τ_{i-1} in G_{i-1} survives as the k-tangle τ_i in G_i for every $i \in [m]$;*
- *G_m is connected and has fewer than $M(k)$ edges.*

Before we proceed to explain the crucial term 'survive' in this theorem, we remark that it is necessary to allow the suppression of a vertex in Theorem 24. Indeed, there are connected graphs G with arbitrarily many edges which have a k-tangle such that every graph obtained from G by deleting an edge has no k-tangles at all (Example 8.3.7).

What does it mean that the k-tangle τ 'survives' as a k-tangle τ' in G' obtained from G by deleting an edge, suppressing a vertex, or taking a proper component? Let us here consider the first case that $G' = G - e$ is obtained from G by deleting an edge e of G. Then every separation of G is also a separation of its subgraph G'. But in general G' admits more separations than G; namely those separations $\{A, B\}$ of G' which have an endpoint of e in each of the two sets $A \setminus B$ and $B \setminus A$. A k-tangle τ in G *extends* to a k-tangle τ' in G' if $\tau \subseteq \tau'$. Then we also say that the k-tangle τ *survives* as the k-tangle τ' in G'. We remark that such an extension τ' of τ may or may not exist in G'; if it exists, it need not be unique. If such a τ' exists and is unique, then we also say that τ *induces* τ'.

For the cases that G' is a component of G or obtained from G by suppressing a vertex of G, we refer the reader to Section 8.3.2 for the details. For readers familiar with the fact that a tangle of order $k \geqslant 3$ in a minor of a graph G 'lifts' to a k-tangle in G (cf. [114, (6.1)] and [33, Lemma 2.1]), we remark that 'surviving' is the reverse notion.

8.1.4 Overview of the proof of Theorem 24

It suffices to describe one step of the construction of the above sequence, that is to find G_i and the k-tangle τ_i in it given τ_{i-1} and G_{i-1}. We reduce the argument to several cases. Let τ be a k-tangle in a graph G. The following is an immediate consequence of the well-known correspondence of 1- and 2-tangles in G to the components and blocks of G, respectively:

(i) if $k = 1$, then τ extends to some k-tangle in $G - e$ for every edge e of G (Lemma 8.3.2), and

(ii) if $k = 2$, then τ extends to a k-tangle in $G - e$ for some edge e of G (Lemma 8.3.3).

A rather simple analysis will yield that

(iii) if G is disconnected, then τ induces a k-tangle in a unique component of G (Lemma 8.3.4),

(iv) if $k \geqslant 3$ and e is the unique edge incident to a vertex of degree 1, then τ induces a k-tangle in $G - e$ (Lemma 8.3.5), and

(v) if $k \geqslant 3$, then τ induces a k-tangle in every graph obtained from G by suppressing any vertex of degree 2 (Lemma 8.3.6).

We remark that in each of the above (i) to (v), the tangle τ survives as a k-tangle in a strictly smaller graph (Section 8.3). It remains to consider the case that τ is a tangle of order $k \geqslant 3$ in a connected graph with minimum degree $\geqslant 3$. Recall that a suppressed vertex always has degree 2. So as soon as we restrict to connected graphs of minimum degree at least 3, the statement requires that we find an edge e to delete.

The proof hinges on an argument that this will always be possible: by carefully picking the edge e of G, we may always extend τ in $G' = G - e$ as long as G is sufficiently large by a function in the tangle's order k.

Theorem 25. *For every integer $k \geqslant 3$, there is some $M = M(k) \in O(3^{k^{k^5}})$ such that the following holds:*

For every k-tangle in a connected graph G with minimum degree $\geqslant 3$ and at least M edges, there is an edge e of G such that τ extends to a k-tangle in $G - e$.

Together (i) to (v) along with Theorem 25 complete the proof, since in any given case one of them will ensure that the k-tangle τ in a sufficiently large graph survives in some smaller graph.

8.1.5 Proof sketch of Theorem 25

Our task is to find a suitable edge e of G such that τ extends to some k-tangle τ' in $G' = G - e$. For this, we consider two cases. First, we assume that the graph G contains a tangle $\tilde{\tau}$ of order $> k$ (Section 8.4). If $\tau \subseteq \tilde{\tau}$, then we will observe that for every edge e of G the k-tangle τ extends to the k-tangle τ' in G' which is essentially the restriction of $\tilde{\tau}$ to the separations of order $< k$ of G' (Lemma 8.4.3). Else we may find an edge e far away from τ and close to $\tilde{\tau}$ such that the high order of $\tilde{\tau}$ enables us to define a suitable extension τ' of τ in $G' = G - e$ (Lemma 8.4.4).

Second, we assume that the graph G contains no tangle of order $> k$ (Section 8.7). We remark that this case requires significantly more effort than the first. In the analysis of this case we aim to decompose G in such a way that we can control the separations which arise from the deletion of an edge e in a suitable location (Section 8.6). To obtain the desired decomposition, we start with the tree-decomposition obtained from the tangle-tree duality theorem [41, Theorem 12.5.1] due to the absence of high-order tangles. As G is sufficiently large, the decomposition tree contains a very long path whose structure we may regularise to obtain a *rainbow-cloud decomposition* of G (Theorem 8.5.1).

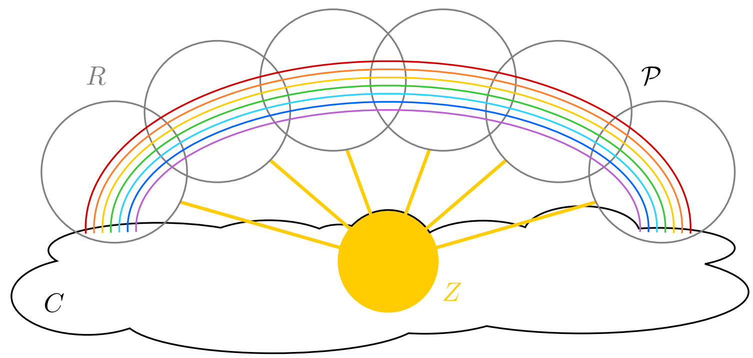

Figure 8.1: Schematic drawing of a rainbow-cloud decomposition

Roughly speaking (see also Figure 8.1), a rainbow-cloud decomposition of G consists of a long *rainbow*, an induced subgraph R of G with a very regular and linear connectivity structure, which has many connections to the *sun* $Z \subseteq V(G)$ and the remaining graph is gathered in a *cloud*, another induced subgraph C with $G = G[Z \cup V(R)] \cup C$ (see Section 8.5 for the precise definition). We show that if we choose the edge e deep inside the rainbow R and far away from τ, then the connectivity structure of the long rainbow R ensures that the new separations which arise from the deletion of e may be oriented in such a way that τ extends to a k-tangle τ' in the graph $G - e$ (Theorem 8.7.1).

8.1.6 How this chapter is organised

We recall the relevant terminology concerning tangles in Section 8.2. We then introduce in Section 8.3 the notion of 'survives' and prove Theorem 24 in several special cases: the case $k \leqslant 2$, the case $\delta(G) \leqslant 2$, and when the graph G is disconnected. In Section 8.4 we then prove Theorem 25, the remaining case of Theorem 24 when G contains a tangle of order $k+1$. In Section 8.5, we introduce the definition of 'rainbow-cloud decompositions' and prove that such a decomposition exists in the absence of $(k+1)$-tangles in G, assuming that G is sufficiently large. Building on several lemmas which we prove in Section 8.6 about the interaction of rainbow-cloud decompositions and separations, we complete the proof of Theorem 25 in Section 8.7. In Section 8.8 we collect all the individual cases to prove Theorem 24, derive Theorem 21, and deduce Corollaries 22 and 23.

8.2 Preliminaries

8.2.1 Separations of sets

While we will only work with separations of graphs in this chapter, we formally introduce the notion of a separation on a set to transfer separations from one graph to another.

Given an arbitrary set V, an *(unoriented) separation* of V is an unordered pair $\{A, B\}$ of subsets A, B of V such that $A \cup B = V$. The *order* of the separation, denoted $|A, B|$, is the cardinality of its *separator* $A \cap B$. Every separation $\{A, B\}$ of V has two *orientations*, (A, B) and (B, A). These orientations of $\{A, B\}$ are *oriented separations* of V, that is, ordered pairs of subsets of V whose union equals V. Given a set V, we write $U(V)$ for the set of all (unoriented) separations of V, and let $S_k(V) := \{\{A, B\} \in U(V) : |A \cap B| < k\}$. The set of all oriented separations of V is denoted by $\vec{U}(V)$ and the set of all oriented separations of V of order less than k by $\vec{S}_k(V)$.

The oriented separations of V have a natural partial order:

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

With this, the set $\vec{U}(V)$ is a lattice with *infimum* $(A, B) \wedge (C, D) := (A \cap C, B \cup D)$ and *supremum* $(A, B) \vee (C, D) := (A \cup C, B \cap D)$. Moreover, $\vec{U}(V)$ is *distributive*, that is, $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$ for all $x, y, z \in \vec{U}(V)$. It is well-known [22, Chapter IX Corollary 1] that a lattice is distributive if and only if every element of the lattice is uniquely determined by its infimum and supremum with any given element.

Let G be a graph. Then $U(G) \subseteq U(V(G))$ by definition. Moreover, $(\vec{U}(G), \leqslant)$ is a sublattice of $(\vec{U}(V), \leqslant)$. If G' is a graph that is obtained from G by deleting some edges, then $U(G) \subseteq U(G')$ and, more precisely, $S_k(G) \subseteq S_k(G')$ for all $k \in \mathbb{N}$. Hence, we may say that a separation of G is a separation of G'.

8.2.2 Topological minors

Let G be a graph. By *suppressing* a degree-2 vertex v of G, we obtain the graph $G' := G - v + uw$ where u, w are the two neighbours of v in G. As this operation is only defined for vertices of degree 2, we often just say that the graph G' is obtained from G by suppressing a vertex. A *topological minor* G' of G is a graph obtained from G by a sequence of deleting edges, suppressing vertices and deleting vertices. Equivalently, G' is obtained from G by a sequence of suppressing vertices in a subgraph of G. We remark that if a topological minor G' of G is connected, we may obtain G' also by a sequence of deleting edges, suppressing vertices and passing to proper components of the current graph.

8.2.3 Lifts of tangles

It is well-known that tangles in minors lift to tangles in the host graph (see [114, (6.1)] or [33, Lemma 2.1]). Here, we introduce the notion only for topological minors.

Let τ' be a k-tangle in a subgraph G' of G. Then the *lift* τ of τ' to G is the set consisting of precisely those $(A, B) \in \vec{S}_k(G)$ whose restriction $(A \cap V(G'), B \cap V(G'))$ is in τ'. It is immediate that τ is a k-tangle in G, since G' is a subgraph of G and thus every forbidden triple in τ restricts to a forbidden triple in τ' by definition.

Now let τ' be a tangle of order $k \geqslant 3$ in a graph G' obtained from G by suppressing a vertex v. Denote the two neighbours of v in G by u_1, u_2. We define the *lift* τ of τ' to G, as follows.

Let $(A, B) \in \vec{S}_k(G)$ be arbitrary. We denote by (A', B') the pair $(A \smallsetminus \{v\}, B \smallsetminus \{v\})$ and by C'_i the set $C' \cup \{u_i\}$ for $C' = A', B'$ and $i = 1, 2$.

(i) If $u_1, u_2 \in A$ or $u_1, u_2 \in B$, then (A', B') is a separation of G', and we include (A, B) in τ if $(A', B') \in \tau'$.

(ii) If, for $\{i, j\} = \{1, 2\}$, $u_i \in A \smallsetminus B$ and $u_j \in B \smallsetminus A$, then (A', B'_i) and (A'_j, B') are separations of G', and we include (A, B) in τ if at least one of (A', B'_i) and (A'_j, B') is in τ'.

Note that in (ii) either both $(A', B'_i), (A'_j, B')$ are in τ' or both their inverses are in τ', which implies that τ does not contain both orientations of a separation, and hence τ is an orientation of $S_k(G)$. Indeed, otherwise $(A', B'_i), (B', A'_j) \in \tau$, as we cannot have $(B'_i, A), (A'_j, B') \in \tau$ because of $(A', B'_i) \leqslant (A'_j, B')$ and the consistency of the tangle τ'. But $(G'[A'] \cup G'[B']) + u_1u_2 = G'$, and thus (A', B'_i), (B', A'_j), $(\{u_1, u_2\}, V(G'))$ would form a forbidden triple in τ', since $(\{u_1, u_2\}, V(G')) \in \tau$ due to the regularity of the k-tangle τ with $k \geqslant 3$. From this it also follows that a forbidden triple in τ would correspond to a forbidden triple in τ', by replacing each (A, B) in the forbidden triple by (A'_1, B') or (A'_2, B') if necessary. Thus, the lift τ of the tangle τ' of G' of order $k \geqslant 3$ to G as defined above is indeed a k-tangle in G.

Given a k-tangle τ' in a topological minor G' of a graph G, we now obtain the *lift* τ of τ' to G by iteratively considering the lifts along the sequence of edge deletions, vertex suppressions and vertex deletions from which G' originated from G.

8.3 Definition of 'survive' and special cases of Theorem 24

In this section we define what it means for a tangle τ in a graph G to 'survive' in a topological minor G' of G. This notion can be seen as a converse to lifting a tangle. But let us emphasise that while a tangle of order at least 3 in G' always lifts to G, a tangle in G need not survive in G'. We first define 'survive' for subgraphs G' of G in Section 8.3.1 and then for graphs G' obtained from G by suppressing a single vertex of degree 2 in Section 8.3.2. Alongside these definitions, we prove several lemmas which all deal with special cases of Theorem 24. Finally, we provide Example 8.3.7, which demonstrates that suppressing vertices of degree 2 needs to be allowed in Theorem 24.

8.3.1 Extending and inducing tangles in subgraphs

Recall that for every subgraph G' of a graph G on the same vertex set, a separation of G is also a separation of G', i.e. $U(G) \subseteq U(G')$ and $S_k(G) \subseteq S_k(G')$ for all integers $k \geqslant 1$. We say that an orientation τ of $S_k(G)$ for some integer $k \geqslant 1$ *extends* to an orientation τ' of $S_k(G')$ if τ' orients every separation in $S_k(G)$ in the same way as τ (equivalently: $\tau \subseteq \tau'$). In this situation, we also sometimes say that the k-tangle τ in G *survives* as the k-tangle τ' in G'. If τ is a k-tangle in G and there exists precisely one k-tangle τ' in G' to which τ extends, then we also say that τ *induces* the k-tangle τ'. In this chapter, G' will often either be a component of G or arise from G by the deletion of a single edge $e \in G$. We remark that if a k-tangle τ in G extends to a k-tangle τ' in a subgraph G' of G, then τ is obviously the lift of τ' to G.

It is well-known [114, (2.6)] that tangles of order 1 and 2 are in a one-to-one correspondence to the components and blocks[2], respectively:

Proposition 8.3.1. *For a tangle τ in a graph G, let $X_\tau := \bigcap_{(A,B)\in\tau} B$. Then the map $\tau \mapsto G[X_\tau]$ is a bijection between the tangles of order 1 and the components of G as well as between the tangles of order 2 and the set of all blocks of G.* □

With this proposition at hand, we can prove Theorem 24 for 1-tangles and 2-tangles.

Lemma 8.3.2. *Let τ be a 1-tangle in a graph G. Then τ extends to a 1-tangle τ' in $G - e$ for every edge $e \in G$.*

Proof. Let e be an arbitrary edge of G, and let $G' := G - e$. By Proposition 8.3.1, the induced subgraph $G[X_\tau]$ of G on the vertex set $X_\tau := \bigcap_{(A,B)\in\tau} B$ is a component of G corresponding to the 1-tangle τ. Let C' be a component of G' whose vertex set is contained in X_τ. By Proposition 8.3.1 the component C' of G' corresponds to a 1-tangle τ' in G' with $X_{\tau'} = \bigcap_{(A,B)\in\tau'} B = V(C')$. Since every vertex of G lies on precisely one side of every separation of G of order 0, a separation (A, B) of G of order 0 is contained in τ' if and only if $X_{\tau'} \subseteq B$. So for every $(A, B) \in \tau$, we have $X_{\tau'} \subseteq X_\tau \subseteq B$ and thus $(A, B) \in \tau'$. Hence, τ extends to τ'. □

We remark that in the proof of Lemma 8.3.2 the component X' need not be unique, as the deletion of e may create (at most) two components in $G[X_\tau]$. Thus, τ does not necessarily extend to a unique 1-tangle in G'.

[2]Recall that a *block* of a graph is a $\subseteq$-maximal connected subgraph H such that $H - v$ is connected for every vertex $v \in H$. In particular, as the empty graph is not considered to be connected, every block contains some edge.

Lemma 8.3.3. *Let τ be a 2-tangle in a graph G with at least 2 edges. Then τ extends to a 2-tangle τ' in $G - e$ for some edge $e \in G$.*

Proof. By Proposition 8.3.1, $G[X_\tau]$ is a block of G; in particular, it contains an edge f. Let $e \neq f$ be any other edge in G, and consider $G' := G - e$. Let X' be the vertex set of the block of G' containing f, and note that $X' \subseteq X_\tau$ by the definition of G'. By Proposition 8.3.1, the block X' corresponds to a 2-tangle τ' of G'; in particular, $f \in G'[X'] = G'[X_{\tau'}]$. Since every edge in G lies on precisely one side of every separation of G of order $\leqslant 1$, a separation (A, B) of G of order $\leqslant 1$ is contained in τ' if and only if $f \in G[B]$. For every $(A, B) \in \tau$, we have $f \in G[X_\tau] \subseteq G[B]$ and hence $(A, B) \in \tau'$. Thus, τ extends to τ'. □

Similarly, as in the proof of Lemma 8.3.2, also in the proof of Lemma 8.3.3 the block X_τ may split into two (or even more) blocks in G'. Thus, τ does not necessarily extend to a unique 2-tangle in G'.

We now prove Theorem 24 in the special cases where the graph G is disconnected or contains a vertex of degree 1.

Lemma 8.3.4. *Let $k \geqslant 1$ be an integer, and let G be a graph with a k-tangle τ. Then there exists a (unique) component G' of G such that τ extends to a k-tangle τ' in G'. Moreover, τ induces the k-tangle τ' in G'.*

Proof. Let τ_1 be the 1-tangle in G with $\tau_1 \subseteq \tau$. By Proposition 8.3.1, the subgraph $G[X]$ on the vertex set $X := X_{\tau_1} = \bigcap_{(A,B)\in\tau_1} B$ is a component of G. We claim that $G' := G[X]$ is as desired.

We define an orientation τ' of $S_k(G')$ based on τ. For any separation $\{A', B'\}$ of order $< k$ of G', we consider the separation $\{A, B\}$ of G defined by $A := A' \cup (V(G) \smallsetminus X)$ and $B := B'$. This again has order $< k$ and is hence oriented by τ. So if $(A, B) \in \tau$, then we put $(A', B') \in \tau'$, and if $(B, A) \in \tau$, we put $(B', A') \in \tau'$.

We now show that τ extends to τ'. It suffices to check that if we included (A', B') or (B', A') in τ because of (A, B) or (B, A) in τ, respectively, then there exists no $(C, D) \in \tau$ such that $(C \cap V(X), D \cap V(X)) = (B', A')$ or $= (A', B')$, respectively. Suppose for a contradiction that there is such a $(C, D) \in \tau$. The definition of $X := X_{\tau_1}$ ensures that the separation $((V(G) \smallsetminus X), X)$ of order 0 is in $\tau_1 \subseteq \tau$. Thus, the forbidden triple

$$\{(A, B), (C, D)), ((V(G) \smallsetminus X), X)\} \text{ or } \{(B, A), (C, D)), ((V(G) \smallsetminus X), X)\})$$

is contained in the tangle τ, respectively, which is a contradiction. Thus, τ extends to τ'.

It remains to show that the orientation τ' of $S_k(G')$ is even a tangle in G'. Suppose for a contradiction that there exists a forbidden triple $\{(A'_i, B'_i) : i \in [3]\}$ in τ'. As above, $(A_i, B_i) := (A'_i \cup (V(G) \smallsetminus X), B'_i)$ is a separation of G of order $< k$ and by construction contained in τ for all $i \in [3]$. Since G' is a component of G, it is immediate that $\{(A_i, B_i) : i \in [3]\}$ forms a forbidden triple in τ, which is a contradiction.

We remark that the proof immediately ensures that G' is the unique component such that τ extends to a k-tangle in it, and also τ' is the unique such k-tangle. $\square$

Lemma 8.3.5. *Let τ be a tangle in G of order $k \geqslant 3$. Suppose that G has a vertex v of degree 1 and let e be the unique edge incident to v. Then τ induces a k-tangle τ' in $G - e$.*

Proof. Let τ' be the subset of $\vec{S}_k(G')$ which contains $(A, B) \in \tau'$ if $(A, B) \in \tau$, $(A \setminus \{v\}, B \cup \{v\}) \in \tau$ or $(A \cup \{v\}, B \setminus \{v\}) \in \tau$. The regularity of the tangle τ of order $\geqslant 3$ ensures that $(e, V(G)) \in \tau$. Thus, τ' is an orientation of $S_k(G')$, as any violation would form together with $(e, V(G))$ a forbidden triple in τ.

Suppose for a contradiction that $\{(A'_i, B'_i) : i \in [3]\}$ is a forbidden triple in τ'. Since $G' = \bigcup_{i \in [3]} G'[A'_i]$, some A'_j contains the other endvertex $u \neq v$ of e. Thus, $(A_j, B_j) := (A'_j \cup \{v\}, B'_j \setminus \{v\})$ is a separation of G and thus in τ, as τ' is an orientation. Let $(A_i, B_i) \in \tau$ which witnesses that $(A'_i, B'_i) \in \tau'$ for $i \in [3] \setminus \{j\}$. Now the (A_i, B_i) form a forbidden triple in τ, which is a contradiction. $\square$

8.3.2 Inducing tangles in graphs obtained by vertex suppression

Recall that for a given vertex $v \in V(G)$ of degree 2, the graph obtained from G by *suppressing the vertex* v is $G' := G - v + uw$ where u, w are the two neighbours of v in G. Let $\{A, B\}$ be a separation of G'. If $u, w \in A$, then $\{A \cup \{v\}, B\}$ is a separation of G, and analogously if $u, w \in B$, then $\{A, B \cup \{v\}\}$ is a separation of G. In particular, they have the same order as $\{A, B\}$. We remark that at least one of $u, w \in A$ and $u, w \in B$ holds, as uw is an edge of G'.

We say that a k-tangle τ of $S_k(G)$ for some integer $k \geqslant 1$ *induces* the subset $\tau' \subseteq \vec{S}_k(G')$ consisting of those $(A, B) \in \vec{S}_k(G')$ such that at least one of $(A \cup \{v\}, B)$ and $(A, B \cup \{v\})$ is in τ. In particular, if $\{A \cup \{v\}, B\}$ is a separation of G, then one of its orientations is contained in τ, and thus at least one of (A, B) or (B, A) is in τ'; but τ' might contain both if also $\{A, B \cup \{v\}\}$ is a separation of G. Lemma 8.3.6 below ensures that if $k \geqslant 3$, then not only the latter does not happen, but τ' is even a k-tangle. In this situation, we also sometimes say that the k-tangle τ in G *survives* as the k-tangle τ' in the graph G' obtained from G by suppressing a vertex. We remark that if a k-tangle τ in G extends to a k-tangle τ' in the graph G' obtained from G by suppressing a vertex, then τ is obviously the lift of τ' to G.

Lemma 8.3.6. *Let G be a graph, and let $G' = G - v + uw$ be the graph obtained by suppressing a vertex v of degree 2, where its two neighbours in G are u, w. Then for a given k-tangle τ with $k \geqslant 3$ in G the set $\tau' \subseteq \vec{S}_k(G')$ induced by τ is a k-tangle in G'.*

Proof. We claim that $(A \cup \{v\}, B) \in \tau$ if and only if $(A, B \cup \{v\}) \in \tau$, if both $\{A \cup \{v\}, B\}$ and $\{A, B \cup \{v\}\}$ are separations of G. Suppose for a contradiction that this is not the case. The consistency of τ ensures that $(A, B \cup \{v\}) \in \tau$ and $(B, A \cup \{v\}) \in \tau$. Since $k \geqslant 3$ and every tangle is regular, $(\{u, v, w\}, V(G) \setminus \{v\}) \in \tau$. Now $\{(A, B \cup \{v\}), (B, A \cup \{v\}), (\{u, v, w\}, V(G) \setminus \{v\})\}$ forms a forbidden triple in the tangle τ, which is a contradiction.

The claim above shows that τ' is an orientation of $S_k(G')$. It remains to show that it is a k-tangle in G'. Suppose that $\{(A'_i, B'_i) : i \in [3]\}$ is a forbidden triple in τ'. Then there is $j \in [3]$ with $u, w \in A'_j$, since uw is an edge in G'. Hence, $\{A'_j \cup \{v\}, B'_j\}$ is a separation of ‘G, and thus $(A_j, B_j) := (A'_j \cup \{v\}, B'_j) \in \tau$ by definition of τ'. For every $i \in [3] \setminus \{i\}$, we let $(A_i, B_i) \in \tau$ be the separation which witnesses that $(A'_i, B'_i) \in \tau'$. Then $\{(A_i, B_i) : i \in [3]\}$ is a forbidden triple in τ, as $\{(A'_i, B'_i) : i \in [3]\}$ is a forbidden triple in G' and $u, v, w \in A_j$. This is a contradiction to the τ being a tangle in G. □

8.3.3 Suppressing vertices of degree 2 in Theorem 24

We conclude this section with an example that shows that there are connected graphs G with arbitrarily many edges which have a k-tangle such that every graph obtained from G by deleting an edge has no k-tangles at all. In particular, in Theorem 24 it is necessary to allow the suppression of a vertex.

Example 8.3.7. *For every $k, M \in \mathbb{N}$ with $k \geqslant 3$, there exists a connected graph G with at least M edges and which has a k-tangle such that, for every edge e of G, the graph $G - e$ does not have a k-tangle.*

Proof. Let G' be some connected graph which has a k-tangle but which is such that, for every edge e of G', the graph $G' - e$ does not have any k-tangles. Such graphs exist: Take any graph H that has a k-tangle, and let $H =: H_0 \supset H_1 \supset \cdots \supset H_n$ be a maximal sequence such that H_{i+1} is obtained from H_i by either deleting an edge or taking a proper component of H_i, and such that H_n still has a k-tangle. Set $G' := H_n$. By the maximal choice of the sequence $(H_i)_{i \in [n]}$ and because of Lemma 8.3.4, G' is connected and no graph obtained from G' by deleting an edge has a k-tangle.

Since edgeless graphs have no tangles of order $\geqslant 2$, the graph G' contains an edge uv. Let G be obtained from G' replacing uv by a u–v path of length at least $M + 1$. Then G has at least M edges, and it has a k-tangle as every tangle of order $k \geqslant 3$ in G' lifts to a k-tangle in G. But $G - e$ has no k-tangle for every edge e of G, since any such a tangle would induce a k-tangle in $G' - e'$ by Lemmas 8.3.5 and 8.3.6 where $e' := e$ if $e \in E(G)$ or $e' := uv$ otherwise. □

8.4 Proof of Theorem 25 if G has a higher-order tangle

In this section we prove Theorem 25 for graphs which contain a tangle of order $> k$.

Theorem 8.4.1. *Let τ be a tangle in G of order $k \geqslant 2$. Suppose further that there exists a $(k+1)$-tangle τ^* in G. Then there is an edge $e \in E(G)$ such that τ extends to some k-tangle τ' in $G - e$.*

We distinguish two cases: First, we consider the case where τ itself *extends* to a $(k+1)$-tangle in G, i.e. $\tau \subseteq \tilde{\tau}$ for some $(k+1)$-tangle $\tilde{\tau}$ in G. We then show that τ extends to a k-tangle in $G - e$ for *every* edge $e \in E(G)$ (Lemma 8.4.3). Otherwise, there exists a $(k+1)$-tangle $\tilde{\tau}$ in G which does not extend τ. Here, τ does not extend to a k-tangle in $G - e$ for every edge $e \in G$, but only for some such edges e and we will need some care to find them (Lemma 8.4.4).

For both cases, we will make use of the following observation.

Lemma 8.4.2. *Let $k \geqslant 3$ be an integer, and let G be a graph with a k-tangle τ. Further, let e be an edge of G and let $\{A, B\}$ be a separation of $G - e$ with $e \in E_G(A \setminus B, B \setminus A)$. If $|A \cap B| < k-1$, then $(A \cup e, B) \in \tau$ if and only if $(A, B \cup e) \in \tau$.*

Proof. The assumptions immediately imply that $\{A \cup e, B\}$ and $\{A, B \cup e\}$ are separations of G of order less than k, and so τ orients both of them. The consistency of the tangle τ together with $(A, B \cup e) \leqslant (A \cup e, B)$ yields the forwards implication. The backwards implication follows from the fact that $\{(A, B \cup e), (B, A \cup e), (e, V(G))\}$ is a forbidden triple but $(e, V(G)) \in \tau$, since τ is a k-tangle with $k \geqslant 3$ and $V(G)$ cannot be the small side of any separation in τ. ☐

We start with the case in which τ extends to a $(k+1)$-tangle in G:

Lemma 8.4.3. *Let $k \geqslant 2$ be an integer, and let G be a graph with a k-tangle τ. If τ extends to a $(k+1)$-tangle $\tilde{\tau}$ in G, then τ extends to a k-tangle τ' in $G - e$ for every edge $e \in E(G)$.*

Proof. Consider $G' := G - e$. We define an orientation τ' of $S_k(G')$ as follows: If $\{A, B\} \in S_k(G')$ is also a separation of G, then we put $(A, B) \in \tau'$ if and only if $(A, B) \in \tau$. Otherwise, the edge e has one endvertex in $A \setminus B$ and the other one in $B \setminus A$. So both $\{A \cup e, B\}$ and $\{A, B \cup e\}$ are separations of G of order $|A, B| + 1 \leqslant k$, and these two separations are oriented by the $(k+1)$-tangle $\tilde{\tau}$, and we have $(A \cup e, B) \in \tau$ if and only if $(A, B \cup e) \in \tilde{\tau}$ by Lemma 8.4.2. Then we put $(A, B) \in \tau'$ if and only if $(A \cup e, B) \in \tilde{\tau}$ (equivalently: $(A, B \cup e)$). The first part of the definition guarantees that τ extends to τ'.

It remains to show that τ' is a k-tangle in G'. If there exists a forbidden triple $\{(A_i, B_i) : i \in [3]\}$ in G' for τ', then we obtain a forbidden triple for $\tilde{\tau}$ in G by replacing those (A_i, B_i) that are not already separations of G with $(A_i \cup e, B_i)$. The arising triple is then by the definition of τ' a forbidden triple in $\tilde{\tau}$. This contradicts that $\tilde{\tau}$ is a tangle in G by assumption. ☐

We remark that one can also show a vertex-version of Lemma 8.4.3 along the same lines: if a k-tangle τ in G extends to a $(k+1)$-tangle in G, then τ extends to a k-tangle in $G' := G - v$ for every vertex $v \in V(G)$.

Now we turn to the case that G has a $(k+1)$-tangle $\tilde{\tau}$ which does not extend τ. Let us briefly describe our proof strategy: First, we observe that there exists a separation $(B, A) \in \tilde{\tau}$ which is $\leqslant$-maximal in $\tilde{\tau} \cap \vec{S}_k(G)$ and distinguishes τ and $\tilde{\tau}$. We will then delete an arbitrary edge e on the side of $\{A, B\}$ which is small with respect to τ, i.e. $e \in G[A \setminus B]$. To define the desired k-tangle τ' of G' to which τ shall extend, we first orient all separations of G' that are 'forced' by τ in that they have an orientation which was either already in τ or which must be in τ' to achieve the desired consistency of τ'. The remaining separations are then oriented according to $\tilde{\tau}$; for this, we draw on the fact that $\tilde{\tau}$ has order $k+1$ and hence naturally defines an orientation of all the separations in $S_k(G')$ using Lemma 8.4.2. While τ extends to this orientation τ' by construction, the main part of the proof is devoted to show that τ' is indeed a tangle. Intuitively speaking, the construction of τ' ensures that neither the separations forced by τ nor those oriented according to $\tilde{\tau}$ contain a forbidden triple. The maximal choice of (B, A) together with submodularity arguments then ensures that there is also no forbidden triple consisting of both kinds of separations: it allows us to transfer any such forbidden triple in τ' either into one in τ or into one in $\tilde{\tau}$.

Lemma 8.4.4. *Let $k \geqslant 2$ be an integer, and let G be a graph with a k-tangle τ. If there exists a $(k+1)$-tangle $\tilde{\tau}$ in G with $\tau \not\subseteq \tilde{\tau}$, then there exists an edge $e \in E(G)$ such that τ extends to a k-tangle τ' in $G - e$.*

Proof. We first find an edge e of G that we afterwards prove to be as desired. For this, let $(B, A) \in \tilde{\tau}$ be a separation of G which distinguishes τ and $\tilde{\tau}$ and is $\leqslant$-maximal in $\tilde{\tau}$ among all such distinguishing separations. Note that (B, A) is even maximal in $\tilde{\tau} \cap \vec{S}_k$: any separation $(C, D) \in \tilde{\tau} \cap \vec{S}_k$ with $(B, A) < (C, D)$ would also distinguish τ and $\tilde{\tau}$ since $(D, C) \in \tau$ by the consistency of τ. We then choose an arbitrary edge e in $G[A \setminus B]$. Let us first show its existence.

There exists a vertex $v \in A \setminus B$, since otherwise $(V(G), A) = (B, A) \in \tilde{\tau}$, contradicting that $\tilde{\tau}$ is a tangle. If the vertex v has a neighbour in $A \setminus B$, then the edge joining them is as desired. So suppose for a contradiction that all neighbours of v are in $A \cap B$. We can then find a forbidden triple in $\tilde{\tau}$: First, we can move v from $A \setminus B$ to $B \setminus A$ to obtain a new separation $\{A \setminus \{v\}, B \cup \{v\}\}$ of G, which has the same order as $\{A, B\}$. Thus, $\tilde{\tau}$ contains an orientation of it, and we must have $(A \setminus \{v\}, B \cup \{v\}) \in \tilde{\tau}$ due to the maximality of (B, A) in $\tilde{\tau} \cap \vec{S}_k(G)$. Secondly, since $|A \cap B| < k$ and $\tilde{\tau}$ is a $(k+1)$-tangle in G, we have $((A \cap B) \cup \{v\}, V(G)) \in \tilde{\tau}$. Hence, $\{(B, A), (A \setminus \{v\}, B \cup \{v\}), ((A \cap B) \cup \{v\}, V(G))\}$ is contained in the tangle $\tilde{\tau}$, but it is also a forbidden triple, which is a contradiction. All in all, v has a neighbour in $A \setminus B$; in particular, $G[A \setminus B]$ contains an edge.

From now on, we prove that the chosen edge $e \in G[A \setminus B]$ is as desired. For this, we consider $G' := G - e$ and construct an orientation τ' of $S_k(G')$ to which τ extends. It then remains to show that τ' is a tangle in G'.

For the construction of τ', note that τ' has to contain not only τ, but also all orientations of separations of $S_k(G')$ that are 'forced' by the request that τ' shall again be a tangle and hence especially consistent. More formally, we say that τ *forces* an orientation of a separation $\{C, D\}$ of G' if there exists a separation $(E, F) \in \tau$ such that $(C, D) \leqslant (E, F)$ or $(D, C) \leqslant (E, F)$. In particular, τ forces an orientation of every separation in $S_k(G')$ which is also a separation of G and the separations in $\vec{S}_k(G')$ which are maximal among all those forced by τ are separations of G. Note that the consistency of τ ensures that at most one orientation of a separation in $S_k(G')$ is forced by τ.

We now define the orientation τ' of $S_k(G')$. If τ forces an orientation of a separation $\{C, D\}$ in $S_k(G')$, then we put the respective orientation in τ'. Otherwise, $\{C, D\}$ is especially not a separation of G, so e has one endvertex in $C \setminus D$ and the other one in $D \setminus C$. Then $\{C \cup e, D\}$ and $\{C, D \cup e\}$ are separations of G of order at most k, as $\{C, D\}$ has order less than k. Thus, both these separations are oriented by the $(k+1)$-tangle $\tilde{\tau}$ in G, and by Lemma 8.4.2, we have $(C \cup e, D) \in \tilde{\tau}$ if and only if $(C, D \cup e) \in \tilde{\tau}$. Now if $(C \cup e, D) \in \tilde{\tau}$ (equivalently: $(C, D \cup e) \in \tilde{\tau}$), then we put $(C, D) \in \tau'$, and if $(D \cup e, C) \in \tilde{\tau}$ (equivalently: $(D, C \cup e) \in \tilde{\tau}$), then we put $(D, C) \in \tau'$. This definition of τ' ensures that τ' indeed is an orientation of $S_k(G')$ and also that τ extends to τ'.

Thus, it remains to show that τ' is indeed a tangle. Suppose for a contradiction that there is a forbidden triple $\{(C_i, D_i) : i \in [3]\}$ in τ'. Without loss of generality, we may assume that all the (C_i, D_i) are $\leqslant$-maximal in τ'. We now aim to use $\{(C_i, D_i) : i \in [3]\}$ together with the construction of τ' to find a forbidden triple in G which is contained in either τ or $\tilde{\tau}$. This then yields a contradiction since both τ and $\tilde{\tau}$ are tangles in G. Towards this, we first give a condition on the (C_i, D_i) which allows us to find a forbidden triple in $\tilde{\tau}$ and prove afterwards that if this condition does not hold, then we can find a forbidden triple in τ.

First, suppose that each $\{C_i, D_i\}$ either crosses $\{A, B\}$ or satisfies $(C_i, D_i) \leqslant (B, A)$. In this case, we aim to find a forbidden triple in $\tilde{\tau}$. Towards this, the following lemma shows that $(C_i, D_i) \in \tilde{\tau}$ if $\{C_i, D_i\}$ crosses $\{A, B\}$ and is also a separation of G.

Claim 1. *Let $\{C, D\}$ be a separation of G' that is also a separation of G and whose orientation $(C, D) \in \tau'$ is $\leqslant$-maximal in τ'. If $\{C, D\}$ crosses $\{A, B\}$, then $(C, D) \in \tilde{\tau}$.*

Proof. Assume that the infimum $(A \cap C, B \cup D)$ of (A, B) and (C, D) has order less than k. By the maximality of (B, A) in $\tilde{\tau} \cap \vec{S}_k(G)$, we then have $(A \cap C, B \cup D) \in \tilde{\tau}$. Since $(B, A), (D, C)$ and $(A \cap C, B \cup D)$ form a forbidden triple in G, this then implies $(C, D) \in \tilde{\tau}$, as desired.

It remains to prove that $\{A \cap C, B \cup D\}$ has order less than k. By submodularity, it suffices to show that $\{A \cup C, B \cap D\}$ has order at least k. Suppose for a contradiction that it has order less than k. Then τ contains an orientation of $\{A \cup C, B \cap D\}$. Since τ extends to τ', we have (C, D) in τ. On the one hand, as (C, D) is also $\leqslant$-maximal in τ, we must have that its supremum $(A \cup C, B \cap D)$ with (A, B) is not in τ. On the other hand, the profile property $(*)$ of τ ensures that $(B \cap D, A \cup C) \notin \tau$, as $(A, B) \in \tau$. This is a contradiction. ■

In this first case, where each $\{C_i, D_i\}$ either crosses $\{A, B\}$ or satisfies $(C_i, D_i) \leqslant (B, A)$, we can use Claim 1 to obtain a forbidden triple $\{(C'_i, D'_i) : i \in [3]\}$ in $\tilde{\tau}$ as follows: For $i \in [3]$, assume first that $\{C_i, D_i\}$ is also a separation of G. If $(C_i, D_i) \leqslant (B, A)$, then $(C'_i, D'_i) := (C_i, D_i) \in \tilde{\tau}$ since $(B, A) \in \tilde{\tau}$ and $\tilde{\tau}$ is consistent. Otherwise, $\{C_i, D_i\}$ crosses $\{A, B\}$ and then $(C'_i, D'_i) := (C_i, D_i) \in \tilde{\tau}$ by Claim 1. Secondly, if $\{C_i, D_i\}$ is not a separation of G, then the maximality of (C_i, D_i) in τ' implies that (C_i, D_i) cannot be forced by τ. Thus, we have $(C'_i, D'_i) := (C_i \cup e, D_i) \in \tilde{\tau}$ by construction. Now $\{(C'_i, D'_i) : i \in [3]\}$ is a forbidden triple in the tangle $\tilde{\tau}$ in G, which is a contradiction.

So we may now assume that some $\{C_i, D_i\}$, say $\{C_1, D_1\}$, neither crosses $\{A, B\}$ nor satisfies $(C_i, D_i) \leqslant (B, A)$. In this case, we aim to find a forbidden triple in τ. We claim that $(A, B) \leqslant (C_1, D_1)$ and that this yields $(C_1, D_1) \in \tau$: By our assumptions, $\{C_1, D_1\}$ is nested with $\{A, B\}$, but we do not have $(C_1, D_1) \leqslant (B, A)$ (equivalently: $(A, B) \leqslant (D_1, C_1)$). If $(D_1, C_1) \leqslant (A, B)$, then $(D_1, C_1) \in \tau'$ is forced by $(A, B) \in \tau$ but τ' is an orientation which already contains (C_1, D_1), which is a contradiction. Furthermore, we cannot have $(C_1, D_1) < (A, B)$, since (C_1, D_1) is maximal in τ'. Thus, $(A, B) \leqslant (C_1, D_1)$; in particular, $e \in G[A \setminus B] \subseteq G[C_1 \setminus D_1]$. Therefore, (C_1, D_1) is not only a separation of $G - e = G'$, but also one of G. Since τ extends to τ' and $(C_1, D_1) \in \tau'$, we obtain $(C_1, D_1) \in \tau$.

As shown, we have $e \in G[C_1]$. So if (C_2, D_2) and (C_3, D_3) are separations of G, then they are not only in τ' but also in τ as τ extends to τ' which yields our desired forbidden triple in τ. So suppose that (C_i, D_i) with $i \in \{2, 3\}$ is not a separation of G. We claim that $\{C_i, D_i\}$ crosses $\{A, B\}$. Indeed, if $\{C_i, D_i\}$ had an orientation that is greater than (A, B), then $\{C_i, D_i\}$ would be a separation of G, as the deleted edge e is contained in $G[A \setminus B]$. If $(C_i, D_i) < (A, B)$, then (C_i, D_i) would in contradiction to its choice not be maximal in τ', since (A, B) is also in τ'. If $(D_i, C_i) \leqslant (A, B)$, then $(D_i, C_i) \in \tau'$ is forced by $(A, B) \in \tau$, which is a contradiction to $(C_i, D_i) \in \tau'$. So $\{C_i, D_i\}$ cannot be nested with $\{A, B\}$, that is, they cross. But then the following lemma shows that the infimum of (C_i, D_i) and (B, A) is in τ.

Claim 2. *Let $\{C, D\}$ be a separation of G' that is not a separation of G and whose orientation $(C, D) \in \tau'$ is maximal in τ'. If $\{C, D\}$ crosses $\{A, B\}$, then $(B \cap C, A \cup D) \in \tau$.*

Proof. Since $e \in G[A \setminus B]$, $\{B \cap C, A \cup D\}$ is a separation of G. Assume that it has order less than k.

Then τ contains an orientation of $\{B \cap C, A \cup D\}$, and this orientation must not be $(A \cup D, B \cap C)$ as τ would otherwise force the orientation (D, C) of $\{C, D\}$ to be in the orientation τ' which already contains (C, D).

It remains to show that $\{B \cap C, A \cup D\}$ indeed has order less than k; suppose for a contradiction otherwise. Since both $\{A, B\}$ and $\{C, D\}$ have order at most $k-1$, this implies that $\{B \cup C, A \cap D\}$ has order less than $k-1$ by submodularity. The edge e is in $G[A \setminus B]$ by its choice. Additionally, it has one endvertex in $C \setminus D$ and the other one in $D \setminus C$ because $\{C, D\}$ is not a separation of G by assumption. Therefore, the order of $\{B \cup C \cup e, A \cap D\}$ increases compared to the order of $\{B \cup C, A \cap D\}$ by exactly one. So $\{B \cup C \cup e, A \cap D\}$ has order $< k$. We now show that none of its orientations can be contained in the tangle $\tilde{\tau}$, which then yields the desired contradiction.

On the one hand, the maximality of (B, A) in $\tilde{\tau} \cap \vec{S}_k$ implies that $(B \cup C \cup e, A \cap D) \notin \tilde{\tau}$. On the other hand, since $\{C, D\}$ is not a separation of G, but (C, D) is maximal in τ', the orientation (C, D) of $\{C, D\}$ cannot be forced by τ. Hence by construction of τ', we put $(C, D) \in \tau'$ because of $(C \cup e, D) \in \tilde{\tau}$. But $\{(B, A), (C \cup e, D), (A \cap D, B \cup C \cup e)\}$ is forbidden triple in G, so $(A \cap D, B \cup C \cup e) \notin \tilde{\tau}$. ■

Using Claim 2, we can now find a forbidden triple $\{(C_1, D_1), (C_2', D_2'), (C_3', D_3')\}$ in τ as follows: As shown above, (C_1, D_1) is in τ and satisfies $(A, B) \leqslant (C_1, D_1)$. If (C_i, D_i) with $i \in \{2, 3\}$ is a separation of G, then it also is in τ, as τ extends to τ', and we set $(C_i', D_i') := (C_i, D_i)$. If it is not a separation of G, then $\{C_i, D_i\}$ must cross $\{A, B\}$, as shown above Claim 2, and Claim 2 yields $(C_i', D_i') := (B \cap C_i, A \cup D_i) \in \tau$. To see that $\{(C_1, D_1), (C_2', D_2'), (C_3', D_3')\} \subseteq \tau$ is indeed a forbidden triple in G, note that

$$G[C_1] \cup G[C_2'] \cup G[C_3'] \supseteq G[C_1] \cup G[C_2 \cap B] \cup G[C_3 \cap B] \supseteq G[C_1] \cup G[C_2 \cap D_1] \cup G[C_3 \cap D_1] = G,$$

where the last equation holds because $e \in G[A] \subseteq G[C_1]$ and $\{(C_i, D_i) : i \in \{1, 2, 3\}\}$ is a forbidden triple in $G' = G - e$. This concludes the proof. □

Proof of Theorem 8.4.1. This follows immediately from Lemmas 8.4.3 and 8.4.4. □

Our proof of Theorem 8.4.1 heavily relies on the fact that the order of the additional tangle $\tilde{\tau}$ is greater than the one of τ. However, we do not know whether similar proof techniques could also work if the order of $\tilde{\tau}$ does not exceed the one of τ:

Problem 8.4.5. *Let τ be a tangle in G of order $k \geqslant 3$. Suppose further that there exists another k-tangle τ^* in G with $\tau^* \nsubseteq \tau$ and G has minimum degree at least 3. Is there an edge $e \in G$ such that τ extends to a k-tangle τ' in $G - e$?*

8.5 Rainbow-Cloud-Decompositions in the absence of high-order tangles

Recall that Theorem 8.4.1 in Section 8.4 immediately yields Theorem 25 if the graph has a $(k+1)$-tangle. So from now on we work towards a proof of Theorem 25 for graphs without tangles of order $\geqslant k+1$. In this section, we show that, in the absence of $(k+1)$-tangles, a large graph admits a certain type of decomposition, which we will call 'rainbow-cloud decomposition'; this decomposition is inspired by [92]. We will later use that this decomposition exhibits a substructure of the graph, the 'rainbow', which is a long linear structure that is fairly independent of the rest of the graph and internally made up of similar enough parts such that deleting an edge in one of the parts does not change the overall structure of the graph. In particular, it will allow us to understand how the separations of the graph change after deleting such an edge and hence how to find a tangle of this smaller graph to which our given tangle extends.

We begin by building up to the definition of a 'rainbow-cloud decomposition'. Let G be a graph. First, a *linear decomposition*[3] of G of *length* $M \in \mathbb{N}$ of G is a family $\mathcal{W} = (W_0, W_1, \dots, W_M)$ of sets W_i of vertices of G such that

(L1) $\bigcup_{i=0}^{M} G[W_i] = G$,

(L2) if $0 \leqslant i \leqslant j \leqslant k \leqslant M$, then $W_i \cap W_k \subseteq W_j$, and

(L3) there is an integer ℓ such that $|W_{i-1} \cap W_i| = \ell$ for every $i \in [M]$, and

(L4) $W_{i-1} \neq W_{i-1} \cap W_i \neq W_i$ for all $i \in [M]$.

We call the sets W_i the *bags* and and the induced subgraphs $G[W_i]$ the *parts* of the linear decomposition $\mathcal{W}$. Note that the bags of a linear decomposition of length at least 1 are non-empty by (L4). The *adhesion sets* of a linear decomposition $\mathcal{W}$ are the sets $U_i := W_{i-1} \cap W_i$ for $i \in [M]$. The size of the adhesion sets U_i is the *adhesion* of $\mathcal{W}$. We emphasise that adhesion 0 is allowed. Whenever we introduce a linear decomposition as $\mathcal{W}$ without specifying its bags, then we will tacitly assume the bags to be denoted by $W_0, \dots, W_M$ and the adhesion sets by $U_1, \dots, U_M$.

Next we turn to the definition of 'rainbow-decompositions' which are special linear decompositions whose adhesion sets are minimal U_1–U_M separators of G as witnessed by respective families of disjoint paths. To make this formal, a *linkage* in a graph G is a set $\mathcal{P}$ of disjoint paths in G. If A and B are sets of vertices of G such that $\mathcal{P}$ consists of A–B paths, i.e. such paths that meet A precisely in one endvertex and B precisely in the other endvertex, then $\mathcal{P}$ forms an A–B *linkage*. A linear decomposition $\mathcal{W}$ of adhesion ℓ and length M is called a *rainbow-decomposition* of *adhesion* ℓ and *length* M if it has the following three properties:

[3] Decompositions satisfying (L1) and (L2) are often known as *path-decompositions* (cf. [41, §12.6]). In [92] these are refereed to as *linear decompositions*. In this chapter, linear decompositions will always not only satisfy (L1) and (L2) but also (L3) and (L4).

(R1) There exists a U_i–U_{i+1} linkage of cardinality ℓ in $G[W_i]$ for every $i \in [M-1]$.

(R2) Every part $G[W_i]$ of $\mathcal{W}$ is connected.

(R3) Every two consecutive adhesion sets U_i, U_{i+1} are disjoint.

We may combine the linkages of cardinality ℓ from (R1) to obtain a U_1–U_M linkage $\mathcal{P}$ of cardinality ℓ in G. We call such a linkage $\mathcal{P}$ a *foundational linkage* of the rainbow-decomposition.

Finally, we define 'rainbow-cloud-decompositions', which consist of a rainbow-decomposition of a subgraph of G that interacts with the remainder of G, the 'cloud', in a very controlled manner (see Figure 8.2 for an illustration). Formally, a *rainbow-cloud-decomposition* (or *RC-decomposition* for short) of G of *adhesion* ℓ and *length* M is a quadruple $(R, \mathcal{W}, Z, C)$ consisting of two induced subgraphs R and C of a graph G, a vertex set $Z \subseteq V(C)$ disjoint from $V(R)$ such that $G[V(R) \cup Z] \cup C = G$ and a rainbow-decomposition $\mathcal{W} = (W_0, \dots, W_M)$ of R of adhesion ℓ and length M with adhesion sets $U_1, \dots, U_M$ and two additional adhesion sets $U_0 := V(C) \cap W_0$ and $U_{M+1} := V(C) \cap W_M$ such that

(RC1) $V(R) \cap V(C) = U_0 \cup U_{M+1}$,

(RC2) $|U_0| = \ell = |U_{M+1}|$ and $U_0 \cap U_1 = \emptyset = U_M \cap U_{M+1}$,

(RC3) there exists a U_0–U_1 linkage in $G[W_0]$ and a U_M–U_{M+1} linkage in $G[W_M]$, both of cardinality ℓ, and

(RC4) $Z \subseteq N_G(W_i)$ for every $i \in \{0, \dots, M\}$.

We refer to R as the *rainbow*, to C as the *cloud* of the RC-decomposition and to Z as the *sun* of the RC-decomposition. Whenever we introduce an RC-decomposition $(R, \mathcal{W}, Z, C)$, we tacitly assume that $\mathcal{W} = (W_0, \dots, W_M)$ and $U_0, \dots, U_{M+1}$ are defined as above.

With the definition of rainbow-cloud-decompositions at hand, we can state the main result of this section:

Theorem 8.5.1. *For every two integers $k, M \geqslant 1$, there exists some integer $N = N(k, M) \geqslant 1$ such that every connected graph G with at least N vertices and no $(k+1)$-tangle admits an RC-decomposition $(R, \mathcal{W}, Z, C)$ of length at least M and adhesion ℓ such that $|Z| + \ell \geqslant 1$.*

The remainder of this section is devoted to the proof of Theorem 8.5.1, which roughly proceeds as follows: We will start by using the tangle-tree duality theorem, one of the two central tangle theorems, to get a long nested sequence of separations which will allow us to construct a still long linear decomposition. This linear decomposition can subsequently be refined into an RC-decomposition, by putting the 'unnecessary bits' of its bags into the cloud.

As a first step, Lemma 8.5.3 asserts that every sufficiently long sequence of separations contains a long subsequence which induces a linear decomposition of some adhesion ℓ satisfying (R1). In other words, we find a subsequence such that all its elements have the same order and such that

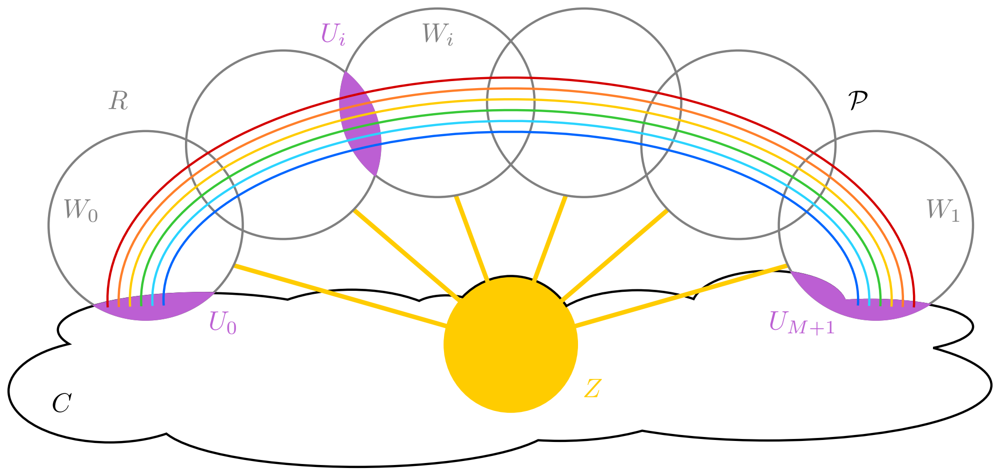

FIGURE 8.2: A rainbow-cloud decomposition $(R, \mathcal{W}, Z, C)$: the rainbow R together with its decomposition $\mathcal{W}$ is indicated in grey, the foundational linkage $\mathcal{P}$ of $\mathcal{W}$ is depicted in rainbow colours, and the cloud C is depicted in black. Further, indicated in yellow, is the sun $Z \subseteq V(C)$ together with the Z–R edges required for (RC4). The adhesion sets U_0, U_i, U_{M+1} are depicted in brown.

for every two successive separations, there exists a linkage between its separators. By Menger's theorem (e.g. [41, Theorem 3.3.1]), this second property is equivalent to the absence of any separation of smaller order between two successive separations.

We start with a lemma about sequences of positive integers, which we will later apply to the order of the separations in the sequence. In what follows all sequences will be finite.

Lemma 8.5.2. *Let $n, m, p \geqslant 1$ be integers with $p \geqslant n^m$. Then every sequence $(a_1, \ldots, a_p)$ of integers of length p with $a_i \in \{0, \ldots, m-1\}$ for every $i \in [p]$ has a subsequence $(a_{i_1}, \ldots, a_{i_n})$ of length n such that*

(i) *$\ell := a_{i_1} = \cdots = a_{i_n}$, and*

(ii) *$a_j \geqslant \ell$ for all $i_1 \leqslant j \leqslant i_n$.*

Proof. We proceed by induction on m. For $m = 1$, we have $a_i = 0$ for all $1 \leqslant i \leqslant n \leqslant p$ which immediately yields the statement. So consider $m \geqslant 2$. If at least n of the a_i equal 0, then any n of them form the desired sequence. So suppose that at most $n' < n$ of the a_i equal 0 and let $i_1, \ldots, i_{n'}$ be the respective indices. Now consider the $n' + 1$ subsequences

$$(a_1, \ldots, a_{i_1 - 1}), (a_{i_1+1}, \ldots, a_{i_2-1}), \ldots, (a_{i_{n'}+1}, \ldots, a_p) \tag{8.1}$$

of consecutive $a_1, \ldots, a_p$; note that some of these subsequences may be empty. If each of these subsequences has length less than n^{m-1}, then we obtain a contradiction via

$$n^m \leqslant p \leqslant (n'+1)(n^{m-1}-1)+n' \leqslant n(n^{m-1}-1)+(n-1) = n^m - 1.$$

So one of the subsequences in (8.1), let us call it $(b_1,\dots,b_{p'})$, must have length at least $p' \geqslant n^{m-1}$, and we thus can apply the induction hypothesis to $(b_1-1,\dots,b_{p'}-1)$ with integers $n, m-1, p'$ to obtain a subsequence of $(b_1,\dots,b_{p'})$ which is as desired. □

A sequence $((A_i,B_i))_{i\in[p]}$ of (oriented) separations of a graph G is *strictly increasing* if $(A_i,B_i) < (A_j,B_j)$ for every two elements $i<j$ of $[p]$.

Lemma 8.5.3. *Let G be a graph and $n,m,p \geqslant 1$ be integers with $p \geqslant n^m$. If there is a strictly increasing sequence of length p in $\vec{S}_m(G)$, then there is also a strictly increasing sequence $((A_i,B_i))_{i\in[n]}$ of length n in $\vec{S}_m(G)$ such that*

(i) *$\ell := |A_1,B_1| = \dots = |A_n,B_n|$, and*

(ii) *for every separation $(A,B)\in\vec{S}_\ell(G)$ there is no $i\in[n-1]$ with $(A_i,B_i) < (A,B) < (A_{i+1},B_{i+1})$.*

Proof. With each strictly increasing sequence T in $\vec{S}_m(G)$, we associate a sequence $(n_i(T))_{i\in[m-1]}$ of integers where $n_i(T)$ denotes the number of separations in the strictly increasing sequence that have exactly order i. Whenever we will compare sequences of integers in this proof, we do so with respect to the lexicographic order. Let $\mathcal{T}$ denote the set of all strictly increasing sequence in $\vec{S}_m(G)$ of length at least p, and let $T = ((C_i,D_i))_{i\in[q]}$ be an element in $\mathcal{T}$ whose associated sequence of integers is maximal among all such sequences associated with elements in $\mathcal{T}$. By assumption there exists some such strictly increasing sequence of length at least p and, as the graph G is finite, there is a maximal sequence of integers among such associated to elements in $\mathcal{T}$.

Applying Lemma 8.5.2 to $(|C_i,D_i|)_{i\in[q]}$, we obtain a subsequence $((C_{i_j},D_{i_j}))_{j\in[n]}$ which we will show to be as desired. A subsequence of a strictly increasing sequence is again strictly increasing. It also satisfies (i), as all the separations in the subsequence have the same order ℓ by Lemma 8.5.2. It remains to show that (ii) holds as well. To do so, we show that if (ii) does not hold for $((C_{i_j},D_{i_j}))_{j\in[n]}$, then we find an element in $\mathcal{T}$ whose associated sequence of integers is larger than the one associated to $((C_i,D_i))_{i\in[q]}$, which contradicts our choice.

So suppose that there exists a separation $(A,B)\in\vec{S}_m(G)$ of order less than ℓ and an integer $j\in[n-1]$ such that $(C_{i_j},D_{i_j}) < (A,B) < (C_{i_{j+1}},D_{i_{j+1}})$. Consider the subsequence

$$((C_{i_j},D_{i_j})),(C_{i_j+1},D_{i_j+1}),\dots,(C_{i_{j+1}},D_{i_{j+1}}))$$

of $((C_i,D_i))_{i\in[p']}$; for notational simplicity, we also denote this subsequence by $R = ((A_i,B_i))_{i\in[r]}$.

Then we can obtain a new strictly increasing sequence $R' = ((A'_i, B'_i))_{i\in[r']}$ by removing duplicates, which will only appear consecutively, from the sequence

$$(A_1 \cap A, B_1 \cup B), \ldots, (A_q \cap A, B_q \cup B), (A, B), (A_1 \cup A, B_1 \cap B), \ldots, (A_r \cup A, B_r \cap B).$$

We now consider the strictly increasing sequence T' which is obtained from T by replacing its subsequence R with R'. As every element of the distributive lattice $\vec{U}(G)$ is uniquely determined by its infimum and supremum with any given other element in $\vec{U}(G)$, T' has length at least $q+1 \geqslant p$; thus, $T' \in \mathcal{T}$. Moreover, the sequence of integers associated to T' is larger than the one associated with T, which contradicts the choice of T. Indeed by the choice of R, all separations (C_g, D_g) with $i_j \leqslant g \leqslant i_{j+1}$ have order at least ℓ. Hence, $n_i(T') \geqslant n_i(T)$ for every $i < \ell$ and moreover $n_{|A,B|}(T') \geqslant n_{|A,B|}(T)+1$, since the new sequence T' additionally contains (A, B). $\square$

For the proof of Theorem 8.5.1 we also need the tangle-tree duality theorem, which we rephrase here in the version which we need later:

Theorem 8.5.4 (e.g. [41, Theorem 12.5.1])**.** *Every graph G with no $(k+1)$-tangle admits a tree-decomposition $(T, \mathcal{V})$ of adhesion at most k such that every two separations induced by distinct edges of T are distinct, and for every node $t \in T$ the set of (oriented) separations induced by the oriented edges (s,t) of T is in $\mathcal{T}$. In particular, every node of T has degree at most 3 and the width of $(T, \mathcal{V})$ is less than $3k$.*

The next lemma assets that in the absence of high-order tangles, the graph admits a certain type of linear decomposition. To prove it, we first apply the tangle-tree duality theorem (Theorem 8.5.4) to obtain a tree-decomposition of small width and then sort out a suitable linear decomposition by Lemma 8.5.3.

Lemma 8.5.5. *For every two integers $k, M \geqslant 1$, there exists an integer $N_1 = N_1(k, M) \geqslant 1$ such that if a graph G with more than N_1 vertices has no $(k+1)$-tangle, then there exists a linear decomposition $\mathcal{W}$ of G of length at least M and adhesion at most k such that $\mathcal{W}$ satisfies (R1).*

We remark that the proof shows that $N_1(k, M) = 3k \cdot 3^{(M+2)^{k+1}}$ suffices.

Proof of Lemma 8.5.5. We set $M_1 := M+2$, $M_2 := {M_1}^{k+1}$, $M_3 := 3^{M_2}$ and $N_1 := N_1(k, M) := 3kM_3$. Let G be a graph with more than N_1 vertices and no $(k+1)$-tangle. Since G has no $(k+1)$-tangle, it admits a tree-decomposition $(T, \mathcal{V})$ of width at most $3k$ and adhesion at most k such that T has maximum degree $\leqslant 3$. Moreover, we may choose $(T, \mathcal{V})$ such that all its induced separations are distinct. Then T contains at least $|G|/3k \geqslant N_1/3k \geqslant M_3$ vertices. Hence, T contains a path of length at least M_2, as all the nodes of T have degree at most 3.

Fix a path $P = p_0 \dots p_{M_2}$ in T, and let (A'_i, B'_i) be the separation of G induced by the oriented edge (p_{i-1}, p_i) of T for all $i \in [M_2]$. As $(T, \mathcal{V})$ has adhesion at most k, all these separations have order at most k. It is also immediate from the definition of inducing a separation that $(A'_i, B'_i) < (A'_{i+1}, B'_{i+1})$ for all $i \in [M_2 - 1]$. Thus, $((A'_i, B'_i))_{i \in [M_2]}$ is a strictly increasing sequence of length M_2 in $\vec{S}_k(G)$.

Hence, by Lemma 8.5.3, we obtain a new strictly increasing sequence $((A_i, B_i))_{i \in [M_1]}$ of length M_1 in $\vec{S}_k(G)$ whose elements all have the same order $\ell \leqslant k$ and such that there is no separation $(A, B) \in \vec{S}_\ell(G)$ with $(A_i, B_i) < (A, B) < (A_{i+1}, B_{i+1})$ for all $i \in [M_1 - 1]$. From this sequence, we now construct a linear decomposition $\mathcal{W} = (W_0, \dots, W_{M_1})$ via

$$W_0 := A_1, W_i := B_i \cap A_{i+1} \text{ for } i \in [M_1 - 1], \text{ and } W_{M_1} := B_{M_1}.$$

As we have noted above for tree-decompositions, $\mathcal{W}$ indeed satisfies (L1) and (L2). Note that the adhesion set U_i equals $W_{i-1} \cap W_i = A_i \cap B_i$; thus, (L3) holds as well. Moreover, $\mathcal{W}$ has length M_1 and adhesion $\ell \leqslant k$.

Before we prove (L4), let us check that $\mathcal{W}$ satisfies (R1). By Menger's theorem (e.g. [41, Theorem 3.3.1]), it suffices to show that for $i \in [M_1 - 1]$, the part $G[W_i]$ contains no separation (C', D') of order less than ℓ with $U_i \subseteq C'$ and $U_{i+1} \subseteq D'$. Suppose for a contradiction that such a separation exists for some $i \in [M_1 - 1]$. Since $\mathcal{W}$ satisfies (L1) and (L2), the ordered pair

$$(C, D) := (C' \cup \bigcup_{j<i} W_j, D' \cup \bigcup_{j>i} W_j)$$

is a separation of G and has the same order as (C', D'); in particular, its order is less than ℓ. The construction of $\mathcal{W}$ ensures that for $i \in [M_1]$ we have

$$(A_i, B_i) = (\bigcup_{j<i} W_j, \bigcup_{j \geqslant i} W_j).$$

Hence, $(A_i, B_i) < (C, D) < (A_{i+1}, B_{i+1})$, which contradicts our choice of the (A_j, B_j) via Lemma 8.5.3.

We finally turn to (L4). If $W_{i-1} \subseteq W_i$ for some $i \in \{2, \dots, M_1\}$, then $W_{i-1} \subseteq A_i \cap B_i$, as we have noted above that $W_i \subseteq A_i \cap B_i$. Since $\ell = |A_{i-1}, B_{i-1}| = |A_i, B_i|$, we thus have $A_{i-1} \cap B_{i-1} = W_{i-1} = A_i \cap B_i$. Together with $W_{i-1} = B_{i-1} \cap A_i$, this implies $(A_{i-1}, B_{i-1}) = (A_i, B_i)$, which contradicts that the (A_i, B_i) are distinct by choice. A symmetrical argument shows that $W_{i-1} \not\supseteq W_i$ for all $i \in \{1, \dots, M_1 - 1\}$. Thus, we have $W_i \neq W_i \cap W_{i+1} \neq W_{i+1}$ for all $i \in \{1, \dots, M_1 - 2\}$. However, we still might have $W_0 \subseteq W_1$ and $W_{M_1 - 1} \supseteq W_{M_1}$. In any of these cases, we remove W_0 or W_{M_1}, respectively, from $\mathcal{W}$. This operation does not affect any of the properties of $\mathcal{W}$, except that its length might decrease by at most 2; however, we have $M_1 = M + 2$, so the length of the obtained linear decomposition of G is still at least M_1, as desired. This completes the proof. □

We will now transform a linear decomposition such as the one in Lemma 8.5.5 into the desired rainbow-cloud-decomposition. As an intermediate step, we find a linear decomposition $\mathcal{W} = (W_0, W_1, \ldots, W_M)$ with a foundational linkage $\mathcal{P}$ which satisfies two additional properties:[4]

(FL1) For every $P \in \mathcal{P}$, if there exists $i \in [M-1]$ such that $P[W_i] = P \cap G[W_i]$ is a trivial path, then $P[W_i]$ is a trivial path for all $i \in [M-1]$.

(FL2) For every two distinct $P, P' \in \mathcal{P}$, if there exists $i \in [M-1]$ such that there is path in $G[W_i]$ with one endvertex in P and the other in P' and whose internal vertices avoid every path in $\mathcal{P}$, then this holds for every $i \in [M-1]$.

The next lemma yields that the existence of a linear decomposition of some suitable length whose foundational linkage satisfies (FL1) and (FL2) is ensured by a long enough linear decomposition as in Lemma 8.5.5.

Lemma 8.5.6 ([92, Lemma 3.5][5])**.** *For every two integers $M \geqslant 1$ and $\ell \geqslant 0$, there exists an integer $M_1 = M_1(\ell, M) \geqslant 1$ such that if a linear-decomposition $\mathcal{W} = (W_0, \ldots, W_{M_1})$ of a graph G has length M_1, adhesion ℓ and pairwise distinct W_i, and $\mathcal{W}$ satisfies (R1), then G has a linear decomposition $\mathcal{W}'$ of length at least M which additionally has a foundational linkage satisfying (FL1) and (FL2).*

We remark that the proof of Lemma 8.5.6 in [92] shows that $M_1(\ell, M) = (M^{\binom{\ell}{2}+1} \cdot \binom{\ell}{2}!)^{\ell+1} \cdot \ell!$ suffices.

We now turn to the proof of Theorem 8.5.1. For this, we use the previous lemmas to find a linear decomposition with strong structural properties, which we then refine into the desired rainbow-cloud-decomposition.

Proof of Theorem 8.5.1. Set $N(k, M) := N_1(k, M_1(k, M+2))$, where N_1 and M_1 are as in Lemmas 8.5.5 and 8.5.6, respectively. Then Lemma 8.5.5 yields a linear-decomposition of G of length at least $M_1(k, M+2)$ and adhesion $\ell \leqslant k$ which satisfies (R1); by merging the first i bags for some suitable integer i, we may assume that this linear decomposition has length exactly $M_1(\ell, M+2) \leqslant M_1(k, M+2)$. By Lemma 8.5.6 ,there is a linear-decomposition $\mathcal{W}''$ of G of length $M' \geqslant M + 2$ which additionally has a foundational linkage $\mathcal{P}''$ satisfying (FL1) and (FL2). We note that $\ell > 0$ since G is a connected graph and $M' \geqslant 1$.

Consider the adhesion sets U''_i of $\mathcal{W}''$ and let $Z' := \bigcap_{i \in \{1,\ldots,M'\}} U''_i$. By (FL1), the set Z' consists precisely of all trivial paths in $\mathcal{P}''$. Now $\mathcal{W}''$ induces the linear decomposition $\mathcal{W}'$ of $G - Z$ by

[4](FL1) and (FL2) are (L7) and (L8) in [92], respectively.

[5]More precisely, it follows from the proof of [92, Lemma 3.5], as (L6) is only assumed to guarantee that the resulting decomposition also satisfies (L6).

setting $W_i' := W_i'' \setminus Z'$ together with its foundational linkage $\mathcal{P}' \subseteq \mathcal{P}''$. Now the adhesion ℓ' of $\mathcal{W}'$ satisfies $\ell = \ell' + |Z'| > 0$; in particular, we might have $\ell' = 0$. Clearly, $\mathcal{P}''$ again satisfies (FL1) and (FL2). Moreover, $\mathcal{W}'$ still satisfies (R1), and it now also satisfies (R3) by (FL1). Finally, (FL2) implies that if $z' \in Z'$ is adjacent to some W_i in G, then it is adjacent to all W_i.

We now define an auxiliary graph $H_{\mathcal{P}'}$ with vertex set $\mathcal{P}'$ and an edge joining two distinct paths $P, P' \in \mathcal{P}'$ if there exists a path in some, and by (FL2) hence every, $G[W_i']$ with one endvertex in P and the other in P' and whose internal vertices avoid every path in $\mathcal{P}'$. Let C_H be an arbitrary component of $H_{\mathcal{P}'}$. For $i \in \{1, \dots, M'-1\} = [M'-1]$ we then let W_i be the vertex set of the component of $G[W_i']$ which contains all $V(P) \cap W_i'$ for $P \in C_H$; note that this is well-defined by the construction of $H_{\mathcal{P}'}$. If C_H, and thus $H_{\mathcal{P}'}$, is empty, we let $Z := Z'$; otherwise, we let Z consist of all those $z \in Z'$ which are adjacent to some W_i, and hence all W_i by (FL2). We further set

$$R := G\left[\bigcup_{i \in [M'-1]} W_i\right] \text{ and}$$

$$C := G\left[Z' \cup \bigcup_{i \in [M'-1]} (W_i' \setminus W_i) \cup W_0' \cup W_{M'}'\right].$$

This ensures that $Z \subseteq V(C)$ as well as $G = G[V(R) \cup Z] \cup C$, as we chose the W_i as components of the $G[W_i']$. We now claim that $(R, \mathcal{W}, Z, C)$ where $\mathcal{W} := (W_1, \dots, W_{M'-1})$ is the desired rainbow-cloud-decomposition of G.

Clearly, $\mathcal{W}$ is a linear decomposition of R of adhesion $|C_H|$ and length $M' - 2 \geqslant M$; note that $|C_H| > 0$ if $\ell' > 0$ and hence $|Z| + |C_H| > 0$ by the choice of Z. We now verify that $\mathcal{W}$ is even a rainbow-decomposition of R. The foundational linkage $\mathcal{P}'$ of $\mathcal{W}'$ induces a foundational linkage of $\mathcal{W}$, as we have chosen a component C_H of $H_{\mathcal{P}'}$; thus, (R1) holds. For (R2), note that the W_i are components of the $G[W_i']$ and hence connected by construction. Finally, (R3) transfers from $\mathcal{W}'$.

It remains to check (RC1) to (RC4). So let $U_1 := V(C) \cap W_1$ and $U_{M'} := V(C) \cap W_{M'-1}$; note that our construction implies $U_1 = W_0' \cap W_1$ and $U_{M'} = W_{M'}' \cap W_{M'-1}$. By definition, we have $V(R) \cap V(C) = (W_1 \cup W_{M'-1}) \cap V(C) = U_1 \cup U_{M'}$, so (RC1) holds, and (R3) for $\mathcal{W}'$ implies (RC2) as well as (R1) implies (RC3). For (RC4), we recall that this holds by the definition of Z. This completes the proof. $\square$

8.6 RC-decompositions and separations

In this section, we investigate how a fixed RC-decomposition of a graph G interacts with the separations of G. On the one hand, we will analyse in what ways a separation of G may meet the rainbow of an RC-decomposition. On the other hand, we look at the separations of G induced by the structure of the RC-decomposition. With this we build a set of tools that we will later apply in the proof of our structural main result, Theorem 25. These tools will allow us to control the new low-order separations that arise when we delete an edge deep inside the rainbow.

8.6.1 Separations and bags

In this subsection, we prove two general lemmas that describe in which ways a separation of a graph meets the bags of a linear decomposition or an RC-decomposition. The first lemma asserts that if a linear decomposition satisfies (R2) and (R3), then the strict sides of a separation contain most of its bags.

Lemma 8.6.1. *Let $\{A, B\}$ be a separation of a graph G of order k and let $\mathcal{W}$ be a linear decomposition of a subgraph R of G that satisfies (R3). Then the separator $A \cap B$ meets at most $2k$ bags of $\mathcal{W}$. Moreover, if $\mathcal{W}$ additionally satisfies (R2), then at most $2k$ bags of $\mathcal{W}$ are not contained in either $A \setminus B$ or $B \setminus A$.*

Proof. Each vertex in R is contained in at most one adhesion set of $\mathcal{W}$ by (R3) and thus in at most two bags of $\mathcal{W}$. So since $\{A, B\}$ has order k, the separator $A \cap B$ meets at most $2k$ bags of $\mathcal{W}$. As each part $R[W_i]$ of $\mathcal{W}$ is connected by (R2), every bag of $\mathcal{W}$ that is disjoint from $A \cap B$ is included in precisely one of $A \setminus B$ and $B \setminus A$. Hence, at most $2k$ bags of $\mathcal{W}$ are not contained in either $A \setminus B$ or $B \setminus A$. □

The second lemma states that if a strict side of a separation has a component which meets both the cloud and a bag W_i of a fixed RC-decomposition, then it also contains most bags of the rainbow along one of the two connections from W_i to the cloud.

Lemma 8.6.2. *Let $(R, \mathcal{W}, Z, C)$ be an RC-decomposition of a graph G of length M, and let $\{A, B\}$ be a separation of G of order k. Suppose that $D \subseteq G[A \setminus B]$ is a component of $G - (A \cap B)$ which meets both C and some bag W_i. Then D contains all but at most $2k$ of $W_0, \dots, W_i$ or $W_i, \dots, W_M$. Moreover, D meets either Z or every bag of $W_0, \dots, W_i$ or of $W_i, \dots, W_M$, respectively.*

Proof. By Lemma 8.6.1 applied to the restriction $\{A \cap V(R), B \cap V(R)\}$ of $\{A, B\}$ to the rainbow R at most $2k$ of the bags W_i meet $A \cap B$. If D contains a vertex z of Z, then every bag which does not meet $A \cap B$ is contained in D by (RC4) and (R2) because D is a component of $G - (A \cap B)$, as desired. So let us now assume that $V(D)$ is disjoint from Z. Since D is connected and meets both C and W_i, the component D contains a path P from W_i to C. This path P has to first enter C either through U_0 or U_{M+1}, as $V(D) \cap Z = \emptyset$ and also $G[V(R) \cup Z] \cup C = G$ and $V(R) \cap V(C) = U_0 \cup U_{M+1}$ by the definition of an RC-decomposition. Thus, we may assume that the only vertex of P in C is its endvertex which is in U_0, as the other case is symmetrical. Hence, it meets all bags W_j with $0 \leqslant j \leqslant i$, because $\mathcal{W}$ is a linear decomposition of R. Thus, every bag W_j with $0 \leqslant j \leqslant i$ which avoids $A \cap B$ is contained in the component D of $G - (A \cap B)$ by (R2), as desired. □

8.6.2 Separating along the rainbow

We can easily turn an RC-decomposition $(R, \mathcal{W}, Z, C)$ of a graph G of length M into a new one of shorter length and of the same adhesion by restricting the linear decomposition $\mathcal{W}$ of R to some interval in $\{0, \ldots, M\}$ and adding the remaining bags to C as follows. For $0 \leqslant i \leqslant j \leqslant M$, we set

$$\begin{aligned} R_{i,j} &:= G\left[\bigcup_{h=i}^{j} W_h\right] = \bigcup_{h=i}^{j} G[W_h], \\ \mathcal{W}_{i,j} &:= (W_i, \ldots, W_j), \\ C_{i,j} &:= G\left[V(C) \cup \left(\bigcup_{h=0}^{i-1} W_h\right) \cup \left(\bigcup_{h=j+1}^{M} W_h\right)\right], \text{ and} \\ (R, \mathcal{W}, Z, C)_{i,j} &:= (R_{i,j}, \mathcal{W}_{i,j}, Z, C_{i,j}). \end{aligned}$$

We remark that we follow the convention that an empty union, such as $\bigcup_{h=0}^{-1} W_h$ or $\bigcup_{h=M+1}^{M} W_h$, is the empty set. It is straight-forward from the definition that $(R, \mathcal{W}, Z, C)_{i,j}$ is again an *RC-decomposition* of G:

Lemma 8.6.3. *Let G be a graph with an RC-decomposition $(R, \mathcal{W}, Z, C)$ of length M and of adhesion ℓ. Then $(R, \mathcal{W}, Z, C)_{i,j}$ is an RC-decomposition of G of length $j - i$ and of adhesion ℓ for all $0 \leqslant i \leqslant j \leqslant M$.* □

It follows immediately that we obtain separations of G along the rainbow of a fixed RC-decomposition:

Lemma 8.6.4. *Let G be a graph with an RC-decomposition $(R, \mathcal{W}, Z, C)$ of length M and adhesion ℓ, and let $0 \leqslant i \leqslant j \leqslant M$. Then $\{V(R_{i,j}) \cup Z, V(C_{i,j})\}$ is a separation of G with separator $U_i \cup U_{j+1} \cup Z$. In particular, its order is $2\ell + |Z|$.* □

8.6.3 Rainbow-crossing separations

We now investigate separations that 'cross' the rainbow of a given RC-decomposition in that both strict sides of the separation contain bags of the linear decomposition of the rainbow: Let $(R, \mathcal{W}, Z, C)$ be an RC-decomposition of a graph G of length M. An oriented separation (A, B) of G of order k *crosses the rainbow R clockwise* if there exist integers $i \in \{0, \ldots, 2k\}$ and

$j \in \{M-2k, \dots, M\}$ such that $W_i \subseteq A \setminus B$ and $W_j \subseteq B \setminus A$. We denote the minimal such i by $i_{A,B}$ and the maximal such j by $j_{A,B}$. If an oriented separation (A, B) crosses the rainbow clockwise, we say that its other orientation (B, A) *crosses the rainbow R counter-clockwise* and the underlying unoriented separation $\{A, B\}$ *crosses the rainbow R*. We remark that a separation $\{A, B\}$ which crosses the rainbow contains Z in its separator $A \cap B$ by (RC4), $W_i \subseteq A \setminus B$ and $W_j \subseteq B \setminus A$. Note that in the above definition, we may have $j \leqslant i$, since we do not assume any lower bound on M. In most applications, however, we will have $M \geqslant 4k$ and thus $i < j$.

Given a separation which crosses the rainbow, we may construct several separations that separate the cloud in the same way, but split the rainbow along the adhesion sets of its linear decomposition (see Figure 8.3): More precisely, let $(R, \mathcal{W}, Z, C)$ be an RC-decomposition of a

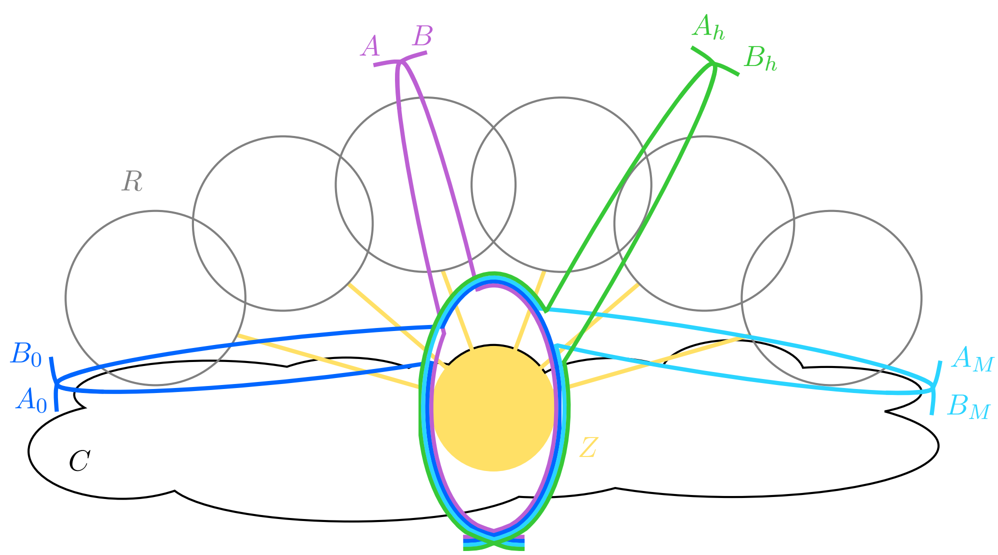

FIGURE 8.3: A rainbow-crossing separation (A, B) and the arising separations (A^0, B^0), (A^h, B^h) and (A^M, B^M).

graph G. For a separation (A, B) of G of order k that crosses the rainbow clockwise, we let $i := i_{A,B} \in \{0, \dots, 2k\}$ be the smallest and $j := j_{A,B} \in \{M-2k, \dots, M\}$ the largest number such that $W_i \subseteq A \setminus B$ and $W_j \subseteq B \setminus A$, which are well-defined by the definition of crossing the rainbow clockwise. For $h \in \{i+1, \dots, j\}$ we set

$$(A^h, B^h) := ((A \cap V(C_{i,j})) \cup V(R_{i,h-1}),\ V(R_{h,j}) \cup (B \cap V(C_{i,j}))).$$

If (A, B) crosses the rainbow counter-clockwise, we use the fact that (B, A) crosses the rainbow clockwise to define (A^h, B^h) accordingly.

It is immediate from the definition that $(A^{i+1}, B^{i+1}) \leqslant (A, B) \leqslant (A^j, B^j)$ and that the (A^h, B^h) form a strictly increasing sequence. Indeed, the (A^h, B^h) also are separations of G:

Lemma 8.6.5. *Let $(R, \mathcal{W}, Z, C)$ be an RC-decomposition of a graph G of adhesion ℓ, and let (A, B) be a separation of G of order k which crosses the rainbow R clockwise. Then (A^h, B^h) is a separation of G of order k which crosses the rainbow R clockwise for every $h \in \{i_{A,B} + 1, \dots, j_{A,B}\}$.*

Proof. We abbreviate $i := i_{A,B}$ and $j := j_{A,B}$. As $C_{i,j}$ is a subgraph of G, the separation $\{A, B\}$ of G induces the separation $\{A \cap V(C_{i,j}), B \cap V(C_{i,j})\}$ of $C_{i,j}$. Note that $Z \subseteq A \cap B$, since Z is contained in the neighbourhood of both $W_i \subseteq A \setminus B$ and $W_j \subseteq B \setminus A$ by (RC4). Moreover, by Lemma 8.6.4, both $U_i \cup Z \cup U_h$ and $U_h \cup Z \cup U_{j+1}$ separate $R_{i,h-1}$ and $R_{h,j}$, respectively, from the rest of G. As $U_i \cup Z \cup U_h \subseteq A^h$ and $U_h \cup Z \cup U_{j+1} \subseteq B^h$, it follows that $\{A^h, B^h\}$ is a separation of G.

It remains to show that $|A^h \cap B^h| \leqslant k$. The foundational linkage of the RC-decomposition induces a U_{i+1}–U_j linkage of cardinality ℓ in $R_{i+1,j-1}$. Since $\{A, B\}$ is a separation of G and $U_{i+1} \subseteq W_i \subseteq A \setminus B$ and $U_j \subseteq W_j \subseteq B \setminus A$, the separator $A \cap B$ has to contain one vertex of every U_{i+1}–U_j path. But since $V(C_{i,j}) \cap V(R_{i+1,j-1}) = \emptyset$, these ℓ vertices are contained in $(A \cap B) \setminus V(C_{i,j})$. Hence, we obtain $|(A \cap B) \cap V(C_{i,j})| \leqslant |A \cap B| - \ell = k - \ell$. With this, we may now bound the order of $\{A^h, B^h\}$:

$$
\begin{aligned}
|A^h \cap B^h| &= |((A \cap V(C_{i,j})) \cup V(R_{i,h-1})) \cap (V(R_{h,j}) \cup (B \cap V(C_{i,j})))| \\
&\leqslant |(A \cap B) \cap V(C_{i,j})| + |A \cap (V(C_{i,j}) \cap V(R_{h,j}))| \\
&\qquad\qquad + |B \cap (V(C_{i,j}) \cap V(R_{i,h-1}))| + |V(R_{i,h-1}) \cap V(R_{h,j})| \\
&\leqslant (k - \ell) + |A \cap W_j| + |B \cap W_i| + |U_h| = (k - \ell) + 0 + 0 + \ell = k,
\end{aligned}
$$

where the penultimate equation holds because $W_i \subseteq A \setminus B$ and $W_j \subseteq B \setminus A$. This completes the proof. □

8.6.4 Rainbow-slicing separations

In this section, we study separations that 'slice' the rainbow in that they cut bags in the middle of the rainbow away.

Let $(R, \mathcal{W}, Z, C)$ be an RC-decomposition of a graph G of length M. A separation $\{A, B\}$ of G of order k *slices the rainbow* R if there are integers $i < h < j$ with $i \in \{0, \dots, 2k\}$ and $j \in \{M - 2k, \dots, M\}$ such that $W_i, W_j \subseteq A \setminus B$ and $W_h \subseteq B \setminus A$ or such that $W_i, W_j \subseteq B \setminus A$ and $W_h \subseteq A \setminus B$.

We now prove a lower bound on the order of separations that slice the rainbow. Towards this, we first show such a lower bound for a slightly more general class of separations for later use, which in particular yields the desired bound for rainbow-slicing separations.

Lemma 8.6.6. *Let G be a graph with an RC-decomposition $(R, \mathcal{W}, Z, C)$ of adhesion ℓ, and let $\{A, B\}$ be a separation of G. Suppose that there are integers $0 \leqslant i < h < j \leqslant M+1$ such that $B \setminus A$ contains W_h and A contains the adhesion sets U_i and U_j. Then $|A \cap B \cap V(R)| \geqslant 2\ell$. Moreover, if additionally $Z \subseteq A$, then $|A \cap B| \geqslant 2\ell + |Z|$.*

Proof. The fundamental linkage induces a U_i–U_h linkage $\mathcal{P}_i$ and a U_h–U_j linkage $\mathcal{P}_j$; both have cardinality ℓ and are in R. We remark that the paths in $\mathcal{P}_i$ meet the paths in $\mathcal{P}_j$ only in U_h. Hence, $W_h \subseteq B \setminus A$ and $U_i, U_j \subseteq A$ yields that the separator $A \cap B$ contains at least one vertex of each of these 2ℓ paths. As every two paths in $\mathcal{P}_i$ and also every two paths in $\mathcal{P}_j$ are pairwise disjoint, we thus have $|A \cap B \cap V(R)| \geqslant 2\ell$.

If additionally $Z \subseteq A$, then $Z \subseteq A \cap B$ since $Z \subseteq N_G(W_h)$ by (RC4) and $W_h \subseteq B \setminus A$. As all the above 2ℓ paths lie in R and hence avoid Z, the claim follows. □

Since every rainbow-slicing separation contains Z in its separator by (RC4), we obtain the following lower bound on their order:

Corollary 8.6.7. *Let G be a graph with an RC-decomposition $(R, \mathcal{W}, Z, C)$ of adhesion ℓ. If a separation $\{A, B\}$ of G slices the rainbow R, then $|A \cap B \cap V(R)| \geqslant 2\ell$ and $|A \cap B| \geqslant 2\ell + |Z|$.* □

We conclude this section of preparatory lemmas by showing that every separation which separates two bags of an RC-decomposition either crosses or slices its rainbow.

Lemma 8.6.8. *Let G be a graph with an RC-decomposition $(R, \mathcal{W}, Z, C)$ of length M, and let $\{A, B\}$ be a separation of G. If there are bags W_i and W_j of $\mathcal{W}$ such that $W_i \subseteq A \setminus B$ and $W_j \subseteq B \setminus A$, then $\{A, B\}$ either crosses or slices the rainbow.*

Proof. Let $k := |A, B|$. If $M \leqslant 2k$, then W_i and W_j witness that $\{A, B\}$ crosses the rainbow. So we may assume $M \geqslant 2k$. By Lemma 8.6.1, all but at most $2k$ bags of $\mathcal{W}$ are contained in either $A \setminus B$ or $B \setminus A$. In particular, there exist bags W_h with $h \in \{0, \dots, 2k\}$ and W_s with $s \in \{M-2k, \dots, M\}$ that are each contained in either $A \setminus B$ or $B \setminus A$.

If one of W_h and W_s is contained in $A \setminus B$ and the other one in $B \setminus A$, then $\{A, B\}$ crosses the rainbow. Otherwise, both W_h and W_s are contained in the same side of $\{A, B\}$. By symmetry, we may assume $W_h, W_s \subseteq A \setminus B$. Now if $j \in \{0, \dots, 2k\}$, then W_j and W_s witness that (B, A) crosses the rainbow clockwise, and if $j \in \{M-2k, \dots, M\}$, then W_h and W_j witness that (A, B) crosses the rainbow clockwise. Otherwise, $j \in \{2k+1, \dots, M-2k-1\}$; in particular $h < j < s$. Therefore $\{A, B\}$ slices the rainbow. □

8.7 Proof of Theorem 25 if G has no higher-order tangle

In this section, we complete the proof of Theorem 25. Based on the results in Sections 8.3–8.5, it suffices to consider graphs that admit an RC-decomposition with certain properties. More precisely, the main result of this section, which will allow us to finish the proof of Theorem 25, reads as follows:

Theorem 8.7.1. *Let $k \geqslant 1$ be an integer, let G be a graph of minimum degree at least 3, and let τ be a k-tangle in G. Suppose that G admits an RC-decomposition with sun Z which has length $\geqslant 18k$ and adhesion ℓ such that $\ell + |Z| \geqslant 1$. Then there exists an edge $e \in E(G)$ such that τ extends to a k-tangle τ' in $G - e$.*

Most of this section is devoted to the proof of Theorem 8.7.1. In the very end we combine the previous results to prove Theorem 25.

8.7.1 Living in the rainbow

In the proof of Theorem 8.7.1, we will delete an edge e deep inside the rainbow R of the given long RC-decomposition. We will then make use of the regular structure of R to orient the newly arising separations in $G' := G - e$ in such a way that we find a k-tangle τ' in G' to which τ extends. As we want to orient these new separations of G' in line with τ, the tangle τ should ideally be 'regular' or 'monotonic' along R. Then we could use this monotonicity of τ along R to orient the newly arising separations in G' by consistency.

In this section, we extract from a given RC-decomposition another RC-decomposition such that τ has the desired monotonic behaviour along the rainbow. More precisely, the tangle τ will 'point away' from the rainbow, which will turn out to be equivalent to the desired monotonicity, as we will see below. Formally, we enclose this in a definition of when a tangle τ does not 'live in the rainbow R' of an RC-decomposition. The intuition behind this definition is as follows.

Clearly, a tangle τ should live in the rainbow R, if there is a separation $(A, B) \in \tau$ whose strict big side $B \setminus A$ is contained in the rainbow R and does not meet the cloud C. For our proof of Theorem 8.7.1, this case is not the only relevant one: we also need to regard τ as living in R if it orients two separations in such a way that they point towards each other and to a piece of R. It turns out that for this second case, we can even restrict our attention to rainbow-crossing separations, as follows.

Let τ be a k-tangle in a graph G. Consider a given RC-decomposition $(R, \mathcal{W}, Z, C)$ of G and a separation (A, B) of G of order $k' < k$ that crosses the rainbow R clockwise or counter-clockwise. Then τ *orients $\{A, B\}$ monotonically over the rainbow* if τ orients all the $\{A^h, B^h\}$ with $h \in \{i_{A,B} + 1, \ldots, j_{A,B}\}$ such that they are pairwise comparable, i.e. either all as (A^h, B^h)

or all as (B^h, A^h). We remark that if τ does not orient a rainbow-crossing separation $\{A, B\}$ of order less than k monotonically over the rainbow R, then the consistency of τ ensures that there is a (unique) index $h^* \in \{0, \dots, M-1\}$ such that $(A^{h^*}, B^{h^*}) \in \tau$ and $(B^{h^*+1}, A^{h^*+1}) \in \tau$, as the (A^h, B^h) form an increasing or decreasing sequence as (A, B) crosses the rainbow clockwise or counter-clockwise, respectively. We then call h^* its *turning point*; note that $h^* \in \{0, \dots, M-1\}$. Moreover, we say that τ *lives in the rainbow* R if at least one of the following holds:

(LR1) there is a separation $(A, B) \in \tau$ such that $B \setminus A \subseteq V(R) \setminus V(C)$, or

(LR2) τ orients at least one rainbow-crossing separation of order less than k not monotonically over the rainbow.

The next theorem asserts that we can shorten an RC-decomposition of sufficient length so that a given tangle does not live in the shortened rainbow:

Theorem 8.7.2. *Let $k \geqslant 1$, $M \geqslant 6k$ and $\ell \geqslant 0$ be integers, and let τ be a k-tangle in a graph G. If G admits an RC-decomposition $(R, \mathcal{W}, Z, C)$ of G of length at least M and of adhesion ℓ, then there exist $0 \leqslant i < j \leqslant M$ with $j - i \geqslant M/2 - k$ such that τ does not live in the rainbow of the RC-decomposition $(R, \mathcal{W}, Z, C)_{i,j}$.*

Proof. If τ does not live in R, then we are done by setting $i := 0$ and $j := M$; so suppose that τ lives in R. By definition, there are two possible reasons for that, (LR1) and (LR2), which we will treat separately.

First, suppose that τ lives in R because of (LR1), i.e. there is a separation $(A, B) \in \tau$ with $B \setminus A \subseteq V(R \setminus C)$. Then there also exists such a separation in τ with an even stronger property:

Claim 1. *There is a separation $(X, Y) \in \tau$ such that $Y \setminus X \subseteq \bigcup_{t=r}^{s} W_t$ for integers $r \leqslant s$ with $s - r < 2k - 2$.*

Proof. First, suppose that $2\ell + |Z| < k$. Then $\{V(C), V(R) \cup Z\}$ is a separation of G of order $2\ell + |Z| < k$ and hence the k-tangle τ has to contain one of its orientations. By the consistency of the tangle τ, $(V(C), V(R) \cup Z) \leqslant (A, B) \in \tau$ yields $(V(C), V(R) \cup Z) \in \tau$. Every 'slice' $\{V(R_{h,h}) \cup Z, V(C_{h,h})\}$ of the rainbow R with $h \in \{0, \dots, M\}$ is also a separation of G of order $2\ell + |Z| < k$ by Lemma 8.6.4.

Suppose for a contradiction that they are all oriented as $(V(R_{h,h}) \cup Z, V(C_{h,h}))$ by τ. By iteratively using the fact that the tangle τ avoids the triples in $\mathcal{T}$ and for every $0 \leqslant h \leqslant i \leqslant j \leqslant M$ the separations $(V(R_{h,i} \cup Z, V(C_{h,i}))$, $(V(R_{i,j}) \cup Z, V(C_{i,j}))$ and $(V(C_{h,j}), V(R_{h,j} \cup Z)\}$ of order $2\ell + |Z|$ form a triple in $\mathcal{T}$, we obtain that $(V(R_{0,M} \cup Z, V(C_{0,M})) = (V(R) \cup Z, V(C))$ is in τ, which contradicts that its other orientation is also in τ. Therefore, there exists $h \in \{0, \dots, M\}$ with $(V(C) \cup (V(R) \setminus W_h) \cup U_h \cup U_{h+1}, W_h \cup Z) \in \tau$, which is the desired separation (X, Y) with $r := h =: s$.

Secondly, suppose that $2\ell + |Z| \geqslant k$, and let $(A, B) \in \tau$. We may assume that the separation $(A, B) \in \tau$ with $B \smallsetminus A \subseteq V(R) \smallsetminus V(C)$ is $\leqslant$-maximal in τ with that property; in particular, $G[B \smallsetminus A]$ is connected. Then $B \smallsetminus A$ cannot contain a bag of $\mathcal{W}$ by Lemma 8.6.6, since $B \smallsetminus A \subseteq V(R) \smallsetminus V(C)$ yields that $U_0, U_{M+1}, Z \subseteq V(C) \subseteq A$, and $\{A, B\}$ has order $< k \leqslant 2\ell + |Z|$. Thus, every bag which meets $B \smallsetminus A$ also meets the separator $A \cap B$. Thus, Lemma 8.6.1 ensures that at most $2k - 2$ bags of $\mathcal{W}$ meet $B \smallsetminus A$. Further, since $G[B \smallsetminus A] \subseteq R$ is connected and each bag W_i separates R, the bags which are met by $B \smallsetminus A$ are consecutive bags of $\mathcal{W}$. Thus, $(X, Y) := (A, B)$ is the desired separation, as witnessed by the respective indices r and s of the first and last bag of $\mathcal{W}$ which are met by $B \smallsetminus A$. ■

Let $(X, Y) \in \tau$ and $r \leqslant s$ with $s - r < 2k - 2$ be given by Claim 1. Now if $r > \frac{M}{2} - k$, then we set $i := 0$ and $j := r - 1$, and otherwise, if $r \leqslant \frac{M}{2} - k$, then we set $i := s + 1$ and $j := M$. We remark that in the latter case $s < \frac{M}{2} + k$. By Lemma 8.6.3, $(R, \mathcal{W}, Z, C)_{i,j}$ is an RC-decomposition of G of adhesion ℓ and length $j - i \geqslant \frac{M}{2} - k$. We claim that τ does not live in $R_{i,j}$, as desired.

Since $Y \smallsetminus X \subseteq \bigcup_{t=r}^{s} W_t$ and hence $Y \smallsetminus X$ is disjoint from $R_{i,j}$ by definition of i, j, we have $V(R_{i,j}) \subseteq X$. Thus, every two separations $(X_1, Y_1), (X_2, Y_2)$ of G with $(Y_1 \cap Y_2) \smallsetminus (X_1 \cup X_2) \subseteq V(R_{i,j})$ form a forbidden triple together with (X, Y). But every separation as in (LR1) (taken as both (X_i, Y_i)) as well as every pair of separations witnessing the non-monotonicity in (LR2) are such two separations. Therefore, τ cannot live in $R_{i,j}$, as desired.

Secondly, suppose that τ does live in the rainbow R because of (LR2), i.e. there is a separation of G of order $< k$ that crosses the rainbow R and τ does not orient it monotonically over the rainbow R. It turns out that all such separations $\{A, B\}$ have the same turning point:

Claim 2. *If a rainbow-crossing separation $\{A, B\}$ which τ does not orient monotonically over the rainbow R has turning point h^*, then the turning point of every such separation is h^*.*

Proof. Suppose for a contradiction that there is another rainbow-crossing separation $\{X, Y\}$ which τ does not orient monotonically over the rainbow R and has turning point $h' \neq h^*$. Let (A, B) and (X, Y) be the orientations which cross the rainbow clockwise. By possibly interchanging $\{A, B\}$ and $\{X, Y\}$, we may assume $h^* < h'$. Now the definition of the (X^h, Y^h) immediately yields that $W_{h^*} \subseteq X^{h'}$. But it also guarantees that $G = G[A^{h^*}] \cup G[B^{h^*+1}] \cup G[W_{h^*}]$. Thus $(X^{h'}, Y^{h'})$ together with (A^{h^*}, B^{h^*}) and (B^{h^*+1}, A^{h^*+1}) forms a forbidden triple in τ, which is a contradiction. ■

If $h^* \geqslant M/2$, then we set $i = 0$ and $j = h^* - 1$, and if $h^* < M/2$, then we set $i = h^* + 1$ and $j = M$. By Lemma 8.6.3, $(R, \mathcal{W}, Z, C)_{i,j}$ then is an RC-decomposition of G of length $j - i \geqslant M/2 - 1$ and adhesion ℓ. We claim that τ does not live in $R_{i,j}$. Since τ does not live in R because of (LR1) by assumption, τ does neither live in $R_{i,j} \subseteq R$ because of (LR1), either. Moreover, τ can also not live in $R_{i,j}$ because of (LR2): Indeed, every separation that crosses $R_{i,j}$ also crosses R, and if τ does not orient it monotonically over $R_{i,j}$, it does not do so over R, as well. But since all such separations have turning point h^* with respect to R by Claim 2, the choice of i and j ensure that τ orients them all monotonically over $R_{i,j}$ by the choice of i and j, as desired. □

8.7.2 Deleting an edge and orienting the new separations

This subsection is dedicated to the proof of Theorem 8.7.1. We remark that this section is the only part of the proof of Theorem 25 for graphs G which do not have a $(k+1)$-tangle in which we make use of the assumption that G has no vertex of degree $\leqslant 2$. In the remainder of this section, we will always assume that the given graph satisfies the premise of Theorem 8.7.1, which we recall here in the following setting:

Setting 8.7.3. *Let G be a graph of minimum degree at least 3 that admits an RC-decomposition with sun Z which has length M_0 and adhesion ℓ such that $\ell + |Z| \geqslant 1$. Let τ be a k-tangle in G with $k \geqslant 1$.*

First we find our desired edge deep inside some rainbow such that the RC-decomposition is unaffected by its deletion:

Lemma 8.7.4. *If we assume Setting 8.7.3 with $M_0 \geqslant 6k$, then there is an edge e of G and an RC-decomposition $(R, \mathcal{W}, Z, C)$ of G which has even length $M \geqslant M_0/2 - k - 2$ and adhesion ℓ such that τ does not live in the rainbow and e has one endvertex in $W_{M/2}$ and one in $W_{M/2}$ or Z. Moreover, $(R, \mathcal{W}, Z, C)$ is also an RC-decomposition of $G - e$ after deleting e from R, if e has both its endvertices in $W_{M/2}$.*

Proof. Theorem 8.7.2 ensures that G admits an RC-decomposition $(R', \mathcal{W}', Z, C')$ of length $M' \geqslant M_0/2 - k$ and adhesion ℓ such that τ does not live in the rainbow R'; by Lemma 8.6.3, we may assume without loss of generality that M' is even. To find the desired edge e, we merge the middle bags $W'_{M'/2-1}$, $W'_{M'/2}$, $W'_{M'/2+1}$ of $\mathcal{W}'$ into one bag $W_{M/2}$ in $\mathcal{W}$ to make it robust against the deletion of an edge, and keep all other bags of $\mathcal{W}'$ in $\mathcal{W}^*$. It is immediate that this yields an RC-decomposition $(R, \mathcal{W}, Z, C)$ in whose rainbow τ does not live, where $R := R', C := C'$; in particular, it has length $M = M' - 2$ and adhesion ℓ.

First, assume that $|Z| \geqslant 1$. We fix an edge E of G between Z and $W_{M/2}$, which exists by (RC4). Note that we still have $Z \subseteq N_{G'}(W_{M/2})$: Since $(R', \mathcal{W}', Z, C')$ is an RC-decomposition, each of the $Z \subseteq N_G(W'_{M'/2+i})$ for every $i \in \{-1, 0, +1\}$ by (RC4) and $W'_{M'/2-1} \cap W'_{M'/2+1} = \emptyset$ by (R3). Thus, each vertex of Z, in particular the endvertex of e in Z, is joined to $W_{M/2}$ by two distinct edges of G of which one persists in $G - e$. Thus, it is now easy to see that $(R, \mathcal{W}, Z, C)$ is also an RC-decomposition of $G - e$.

Secondly, assume that $Z = \emptyset$. Then $|Z| + \ell \geqslant 1$ ensures that $\ell \geqslant 1$. By (R1), there exists a $U_{M/2}$–$U_{M/2+1}$ linkage $\mathcal{P}$ of cardinality ℓ in $G[W_{M/2}]$; note that $U'_{M'/2+i} \subseteq V(\bigcup \mathcal{P})$ for every $i \in \{-1, 0, +1, +2\}$ by the construction of $W_{M/2}$ and since $(R', \mathcal{W}', Z', C')$ also has adhesion ℓ. We aim to fix an edge e of $G[W_{M/2}] - E(\bigcup \mathcal{P})$ such that $G[W_{M/2}] - e$ is still connected. If this is possible, it is easy to check that $(R - e, \mathcal{W}, Z, C)$ is an RC-decomposition of $G - e$, i.e. e is as desired. We claim that such an edge e exists:

Consider any vertex $v \in U'_{M'/2}$, which exists as $\ell \geqslant 1$. All its neighbours and itself are contained in $W'_{M'/2-1} \cup W'_{M'/2} \subseteq W_{M/2}$, as $Z = \emptyset$ and by the definition of RC-decompositions, in particular by (R3). Since v has degree at least 3 by assumption on G and every vertex has at most 2 incident edges in the linkage $\mathcal{P}$, there is a neighbour w of v with $vw \notin E(\mathcal{P})$. If $G[W^*_{M^*/2}] - vw$ is connected, $e := vw$ is as desired. So assume that $G[W_{M/2}] - vw$ is disconnected. We note that the linkage $\mathcal{P}$ is contained in one component of $G[W_{M/2}] - vw$, since $G[W'_{M'/2+1}] \subseteq G[W_{M/2}]$ is connected by (R2), meets every path in $\mathcal{P}$ as it contains $U'_{M'/2+1}$ and also avoids vw by the choice of vw. Thus, the component C of $G[W_{M/2}] - vw$ containing w does not meet $V(\mathcal{P})$, as $v \in U'_{M'/2} \subseteq V(\mathcal{P})$. Now $U_{M/2} \cup U_{M/2-1} = U'_{M'/2+2} \cup U'_{M'/2-1} \subseteq V(\mathcal{P})$ is the separator of a separation of G with precisely $W_{M/2}$ on one of its sides by Lemma 8.6.4. Thus, C is not only a component of $G[W_{M/2} - vw]$, but also a component of $G - vw$, since C avoids $\mathcal{P}$. In particular, all vertices in C have degree at least 2 in C as G has minimum degree 3. Hence, C contains a cycle, and we fix e as an edge on this cycle. Then $C - e$ is still connected which by the previous description implies that also $G[W_{M/2}] - e$ is connected, as desired. □

We emphasise that in Lemma 8.7.4 the possibly removed edge e from R is the only difference between the two RC-decompositions of G and $G' := G - e$; in particular, all corresponding vertex sets such as $V(R)$ and the $V(R_{i,j})$ are the same, and a separation crosses or slices the rainbow in G' if and only if it does so in G.

For the proof of Theorem 8.7.1, we have to construct an orientation τ' of $S_k(G')$ that is a k-tangle in G' and extends τ. Clearly, any such extension τ' has to contain not only τ, but also all orientations of separations of $S_k(G')$ that are forced by the request that τ' shall again be a tangle and hence especially consistent. More formally, recall that τ *forces* an orientation of a separation $\{A, B\}$ of G' if there exists a separation $(E, F) \in \tau$ such that $(A, B) \leqslant (E, F)$ or $(B, A) \leqslant (E, F)$ (cf. the proof of Lemma 8.4.4). Let us also recall that τ forces an orientation of every separation in $S_k(G')$ that is also a separation of G.

Towards a suitable construction of τ', we now prove that rainbow-crossing and rainbow-slicing separations are forced by τ (Lemmas 8.7.6 and 8.7.7 below). This then allows us to show that an even broader class of separations are forced by τ (Lemma 8.7.8 below). For all these subsequent three lemmas, we assume the following setting:

Setting 8.7.5. *Assume Setting 8.7.3 with* $M_0 \geqslant 18k$. *Fix some edge* e *of* G *and some RC-decomposition* $(R, \mathcal{W}, Z, C)$ *of* G *of length* M *which Lemma 8.7.4 yields. In particular,* $M \geqslant 8k$.

Lemma 8.7.6. *Assume Setting 8.7.5. If* $\{A, B\} \in S_k(G')$ *crosses the rainbow, then there is* $(X, Y) \in \tau$ *with* $(A, B) \leqslant (X, Y)$ *or* $(B, A) \leqslant (X, Y)$ *such that* $V(R_{2k-1,M-2k+1}) \subseteq X$. *In particular,* τ *forces an orientation of* $\{A, B\}$.

Proof. By possibly reversing the rainbow, we may assume that (A, B) crosses the rainbow clockwise. Let $k' < k$ be the order of $\{A, B\}$, and set $i := i_{A,B} \in \{0, \dots, 2k'\}$ and $j := j_{A,B} \in \{M-2k', \dots, M\}$. By Lemma 8.6.5, $(A^{i+1}, B^{i+1}), (A^j, B^j)$ are again separations of G' of order at most k' which cross the rainbow clockwise. They are also separations of G by the choice of e: The bag $W_{M/2}$ is in $B^{i+1} \cap A^j$, since $i \leqslant 2k' \leqslant M/2 \leqslant M - 2k' \leqslant j$ because $M \geqslant 4k$ and $k' \leqslant k - 1$. As both $(A^{i+1}, B^{i+1}), (A^j, B^j)$ cross the rainbow, Z is contained in each of their separators. Thus, both endvertices of e are in any of the cases contained in $B^{i+1} \cap A^j$, as claimed.

Since τ does not live in the rainbow by the choice of $(R, \mathcal{W}, Z, C)$, the separation $\{A, B\}$ of G is oriented by τ monotonically over the rainbow. In particular, either $(A^{i+1}, B^{i+1}), (A^j, B^j) \in \tau$ or $(B^{i+1}, A^{i+1}), (B^j, A^j) \in \tau$. By definition of the (A^h, B^h) and the choice of i, j, we have $(A^{i+1}, B^{i+1}) \leqslant (A, B) \leqslant (A^j, B^j)$ and $V(R_{2k'+1,M-2k'-1}) \subseteq B^{i+1}, A^j$. So either (A^j, B^j) or (B^{i+1}, A^{i+1}) can be chosen as the desired separation $(X, Y) \in \tau$, as $k' < k$. □

Lemma 8.7.7. *Assume Setting 8.7.5. If* $\{A, B\} \in S_k(G')$ *slices the rainbow, then there exists* $(X, Y) \in \tau$ *with* $(A, B) \leqslant (X, Y)$ *or* $(B, A) \leqslant (X, Y)$ *such that* $V(R) \subseteq X$. *In particular,* τ *forces an orientation of* $\{A, B\}$.

Proof. As the separation $\{A, B\}$ of order $k' < k$ slices the rainbow, we may, by possibly renaming the sides of $\{A, B\}$, assume that there exists integers $i < h < j$ with $i \in \{0, \dots, 2k'\}$ and $j \in \{M - 2k', \dots, M\}$ with $W_i, W_j \subseteq A \setminus B$ and $W_h \subseteq B \setminus A$. If $\{A, B\}$ is only a separation of G', but not a separation of G, we let D be the component of $G'[B \setminus A]$ that contains an endvertex of e; otherwise, we let D be the empty graph. Since $\{A, B\}$ slices the rainbow, we have $Z \subseteq A \cap B$ and thus $Z \cap V(D) = \emptyset$. So by the choice of the edge e that we deleted from G to get G', we observe that if $V(D) \neq \emptyset$, then it meets $W_{M/2}$. By Lemma 8.6.2 applied in G', we thus obtain that either $V(D) \subseteq V(R) \setminus V(C)$ (which is in particular the case if $V(D) = \emptyset$) or D meets every bag of either $W_0, \dots, W_{M/2}$ or $W_{M/2}, \dots, W_M$. But since $W_i, W_j \subseteq A \setminus B$ and $i \leqslant 2k' \leqslant M/2 \leqslant M - 2k' \leqslant j$, the second case cannot occur; so we have $V(D) \subseteq V(R) \setminus V(C)$.

Further, Corollary 8.6.7 yields $2\ell + |Z| \leqslant |A, B| < k$, and hence $|V(R) \cup Z, V(C)| = 2\ell + |Z| < k$. Thus, τ orients $\{V(R) \cup Z, V(C)\}$, and we have $(V(R) \cup Z, V(C)) \in \tau$ since τ does not live in the rainbow by construction of the RC-decomposition $(R, \mathcal{W}, Z, C)$.

By the choice of D, the pair $(A', B') := (A \cup V(D), B \setminus V(D))$ is a separation of G, which again has the same order $< k$ as $\{A, B\}$. Thus, τ contains an orientation of $\{A', B'\}$. If $(A', B') \in \tau$, we define

$$\begin{aligned}(E,F) &:= (A',B') \vee (V(R)\cup Z, V(C)) = ((A\cup V(D))\cup V(R), (B\setminus V(D))\cap V(C))\\ &= (A\cup V(R), B\cap V(C)).\end{aligned}$$

If $(B',A')\in\tau$, we define

$$\begin{aligned}(E,F) &:= (B',A') \vee (V(R)\cup Z, V(C)) = ((B\setminus V(D))\cup V(R), (A\cup V(D))\cap V(C))\\ &= (B\cup V(R), A\cap V(C)).\end{aligned}$$

We remark that in both definition the second equality holds due to $Z\subseteq A\cap B\setminus V(D)$ and the third holds due to $V(D)\subseteq V(R)\setminus V(C)$. We remark that the proof in the case of $(A',B')\in\tau$ is analogous to the one below in the case of $(B',A')\in\tau$ by swapping A and B; thus, one obtains $(A,B)\leqslant (A\cup V(R), B\cap V(C) = (E,F)\in\tau$, if $(A',B')\in\tau$.

So let us assume that $(B',A')\in\tau$. The pair (E,F) is again a separation of G, since it is the supremum of two separations of G. We claim that $\{E,F\}$ has order $<k$. Indeed, we have

$$\begin{aligned}|E,F| &= |(B\cup V(R))\cap(A\cap V(C))| = |(B\cap A\cap V(C))\cup(V(R)\cap A\cap V(C))|\\ &= |B\cap A\cap(V(C)\setminus V(R))| + |V(R)\cap V(C)\cap A|\\ &\leqslant |B\cap A\cap(V(C)\setminus V(R))| + |V(R)\cap V(C)|\\ &= |B\cap A\cap(V(C)\setminus V(R))| + 2\ell \leqslant |B\cap A\cap(V(C)\setminus V(R))| + |B\cap A\cap V(R)|\\ &= |B\cap A| = |B,A| < k,\end{aligned}$$

where $2\ell\leqslant |B\cap A\cap V(R)|$ holds by Corollary 8.6.7. Hence, τ has to orient $\{E,F\}$ and it does so as (E,F) by the profile property $(*)$ of the tangle τ, since (E,F) is the supremum of the two separations (B',A') and $(V(R)\cup Z, V(C))$, which are both contained in τ. Thus, $(B,A)\leqslant (B\cup V(R), A\cap V(C)) = (E,F)\in\tau$ and $V(R)\subseteq E$, so $(X,Y):=(E,F)$ is as desired. □

Lemma 8.7.8. *Assume Setting 8.7.5. Let $\{A,B\}\in S_k(G')\setminus S_k(G)$, and let $C_A\subseteq G'[A\setminus B]$ and $C_B\subseteq G'[B\setminus A]$ be the two components of $G'-(A\cap B)$ that contain an endvertex of e. If either both or none of C_A and C_B meet C, then τ forces an orientation of $\{A,B\}$.*

Proof. Let $k'<k$ be the order of $\{A,B\}$. We first argue that we are in the case in which both endvertices of e are in $W_{M/2}$. Suppose for a contradiciton that e otherwise has one endvertex in $W_{M/2}$ and the other one in Z. Then one of C_A and C_B, say C_A, contains the endvertex z of e in Z and thus meets $V(C)\supseteq Z$. Then C_B meets C as well by assumption. Since C_B also contains the other endvertex of e, i.e. the one in $W_{M/2}$, and $M\geqslant 4k'$, Lemma 8.6.2 yields that C_B contains

a bag W_i. By (RC4), G' contains an edge joining $z \in V(C_A) \subseteq A \setminus B$ and $W_i \subseteq V(C_B) \subseteq B \setminus A$, which is a contradiction. So we may assume by the choice of e that both endvertices of e lie in $W_{M/2}$; in particular, both C_A and C_B meet $W_{M/2}$.

The assumptions on C_A, C_B ensure that $\{A \cup V(C_B), B \setminus V(C_B)\}$ and $\{A \setminus V(C_A), B \cup V(C_A)\}$ are both separations of G. Also their order $|A \cap B|$ is $< k$; thus, τ orients them. If τ orients one of them as $(A \cup V(C_B), B \setminus V(C_B)) \geqslant (A, B)$ or $(B \cup V(C_A), A \setminus V(C_A)) \geqslant (B, A)$, then τ forces an orientation of $\{A, B\}$, as desired. So assume that τ orients them as $(B \setminus V(C_B), A \cup V(C_B))$ and $(A \setminus V(C_A), B \cup V(C_A))$. Then their supremum $(V(G) \setminus (V(C_A) \cup V(C_B)), (A \cap B) \cup V(C_A) \cup V(C_B))$ is also contained in τ, as its order is the same as $|A, B| = k' < k$ and the tangle τ has the profile property $(*)$. Since τ does not live in R by construction of $(R, \mathcal{W}, Z, C)$, it follows that $(V(C_A) \cup V(C_B)) \cap V(C) \neq \emptyset$. As either both or none of C_A and C_B meets C by assumption, both C_A and C_B meet C.

All in all, both C_A and C_B meet C and $W_{M/2}$. Hence, Lemma 8.6.2 yields that each of C_A and C_B contains all but at most $2k'$ bags of $W_0, \dots, W_{M/2}$ or of $W_{M/2}, \dots, W_M$; in particular, they contain some W_i and some W_j respectively, as $M \geqslant 4k'$. Now Lemma 8.6.8 ensures that $\{A, B\}$ either crosses or slices the rainbow in G'. Thus, the previous Lemmas 8.7.6 and 8.7.7 guarantee that τ forces an orientation of $\{A, B\}$. □

With these tools at hand, we are now ready to prove Theorem 8.7.1:

Proof of Theorem 8.7.1. We may assume Setting 8.7.5, as Setting 8.7.3 with $M_0 \geqslant 18k$ is precisely the premise of Theorem 8.7.1. We claim that τ extends to a k-tangle in G'. For this, we begin by defining an orientation τ' of $S_k(G')$. So let $\{A, B\}$ be an arbitrary separation in $S_k(G')$. If $\{A, B\}$ is not a separation of G, then we let $C_A \subseteq G'[A \setminus B]$ and $C_B \subseteq G'[B \setminus A]$ be the components of $G - (A \cap B)$ which contain the respective endvertex of e.

(1) If τ forces an orientation of $\{A, B\}$, then we let $(A, B) \in \tau'$ if and only if (A, B) is forced by τ.

(2) If τ does not force an orientation of $\{A, B\}$, then $\{A, B\} \notin S_k(G)$. If C_B meets C, then we let $(A, B) \in \tau'$, and if C_A meets C, we let $(B, A) \in \tau'$.

Note that τ' contains at least one orientation of every separation in $S_k(G')$. Moreover, τ' is an orientation of $S_k(G')$: for a separation in case (1) τ forces at most one of its orientations due to the consistency of τ, and Lemma 8.7.8 ensures that precisely one of C_A and C_B meet C for a separation $\{A, B\}$ in case (2). It is immediate form (1)that τ' contains τ; so τ extends to τ'. It remains to show that τ' is indeed a k-tangle in G'.

Suppose for a contradiction that τ' is not a tangle in G', i.e. the orientation τ' of $S_k(G')$ contains a forbidden triple $\{(A'_i, B'_i) : i \in [3]\} \in \mathcal{T}(G')$. By the definition of $\mathcal{T}(G')$, we may assume without loss of generality that all the (A'_i, B'_i) are $\leqslant$-maximal in τ'. In the remainder of the proof, we will construct from the forbidden triple $\{(A'_i, B'_i) : i \in [3]\} \subseteq \tau'$ a forbidden triple $\{(A_i, B_i) : i \in [3]\} \subseteq \tau$, which then contradicts that τ is a tangle in G.

To this end, we will make use of the RC-decomposition $(R, \mathcal{W}, Z, C)$. We first show that one of the (A'_i, B'_i) contains $V(R_{2k-1,M-2k+1}) \cup Z$ in its respective small side A'_i; in particular, the edge e is contained in $G[A'_i]$ and thus it is a separation of G. We then show for the other two separations (A'_i, B'_i) in the forbidden triple that either they are separations of G, too, or the component $C_{A'_i}$ of $G[A'_i \setminus B'_i]$ containing an endvertex of e in fact contains also $V(R_{2k-1,M-2k+1})$. Moving the $C_{A'_i}$ to the respective big side B'_i if necessary, these three separations of G will then yield a forbidden triple in τ. So we first show the following:

Claim 1. *If $\{(A'_i, B'_i) : i \in [3]\} \in \mathcal{T}(G')$ is contained in τ' where each (A'_i, B'_i) is $\leqslant$-maximal in τ', then some (A'_j, B'_j) is also a separation of G with $V(R_{2k-1,M-2k+1}) \cup Z \subseteq A'_j$. In particular, $(A'_j, B'_j) \in \tau$.*

Proof. For every (A'_i, B'_i), at most $2k-2$ bags of $\mathcal{W}$ are contained neither in $A'_i \setminus B'_i$ nor in $B'_i \setminus A'_i$ by Lemma 8.6.1. Thus, all but at most $6k-6$ bags are contained in one of the strict sides of every (A'_i, B'_i). As $\{(A'_i, B'_i) : i \in [3]\}$ is a forbidden triple in G', we in particular have that $R[A'_1]$, $R[A'_2]$ and $R[A'_3]$ together cover the rainbow R in G'. Thus, no such bag may be contained in every big side B'_i. Since $M \geqslant 6k-6$, there thus must be some bag W_h of $\mathcal{W}$ that is contained in the strict small side $A'_j \setminus B'_j$ of some (A'_j, B'_j); by renaming the (A'_i, B'_i), we may assume that $j = 1$. Note that as $W_h \subseteq A'_1 \setminus B'_1$ and $Z \subseteq N_{G'}(W_h)$ by (RC4), we have $Z \subseteq A'_1$. We claim that (A'_1, B'_1) is as desired, i.e. it is also a separation of G and $V(R_{2k-1,M-2k+1}) \subseteq A'_1$.

Assume that $B'_1 \setminus A'_1$ also contains some bag W_s. By Lemma 8.6.8, (A'_1, B'_1) then either crosses or slices the rainbow in G', and thus, the orientation (A'_1, B'_1) of $\{A'_1, B'_1\}$ was forced by τ by Lemmas 8.7.6 and 8.7.7. But if a separation is both $\leqslant$-maximal in τ' and forced by τ, then it is also contained in τ, and thus a separation of G. Hence, (A'_1, B'_1) is a separation of G and also $\leqslant$-maximal in τ. Moreover, then Lemmas 8.7.6 and 8.7.7 yield $V(R_{2k-1,M-2k+1}) \subseteq A'_1$, as desired.

So now assume that no bag of $\mathcal{W}$ is contained in $B'_1 \setminus A'_1$. By Lemma 8.6.1, the strict side $A'_1 \setminus B'_1$ then contains all but at most $2k-2$ many bags of $\mathcal{W}$. Suppose for a contradiction that $\{A'_1, B'_1\}$ was not already a separation of G. Since (A'_1, B'_1) is $\leqslant$-maximal in τ', the tangle τ thus did not force an orientation of $\{A'_1, B'_1\}$. Hence we have $(A'_1, B'_1) \in \tau'$ due to (2). Hence, $C_{B'_1}$ meets C. Since we have seen above that $Z \subseteq N_{G'}(W_m) \subseteq A'_1$, the choice of e yields that the endvertex of e contained in $C_{B'_1} \subseteq B'_1 \setminus A'_1$ lies in $W_{M/2}$. Thus, Lemma 8.6.2 yields that $C_{B'_1} \subseteq B'_1 \setminus A'_1$ contains some bag of $\mathcal{W}$, which contradicts our assumption on $B'_1 \setminus A'_1$. Thus, $\{A'_1, B'_1\}$ is also a separation of G. It remains to show $V(R_{2k-1,M-2k+1}) \subseteq A'_1$. Since (A'_1, B'_1) is a separation of G and $\leqslant$-maximal in τ', it is also $\leqslant$-maximal in the tangle $\tau \subseteq \tau'$. Thus, $G[B'_1 \setminus A'_1]$ is connected, and thus equal to $C_{B'_1}$. As $C_{B'_1}$ meets C but $G[B'_1 \setminus A'_1] = C_{B'_1}$ does not contain a bag of $\mathcal{W}$, Lemma 8.6.2 yields that $C_{B'_1} = G[B'_1 \setminus A'_1]$ does meet at most the first and last $2k-2$ bags of $\mathcal{W}$. Thus, $V(R_{2k-1,M-2k+1}) \subseteq A'_1$, as desired. ■

So by Claim 1, there is some $j \in [3]$ such that $(A'_j, B'_j) \in \tau$ and $V(R_{2k-1,M-2k+1}) \cup Z \subseteq A'_j$; by symmetry we may assume $j = 1$. Next, we aim to obtain from the other two separations (A'_2, B'_2) and (A'_3, B'_3) of the forbidden triple two separations (A_2, B_2) and (A_3, B_3) of G that are contained in τ and differ from (A'_2, B'_2) or (A'_3, B'_3), respectively, only in a subset of A'_1. This will then be the desired forbidden triple in τ.

Claim 2. *Every $\leqslant$-maximal separation (A, B) in τ' is either also contained in τ or we have $(A \setminus V(C_A), B \cup V(C_A)) \in \tau$ and $V(C_A) \subseteq V(R_{2k-1,M-2k+1})$.*

Proof. If the tangle τ forces the orientation (A, B) of $\{A, B\}$, then we already have $(A, B) \in \tau$ by the $\leqslant$-maximality of (A, B) in τ'. Thus, we may assume that $(A, B) \in \tau'$ due to (2), i.e. C_A avoids C and C_B meets C. As each of C_A and C_B contains one endvertex of e, $\{A \setminus V(C_A), B \cup V(C_A)\}$ is a separation of G. It also has separator $A \cap B$, and thus is oriented by the k-tangle τ in G. Since τ does not force an orientation of $\{A, B\}$, it follows from the consistency of the tangle τ that $(B \cup V(C_A), A \setminus V(C_A)) \notin \tau$. Hence, $(A \setminus V(C_A), B \cup V(C_A)) \in \tau$. It remains to show $V(C_A) \subseteq V(R_{2k-1,M-2k+1})$.

First we claim that C_B contains a bag of $\mathcal{W}$: Since C_B contains an endvertex x of the edge e by definition, the choice of e ensures that either $x \in W_{M/2}$ or $x \in Z$. If $x \in W_{M/2}$, then Lemma 8.6.2 yields that C_B contains a bag of $\mathcal{W}$, as C_B meets C and $M \geqslant 4k-4$. So assume that $x \in Z$. By (RC4), x has neighbours in every bag of $\mathcal{W}$. In particular, every bag of $\mathcal{W}$ is adjacent to $V(C_B)$. At most $2k-2$ many bags of $\mathcal{W}$ meet $A \cap B$ by Lemma 8.6.1. Thus, the component C_B of $G' - (A \cap B)$ contains some bag of $\mathcal{W}$, since $M \geqslant 2k-2$ all the $G'[W_i]$ are connected by (R2).

Secondly, we claim that C_A does not contain a bag of $\mathcal{W}$. Suppose that C_A contains a bag of $\mathcal{W}$. Then $\{A, B\}$ would either cross or slice the rainbow by Lemma 8.6.8. Thus, τ forces an orientation of $\{A, B\}$ by Lemmas 8.7.6 and 8.7.7, which contradicts our assumption on $\{A, B\}$.

Hence, as all the $G'[W_i]$ is connected by (R2), every bag which meets C_A also meets $A \cap B$. Lemma 8.6.1 shows that there are at most $2k-2$ such bags of $\mathcal{W}$. Since C_A contains an endvertex of e but avoids $V(C) \supseteq Z$, it contains the endvertex of e in $W_{M/2}$. Hence, C_A can only meet W_i with $|M/2 - i| \leqslant 2k-3$, as C_A is connected and $\mathcal{W}$ is a linear decomposition of R. Thus, $M \geqslant 8k$ yields $V(C_A) \subseteq V(R_{2k-1,M-2k+1})$, as desired. ■

If a separation (A, B) of G' is also a separation of G, let C_A be the empty graph. Recall that we have obtained earlier from Claim 1 that $(A'_1, B'_1) \in \tau$ and $V(R_{2k-1,M-2k+1}) \cup Z \subseteq A'_1$. By Claim 2, we now also have $(A_i, B_i) := (A'_i \setminus V(C_{A'_i}), B'_i \cup V(C_{A'_i})) \in \tau$ for $i = 2, 3$.

We claim that the triple $\{(A'_1, B'_1), (A_2, B_2), (A_3, B_3)\} \subseteq \tau$ is a forbidden triple in G. Since

$\{(A'_i, B'_i) : i \in \{1,2,3\}\}$ is a forbidden triple in G', we have

$$G[A'_1] \cup G[A_2] \cup G[A_3] = G[A'_1] \cup G[A'_2 \setminus V(C_{A'_2})] \cup G[A'_3 \setminus V(C_{A'_3})] \supseteq G - e - C_{A'_2} - C_{A'_3}.$$

Since the endvertices of e are contained in $W_{M/2} \cup Z$, and $W_{M/2} \cup Z \subseteq V(R_{2k-1,M-2k+1}) \cup Z \subseteq A'_1$ as $M \geqslant 4k$, we have that

$$G[A'_1] \cup G[A_2] \cup G[A_3] \supseteq G - C_{A'_2} - C_{A'_3}.$$

By Claim 2, each of $V(C_{A'_2})$ and $V(C_{A'_3})$ is either empty or contained in $V(R_{2k-1,M-2k+1})$. So since $V(R_{2k-1,M-2k+1}) \subseteq A'_1$, it follows that both $C_{A'_2}$ and $C_{A'_3}$ are subgraphs of $G[A'_1]$ and hence $G[A'_1] \cup G[A_2] \cup G[A_3] = G$. Thus, $\{(A'_1, B'_1), (A_2, B_2), (A_3, B_3)\}$ is a forbidden triple in τ, which contradicts that τ is a k-tangle in G. □

Proof of Theorem 25. Choose $M(k)$ to be $N(k, 18k)$ as in Theorem 8.5.1. Let $k \geqslant 1$ be an integer, let G be a connected graph with at least $M(k)$ edges, and let τ be a k-tangle in G. We may assume $k \geqslant 3$; otherwise, we are done by Lemmas 8.3.2 and 8.3.3. If there exists a $(k+1)$-tangle in G, then we are done by Theorem 8.4.1. Therefore, we may assume that there is no tangle in G of order $> k$. Then, the connected graph G admits an RC-decomposition with sun Z which has length $\geqslant 18k$ and adhesion ℓ such that $|Z| + \ell \geqslant 1$ by Theorem 8.5.1. If G has a vertex of degree at most 2, then we are done by Lemmas 8.3.5 and 8.3.6. Thus, Theorem 8.7.1 concludes the proof. □

We remark that one may calculate that $M(k) \in O(3^{k^{k^5}})$.

8.8 The inductive proof method and its applications

This section consists of three parts: we first collect all above auxiliary results to conclude the formal proof of our inductive proof method, Theorem 24. Next, we deduce from Theorem 24 our reduction of Problem 8.1.1 to small graphs, Theorem 21, and then derive Corollaries 22 and 23 from it. Finally, we present a further application of Theorem 24 in Corollary 8.8.4, which bounds the size of a subgraph 'witnessing' a k-tangle.

Let us first prove our inductive proof method, which we restate here for the reader's convenience:

Theorem 24. *For every integer $k \geqslant 1$ there is some $M(k) \in O(3^{k^{k^5}})$ such that the following holds: Let τ be a k-tangle in a graph G. Then there exists a sequence $G_0, \ldots, G_m$ of graphs and k-tangles τ_i in G_i for every $i \in \{0, \ldots, m\}$ such that*

- *$G_0 = G$, $\tau_0 = \tau$;*

- *G_i is obtained from G_{i-1} by deleting an edge, suppressing a vertex, or taking a proper component;*
- *the k-tangle τ_{i-1} in G_{i-1} survives as the k-tangle τ_i in G_i for every $i \in [m]$;*
- *G_m is connected and has less than $M(k)$ edges.*

Proof. Let $M(k)$ be given by Theorem 25. Suppose that $G_0, \ldots, G_{i-1}$ and $\tau_0, \ldots, \tau_{i-1}$ are already defined. If G_{i-1} is disconnected, then we apply Lemma 8.3.4 to obtain G_i and τ_i. If $k \leqslant 2$, then we apply Lemma 8.3.2 or Lemma 8.3.3. If $k \geqslant 3$, but G_{i-1} has a vertex of degree $\leqslant 2$, then we apply Lemma 8.3.5 or Lemma 8.3.6. Thus, we may assume that G_{i-1} is a connected graph with minimum degree $\geqslant 3$ and $k \geqslant 3$. In this case, we apply Theorem 25, if G_{i-1} has at least $M(k)$ edges. Otherwise, we set $m := i - 1$, completing the proof. $\square$

8.8.1 Application I: sets and functions inducing tangles

In this section, we address Problem 8.1.1, which we recall here for convenience.

Problem 8.1.1. *Is every tangle in a graph G induced by some set $X \subseteq V(G)$?*

We will use Theorem 24 to prove Theorem 21, our reduction of Problem 8.1.1 for k-tangles to graphs of size bounded in k. Let us briefly recall the relevant definitions, for which we mostly follow [46].

Given a tangle τ in a graph G, a set $X \subseteq V(G)$ *induces* τ if for every separation $(A, B) \in \tau$, we have $|X \cap A| < |X \cap B|$; in this case, we also say that X *induces* the orientation (A, B) of the separation $\{A, B\}$. As a natural relaxation, a *weight function* on $V(G)$ is a map $w \colon V(G) \to \mathbb{N}$, and we say that it *induces* τ if $w(A) < w(B)$ for all $(A, B) \in \tau$; in this case, we also say that w *induces* the orientation (A, B) of the separation $\{A, B\}$. We remark that a set $X \subseteq V(G)$ induces a tangle τ if and only if its indicator function $\mathbb{1}_X$ induces τ. This allows us to focus on weight functions in what follows.

Recall that Theorem 21 reads as follows:

Theorem 21. *For every integer $k \geqslant 1$, there exists $M = M(k) \in O(3^{k^{k^5}})$ such that for every k-tangle τ in a graph G, there exists a k-tangle τ' in a connected topological minor G' of G with less than M edges such that if a weight function w' on $V(G')$ induces the tangle τ', then the weight function w on $V(G)$ which extends w' by zero induces the tangle τ. In particular, a set of vertices which induces τ' also induces τ.*

For the proof of Theorem 21 via Theorem 24, we need to consider the following setting: Let G' be a graph which arises from a graph G by deleting an edge, suppressing a vertex or by passing to a component such that a tangle τ in G survives as a tangle τ' in G'. We now aim to transfer a weight function inducing the tangle τ' of G' to a weight function inducing the tangle τ of G. The subsequent three lemmas show that the extension by zero always works.

Lemma 8.8.1. *If a k-tangle τ in a graph G extends to a k-tangle τ' in $G - e$ for an edge $e \in G$, then every weight function w on $V(G) = V(G')$ which induces τ' also induces τ. In particular, a set of vertices which induces τ' also induces τ.*

Proof. As τ extends to the tangle τ' in $G - e$, we have $\tau \subseteq \tau'$. Thus, w induces τ as well. □

Lemma 8.8.2. *Let τ be a k-tangle in a graph G with $k \geqslant 3$, and let τ' be the induced k-tangle in a graph $G' = G - v + xy$ obtained by suppressing a vertex v with its two neighbours x, y in G. If a weight function w' on $V(G')$ induces τ', then the weight function w on $V(G)$ which extends w' by zero induces τ. In particular, a set of vertices which induces τ' also induces τ.*

Proof. Throughout this proof, we will use that, by definition of τ', a separation (A', B') of G' is in τ' if and only if at least one of $(A' \cup \{v\}, B')$ or $(A', B' \cup \{v\})$ is in τ. Let $(A, B) \in \tau$. Our aim is to find a separation $(A', B') \in \tau'$ such that $A \subseteq A' \cup \{v\}$ and $B' \subseteq B \cup \{v\}$, which then implies that $w(A) \leqslant w'(A') + w(v) = w'(A') < w'(B') = w'(B') + w(v) \leqslant w(B)$; so w induces the orientation (A, B) of $\{A, B\}$ as desired.

If $x, y \in A$ and $v \notin B$, then $(A', B') := (A \smallsetminus \{v\}, B)$ is a separation of G', and $(A, B) \in \tau$ witnesses $(A', B') \in \tau'$, as desired. Similarly, if $x, y \in B$ and $v \notin A$, then $(A', B') := (A, B \smallsetminus \{v\})$ is contained in τ', as desired.

Let us now assume that $x, y \in A$ and $v \in B$. Then $(A, B \smallsetminus \{v\})$ is a separation of G. Since the only neighbours x, y of v are in A, the separations (A, B) and $(B \smallsetminus \{v\}, A)$ form a forbidden tuple in G. Hence, $(A, B \smallsetminus \{v\})$ is in the tangle τ. Now the above described case yields $(A', B') := (A \smallsetminus \{v\}, B \smallsetminus \{v\}) \in \tau'$, as desired. Similarly, if $u, w \in B$ and $v \in A$, then $(A', B') := (A \smallsetminus \{v\}, B \smallsetminus \{v\}) \in \tau'$.

So to conclude the proof, we may assume that $x \in A \smallsetminus B$ and $y \in B \smallsetminus A$ by possibly renaming x, y; in particular $v \in A \cap B$, as $xv, vy \in E(G)$. Hence, $(C, D) := (A \cup \{y\}, B \smallsetminus \{v\})$ is a separation of G. It suffices to show that $(C, D) \in \tau$ because the above described case for $x, y \in C$ and $v \notin D$ yields that $(A', B') := ((A \smallsetminus \{v\}) \cup \{y\}, B \smallsetminus \{v\}) = (C \smallsetminus \{v\}, D) \in \tau'$, as desired. We now show that $(C, D) \in \tau$: The regularity of the tangle τ of order $k \geqslant 3$ yields that the separation $(\{v, y\}, V(G))$ is in τ. Since the separation (D, C) of G together with (A, B) and $(\{v, y\}, V(G))$ forms a forbidden triple in G, we have $(C, D) \in \tau$. □

Lemma 8.8.3. *Let τ be a k-tangle τ in a graph G, and let τ' be its induced k-tangle τ' in some component G' of G. If a weight function w' on $V(G')$ induces τ', then the weight function w on $V(G)$ which extends w' by zero induces τ. In particular, a set of vertices which induces τ' also induces τ.*

Proof. Consider an arbitrary separation $(A, B) \in \tau$. Then $(A', B') := (A \cap V(G'), B \cap V(G')) \in \tau'$, as τ induces the tangle τ' in the component G' of G. Since w' induces τ' and due to the definition of w, we have $w(A) = w'(A') < w'(B') = w(B)$, as desired. □

Proof of Theorem 21. Let τ_m be the k-tangle in the graph G_m with less than $M(k)$ edges as described in Theorem 24. Let w' be a weight function on $V(G')$ which induces the tangle τ_m. Iteratively applying Lemmas 8.8.1 to 8.8.3 yields that the extension of the weight function w' by zero induces the tangle $\tau_0 = \tau$ in the graph $G_0 = G$. □

As direct consequences of Theorem 21, we now deduce Corollaries 22 and 23:

Corollary 22. *For $k \geqslant 1$, there exists $M = M(k) \in O(3^{k^{k^5}})$ such that Problem 8.1.1 holds for k if it holds for all k-tangles in connected graphs G with fewer than M edges.*

Proof. This follows immediately from Theorem 21. □

Corollary 23. *For every integer $k \geqslant 1$, there exists $K = K(k)$ such that for every k-tangle τ in a graph G there exists a weight function $V(G) \to \mathbb{N}$ which induces τ and whose total weight $w(V(G))$ is bounded by K. In particular, the support of w has size $\leqslant K$.*

Moreover, if Problem 8.1.1 holds for k, then every k-tangle in a graph is induced by a set of at most $M(k)$ vertices, where $M(k)$ is given by Theorem 21.

Proof. We enumerate all the finitely many non-isomorphic connected graphs $G_1, \ldots, G_m$ with fewer than M edges. As every finite graph has at most finitely many tangles, there are only finitely many k-tangles τ in any such G_i. Elbracht, Kneip and Teegen showed with Theorem 8.1.2 that every tangle in a graph is induced by a weight function, so we may fix for every such k-tangle τ a weight function w_τ which induces τ. We then set $K(k)$ to be the maximum over all the total weights of these weight functions w_τ. Theorem 21 yields that every k-tangle in a graph G is induced by some weight function which extends one of the weight functions w_τ by zero, and thus has total weight $\leqslant K(k)$.

The moreover-part follows immediately from Theorem 21 by choosing the weight functions as indicator functions of the inducing sets given by the assumed positive answer to Problem 8.1.1. □

We remark that the proof of Corollary 23 in fact shows that the support of the weight function inducing a k-tangle may actually be bounded by $M(k)$ as given in Theorem 21.

8.8.2 Application II: subgraphs witnessing a tangle

In this section, we demonstrate another application of our inductive proof method Theorem 24 by bounding the size of a subgraph 'witnessing' a tangle. We say that a subgraph H of a graph G

witnesses that an orientation τ of $S_k(G)$ is a tangle if $H \not\subseteq \bigcup_{i=1}^{3} G[A_i]$ for every three (not necessarily distinct) $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$. Indeed, τ is a tangle if and only if such a witnessing subgraph H exists, since every tangle in G is witnessed by G itself.

Grohe and Schweitzer [79, Lemma 3.1][6] proved that every k-tangle is witnessed by a set of edges whose size can be bounded in k. However, their bound is defined recursively and yields a power tower of height $k-1$. By Theorem 24, we obtain a new bound which is significantly better for sufficiently large k:

Corollary 8.8.4. *For every integer $k \geqslant 1$, there is an integer $M' = M'(k) \in O(3^{k^{k^5}})$ such that every k-tangle in a graph G is witnessed by some subgraph H of G of size at most M'.*

Proof of Corollary 8.8.4. Let $M(k)$ be given by Theorem 24, and set $M'(k) := 2M(k)$. Let τ be a k-tangle in some graph G. Theorem 24 yields a k-tangle $\tau' := \tau_m$ in a topological minor $G' := G_m$ of G which is connected and has fewer than $M(k)$ edges; in particular, G' has at most $M(k)$ vertices. Now the lift of τ' to G is indeed τ, as one checks by following the lifts along the inductive structure given by Theorem 24.

Let H' be the subdivision of G' in G, i.e. the subgraph of G from which we obtain G' by vertex suppressions. Consider the subgraph H of H' consisting of the branch vertices $V(G')$ together with precisely one edge of H' per $V(G')$-path in H'. We may choose these edges such that they are incident with at least one branch vertex. Since G' witnesses τ', one again checks along the inductive structure given by Theorem 24 that H witnesses τ. Note that H has at most $|E(G')|$ edges and, as each edge of H is incident to a branch vertex of H', H has at most $|V(G')| + |E(G')| \leqslant M'(k)$ vertices. $\square$

[6] We remark that triple covers in their paper are precisely the witnessing sets here. In fact, they proved a more general result about tangles on bipartitions in a more general setting. However, every k-tangle τ in G induces a tangle τ' on the set of bipartitions of the edge set $E(G)$ of order $< k$: let $(C, D) \in \tau'$ if and only if there exists a separation $(A, B) \in \tau$ with $C \subseteq E(G[A])$ and $D \subseteq E(G[B])$. Then their result yields the described conclusion.

Part III

Coarse graph theory

9 Terminology

This chapter gathers all the definitions that we need for this part. Our presentation closely follows [8].

9.1 Distance, radius and balls

Let G be a graph. The *distance* of two vertices u, v in G, denoted by $d_G(u, v)$, is the minimum length of a u–v path in G. If u, v lie in distinct components of G, then $d_G(u, v) := \infty$. For two sets U and U' of vertices of G, we write $d_G(U, U')$ for the minimum distance of two elements of U and U', respectively. If one of U or U' is just a singleton, then we omit the braces, writing $d_G(v, U') := d_G(\{v\}, U')$ for $v \in V(G)$. If X is a subgraph of G, then we abbreviate $d_G(U, V(X))$ as $d_G(U, X)$.

Given a set U of vertices of G, the *ball (in G) around U of radius* $r \in \mathbb{N}$, denoted by $B_G(U, r)$, is the set of all vertices in G of distance at most r from U in G. If $U = \{v\}$ for some $v \in V(G)$, then we omit the braces, writing $B_G(v, r)$ for the ball (in G) around v of radius r. Additionally, we abbreviate the induced subgraph on $B_G(U, r)$ of G with $G[U, r] := G[B_G(U, r)]$.

Further, the *radius* $\mathrm{rad}(G)$ of G is the smallest number $k \in \mathbb{N} \cup \{\infty\}$ such that there exists some vertex $w \in V(G)$ with $d_G(w, v) \leqslant k$ for every vertex v of G. If G is empty, then we define its radius to be 0. We remark that if G is disconnected but not the empty graph, then its radius is ∞. Note that G has radius at most k if and only if there is some vertex v of G with $V(G) = B_G(v, k)$. Additionally, if $U \subseteq V(G)$, then the *radius of U in G*, denoted as $\mathrm{rad}_G(U)$ is the smallest number $k \in \mathbb{N}$ such that there exists some vertex v of G with $U \subseteq B_G(v, k)$ or ∞ if such a $k \in \mathbb{N}$ does not exist.

If Y is a subgraph of G, then we abbreviate $d_G(U, V(Y))$, $\mathrm{rad}_G(V(Y))$, $B_G(V(Y), r)$, and $G[V(Y), r]$ as $d_G(U, Y)$, $\mathrm{rad}_G(Y)$, $B_G(Y, r)$, and $G[Y, r]$, respectively.

A subgraph Y of G is *c-quasi-geodesic (in G)* for some $c \in \mathbb{N}$ if for every two vertices $u, v \in V(Y)$ we have $d_Y(u, v) \leqslant c \cdot d_G(u, v)$. We call Y *quasi-geodesic* if it is c-quasi-geodesic for some $c \in \mathbb{N}$ and *geodesic* if it is 1-quasi-geodesic.

9.2 Quasi-isometries

Let G, H be graphs, and let $M \in \mathbb{R}_{\geqslant 1}$ and $A \in \mathbb{R}_{\geqslant 0}$. An *$(M, A)$-quasi-isometry* from H to G is a map $\varphi\colon V(H) \to V(G)$ such that

(Q1) $M^{-1} \cdot d_H(h, h') - A \leqslant d_G(\varphi(h), \varphi(h')) \leqslant M \cdot d_H(h, h') + A$ for every $h, h' \in V(H)$, and

(Q2) for every vertex v of G, there exists a node h of H with $d_G(v, \varphi(h)) \leqslant A$.

We say that H is (M, A)-*quasi-isometric* to G if there exists an (M, A)-quasi-isometry from H to G.

The following is a well-known fact.

Lemma 9.2.1. *If a graph H is (M, A)-quasi-isometric to a graph G, then G is $(M, 3AM)$-quasi-isometric to H.* □

9.3 Graph-decompositions

Let G, H be graphs and let $\mathcal{V} = (V_h)_{h \in H}$ be a family of subsets of $V(G)$ indexed by the nodes of H. We call $(H, \mathcal{V})$ an H-*decomposition* of G, a *decomposition* of G *modelled on* H, or just a *graph-decomposition* [48], if

(H1) $\bigcup_{h \in H} G[V_h] = G$, and

(H2) for every vertex v of G, the graph $H_v := H[\{h \in V(H) \mid v \in V_h\}]$ is connected.

Whenever a graph-decomposition is introduced as $(H, \mathcal{V})$, we tacitly assume $\mathcal{V} = (V_h)_{h \in H}$.

The sets V_h are called the *bags* of the graph-decomposition, their induced subgraphs $G[V_h]$ are its *parts*, and the graph H is its *decomposition graph*. A graph-decomposition $(H, \mathcal{V})$ is *honest*, if all its bags V_h are non-empty and for every edge $h_0 h_1$ of H the bags V_{h_0} and V_{h_1} intersect non-emptily.

In fact, every graph-decomposition not only satisfies (H2) but also

(H2') for any connected subgraph Y of G, the graph $H_Y := H[\{h \in V(H) \mid V(Y) \cap V_h \neq \emptyset\}]$ is connected.

Proof. (H1) yields for every edge vw in H a part $G[V_h]$ that contains vw. This implies that both H_v and H_w contain h and hence intersect. Since H_v and H_w are connected by (H2), it follows that $H_v \cup H_w$ is connected as well. This implies that as Y is connected, $H_Y = \bigcup_{y \in Y} H_y$ is connected as well. □

A graph-decomposition of an induced subgraph Y of G is a *partial graph-decomposition* of G with *support* Y. Note that every partial graph-decomposition of G with support $Y = G$ is a graph-decomposition of G and vice versa.

9.4 Radial width and radial spread

While the usual width of a tree-decomposition is measured in terms of the cardinality of its bags, the 'radial width' of a graph-decomposition is measured in terms of the radius of its parts as follows.

Let G, H be graphs and let $(H, \mathcal{V})$ be a graph-decomposition of G. The *(inner)-radial width* of $(H, \mathcal{V})$ is

$$\operatorname{iradw}(H, \mathcal{V}) := \sup_{h \in V(H)} \operatorname{rad}(G[V_h]),$$

and the *outer-radial width* of $(H, \mathcal{V})$ is

$$\operatorname{oradw}(H, \mathcal{V}) = \sup_{h \in V(H)} \operatorname{rad}_G(V_h).$$

We remark that the outer-radial width is at most the (inner-)radial width. Note that if $G[V_h]$ is disconnected for some $h \in H$, then $\operatorname{radw}(H, \mathcal{V}) = \infty$; but whenever G is connected and finite, then $\operatorname{oradw}(H, \mathcal{V}) < \infty$.

Given a non-empty class $\mathcal{H}$ of graphs, the *(inner-)radial $\mathcal{H}$-width* of G is

$$\operatorname{iradw}_{\mathcal{H}}(G) := \min \{\operatorname{radw}(H, \mathcal{V}) \mid (H, \mathcal{V}) \text{ is a graph-decomposition of } G \text{ with } H \in \mathcal{H}\}.$$

Note that the (inner-)radial $\mathcal{H}$-width will always be at most $\operatorname{rad}(G)$ for a connected graph G if $\mathcal{H}$ contains at least one non-empty graph. The outer-radial $\mathcal{H}$-width is defined analogously.

Under the name 'tree breadth', the concept 'outer-radial tree-width' has previously been studied [57, 98] with applications to tree-spanners and routing problems.

Let $(H, \mathcal{V})$ be a graph-decomposition of a graph G modelled on a graph H. The *(inner-)radial spread* of $(H, \mathcal{V})$ is

$$\operatorname{irads}(H, \mathcal{V}) := \sup_{v \in V(G)} \operatorname{rad}(H_v).$$

Additionally, we set $\operatorname{irads}(\mathcal{H}, v) := \operatorname{rad}_{H_v}(H_v)$ for a vertex v of G.

We remark that the inner-radial spread of $(H, \mathcal{V})$ is at least the *outer-radial spread* $\operatorname{orads}(\mathcal{H}) = \sup_{v \in Y} \operatorname{rad}_H(H_v)$ of $(H, \mathcal{V})$.

To shorten notation, we say that a partial graph-decomposition $\mathcal{H}$ of a graph G is (R_0, R_1)*-radial* in G for $R_0, R_1 \in \mathbb{N}$, if $\operatorname{oradw}_G(\mathcal{H}) \leqslant R_0$ and $\operatorname{irads}(\mathcal{H}) \leqslant R_1$.

Berger and Seymour [19, 1.5] proved that the (inner-)radial tree-width of a graph is at most two times its outer-radial width, and thus these width measures are qualitatively equivalent. We generalise this to arbitrary decomposition graphs. For this, we need the next lemma, which shows that we can obtain a new graph-decomposition by enlarging all bags of a given graph-decomposition simultaneously.

Lemma 9.4.1 ([8])**.** *Let $(H, \mathcal{V})$ be a graph-decomposition of a graph G, and let $r \in \mathbb{N}$. For every vertex $h \in H$, we let $V'_h := B_G(V_h, r)$. Then $(H, \mathcal{V}')$ is again a graph-decomposition of G.*

A version of Lemma 9.4.1 for tree-decompositions was first proven in [49, Lemma A.1].

Proof ([8]). For notational simplicity, we write $H_X := H[\{h \in V(H) : V_h \cap V(X) \neq \emptyset\}]$ and $H'_X := H[\{h \in V(H) : V'_h \cap V(X) \neq \emptyset]$ for a subgraph X of G.

By definition, $(H, \mathcal{V}')$ satisfies (H1). For (H2) consider any vertex $v \in G$. Then H_v is connected by (H2) and non-empty by (H1). So in order to prove that H'_v is connected as well, it suffices to show that for every $h \in V(H'_v)$, there is an h–$V(H_v)$ path in H'_v.

By the definition of V'_h, there exists $w \in V_h$ with $d_G(v, w) \leqslant r$, and we fix a shortest v–w path P in G. Then every vertex $p \in P$ satisfies $d_G(v, p) \leqslant r$ as witnessed by P. In particular, every node $h' \in H$ with $V_{h'} \cap V(P) \neq \emptyset$ satisfies $v \in V'_{h'} = B_G(V_{h'}, r)$, and hence H_P is a subgraph of H'_v. But now $h \in V(H_P)$ as P meets the vertex $w \in V_h$, and H_v is a subgraph of H_P as $v \in V(P)$. So since H_P is connected by (H2'), there exists a h–$V(H_v)$ path in H_P and hence in H'_v, as desired. □

Lemma 9.4.2 ([8])**.** *Let $(H, \mathcal{V})$ be a graph-decomposition of a graph G of outer-radial width $k \in \mathbb{N}$ and (inner-)radial spread $r \in \mathbb{N}$. Then setting $V'_h := B_G(V_h, k)$ for every node $h \in H$ yields a graph-decomposition $(H, \mathcal{V}')$ of G of (inner-)radial width at most $2k$ and (inner-)radial spread at most $2kr$.*

Proof ([8]). By Lemma 9.4.1, $(H, \mathcal{V}')$ is indeed a graph-decomposition of H. To show that $(H, \mathcal{V}')$ has (inner-)radial width at most $2k$, consider any node $h \in H$. Since $(H, \mathcal{V})$ has outer-radial width at most k, there is a vertex z_h with $V_h \subseteq B_G(z_h, k)$. Then $z_h \in V'_h$ and $G[V'_h]$ contains a shortest path in G between every vertex $x \in V_h$ and z_h. Moreover, there exists for every vertex $x' \in V'_h$ a vertex $x \in V_h$ with $d_{G[V'_h]}(x', x) = d_G(x', x) \leqslant k$. Thus, $d_{G[V'_h]}(z_h, x) \leqslant d_{G[V'_h]}(z_h, x') + d_{G[V'_h]}(x', x) \leqslant k + k = 2k$. This implies that each $G[V'_h]$ has radius at most $2k$, and hence $(H, \mathcal{V}')$ has (inner-)radial width at most $2k$.

To see that $(H, \mathcal{V}')$ has (inner-)radial spread at most $2kr$, consider any vertex v of G and any node h of H such that $v \in V'_h$. By the definition of V'_h, there exists a vertex $u \in V_h$ such that $d_G(v, u) \leqslant k$. Let $P = p_0 \dots p_n$ be a shortest v–u path in G, in particular, $||P|| = n \leqslant k$. By (H2'), the subgraph $H_P := H[\{h \in H \mid V(P) \cap V_h \neq \emptyset\}]$ of H is connected. Since $(H, \mathcal{V})$ has (inner-)radial spread r, the subgraph H_P has diameter at most $2r \cdot ||P|| \leqslant 2rk$. As P starts in v and ends in u, the subgraph H_P includes H_v and H_u. Let h' be the node of H witnessing that H_v has radius at most r (in H_v); in particular $h' \in H_v \subseteq H_P$. Since also $h \in H_u \subseteq H_P$ and H_P is connected, there is a path Q from h to h' in H_P of length at most $2rk$. As $Q \subseteq H_P$ and every bag V_g for $g \in H_P$ contains a vertex w of P, which implies $d_G(w, v) \leqslant ||P|| \leqslant k$, it follows by the definition of the new bags V'_h that $Q \subseteq H_P \subseteq H'_v := H[\{h \in H \mid v \in V'_h\}]$. Hence, h' witnesses that H'_v has radius at most $2rk$, so $(H, \mathcal{V}')$ has (inner-)radial spread at most $2rk$. □

The following observation is immediate from the definitions.

Observation 9.4.3 ([13])**.** *Let G be a graph and, for every component C of G, let $\mathcal{H}^C = (H^C, \mathcal{V}^C)$ be a graph-decomposition of C. Then $\mathcal{H} = (H, \mathcal{V})$ is a graph-decomposition of G where H is the disjoint union of the H^C and $V_h := V_h^C$ for all $h \in V(H)$ and the unique $C \in \mathcal{C}(G)$ such that $h \in V(H^C)$.*

In particular, $\operatorname{oradw}_G(\mathcal{H}) = \sup_{C \in \mathcal{C}(G)} \operatorname{oradw}_C(\mathcal{H}^C)$ *and* $\operatorname{irads}(\mathcal{H}) = \sup_{C \in \mathcal{C}(G)} \operatorname{irads}(\mathcal{H}^C)$. *Moreover, if, for some connected graph X, we have $H^C \in \operatorname{Forb}_{\preccurlyeq}(X)$ for all $C \in \mathcal{C}(G)$, then also $H \in \operatorname{Forb}_{\preccurlyeq}(X)$.* □

9.5 Fat minors

In this part we deviate from the definition of minor in [41], and consider the following equivalent definition: Let G, X be graphs. A *model* $(\mathcal{V}, \mathcal{E})$ of X in G is a collection $\mathcal{V}$ of disjoint sets $V_x \subseteq V(G)$ for vertices x of X such that each $G[V_x]$ is connected, and a collection $\mathcal{E}$ of internally disjoint V_{x_0}–V_{x_1} paths E_e for edges $e = x_0x_1$ of X which are disjoint from every V_x with $x \neq x_0, x_1$.[1] The V_x are its *branch sets* and the E_e are its *branch paths*. A model $(\mathcal{V}, \mathcal{E})$ of X in G is *K-fat* for $K \in \mathbb{N}$ if $d_G(Y, Z) \geqslant K$ for every two distinct $Y, Z \in \mathcal{V} \cup \mathcal{E}$ unless $Y = E_e$ and $Z = V_x$ for some vertex $x \in V(X)$ incident to $e \in E(X)$, or vice versa. Then X is a *(K-fat) minor* of G, denoted by $X \prec G$ ($X \prec_K G$) if G contains a (K-fat) model of X. We remark that the 0-fat minors of G are precisely its minors. By $\operatorname{Forb}_{\preccurlyeq}(X)$, we denote the class of all graphs with no X minor.

[1] By enlarging the branch sets V_x along the 'adjacent' branch paths E_{xy}, we obtain that this notion of model is equivalent to the notion of model from [41].

10 A Menger-type theorem for two induced paths

We give an approximate Menger-type theorem for when a graph G contains two X–Y paths P_1 and P_2 such that $P_1 \cup P_2$ is an induced subgraph of G. More generally, we prove that there exists a function $f(d) \in O(d)$, such that for every graph G and $X, Y \subseteq V(G)$, either there exist two X–Y paths P_1 and P_2 such that $d_G(P_1, P_2) \geqslant d$, or there exists $v \in V(G)$ such that the ball of radius $f(d)$ around v intersects every X–Y path.

This chapter is based on [10] and joint work with Tony Huynh, Raphael W. Jacobs, Paul Knappe and Paul Wollan.

10.1 Introduction

All graphs in this chapter are finite.[1]

Given a graph G and $X, Y \subseteq V(G)$, Menger's classic result says that the maximum number of pairwise disjoint X–Y paths is equal to the minimum size of a set $Z \subseteq V(G)$ intersecting every X–Y path [101]. We consider the problem of finding many distinct paths between X and Y requiring not just that the paths be pairwise disjoint, but that the paths be pairwise far apart in G. Let H_1 and H_2 be subgraphs of a graph G. We say H_1 and H_2 are *anti-complete* if there does not exist an edge of G with one endvertex in $V(H_1)$ and one endvertex in $V(H_2)$. No analogue of Menger's theorem is known for pairwise anti-complete paths linking two sets of vertices. Such an exact characterisation of when such paths exist is unlikely given that it is NP-complete to decide whether there exist two disjoint anti-complete X–Y paths [20, 21]. Note that a graph G contains two disjoint anti-complete X–Y paths if an only if it has two disjoint X–Y paths which are an induced subgraph (take two disjoint anti-complete X–Y paths P_1 and P_2 with $|V(P_1 \cup P_2)|$ minimal).

We will show the following weak characterization of the induced two paths problem.

Theorem 26. *There exists $c \in \mathbb{N}$ such that for all graphs G, and all $X, Y \subseteq V(G)$, either there exist two disjoint X–Y paths P_1, P_2 such that $P_1 \cup P_2$ is an induced subgraph of G, or there exists $z \in V(G)$ such that $B_G(z, c)$ intersects every X–Y path.*

[1] We remark that while we only considered finite graphs, all proofs in this chapter also work for infinite graphs (with obvious changes).

Theorem 26 immediately follows from a more general result where we ask that the two paths be at distance at least d.

Theorem 27. *Let $c = 129$. For all graphs G, all $X, Y \subseteq V(G)$, and all integers $d \geqslant 1$, either there exist two disjoint X–Y paths P_1, P_2 such that $d_G(P_1, P_2) \geqslant d$ or there exists $z \in V(G)$ such that $B_G(z, cd)$ intersects every X–Y path.*

We remark that Georgakopoulos and Papasoglu have independently proved Theorem 27 (with a slightly worse constant of 272) [75].

We conjecture that Theorem 27 should generalize to an arbitrary number of X–Y paths at pairwise distance d.

Conjecture 28. *There exists $c \in \mathbb{N}$ satisfying the following. For all graphs G, all $X, Y \subseteq V(G)$, and all integers $d, k \geqslant 1$, either there exist k disjoint X–Y paths $P_1, \ldots, P_k$ such that $d_G(P_i, P_j) \geqslant d$ for all distinct i, j or there exists a set $Z \subseteq V(G)$ of size at most $k-1$ such that $B_G(Z, cd)$ intersects every X–Y path.*

McCarty and Seymour proved that the $d = 3$ case of Conjecture 28 actually implies the general case.

Theorem 10.1.1 (McCarty and Seymour, 2023)**.** *If Conjecture 28 holds for $d = 3$ with constant c, then it holds for all $d \geqslant 3$ with constant $3c$.*

Proof. Let H be the dth power of G, that is, $V(H) = V(G)$, and u and v are adjacent in H if and only if $d_G(u, v) \leqslant d$. If there exist $k-1$ balls of radius $3c$ in H whose union intersects all X–Y paths in H, then there are $k-1$ balls of radius $3cd$ in G so that their union intersects all X–Y paths in G. Otherwise (by assumption), there are X–Y paths $P_1, \ldots, P_k$ in H such that $d_H(P_i, P_j) \geqslant 3$ for all distinct i, j. Suppose $P_1 = v_0 \ldots v_t$. For each $0 \leqslant i \leqslant t-1$, there is a path S_i of length at most d in G between v_i and v_{i+1}. Thus, there is a $v_0 - v_t$ path P_1^* in G such that $P_1^* \subseteq \bigcup_{i=0}^{t-1} S_i$. We define $P_2^*, \ldots P_k^*$ similarly. We claim that $d_G(P_i^*, P_j^*) \geqslant d$ for all distinct i, j. Towards a contradiction, suppose $d_G(P_i^*, P_j^*) \leqslant d$ for some $i < j$. Then there is a $V(P_i) - V(P_j)$ path in G of length at most $\frac{d}{2} + d + \frac{d}{2} \leqslant 2d$. Hence, $d_H(P_i, P_j) \leqslant 2$, which is a contradiction. $\square$

We would like to point out that Nguyen, Scott, and Seymour [103] have recently found a counterexample to Conjecture 28.

Our proof shows that we may take $c = 129$ in Theorem 27. On the other hand, the following lemma shows that c must be at least $3/2$.

Lemma 10.1.2. *For every integer $d \geqslant 2$, there exists a graph G and $X, Y \subseteq V(G)$ such that there do not exist two disjoint X–Y paths P_1, P_2 such that $dist_G(P_1, P_2) \geqslant d$, nor does there exist a vertex $z \in V(G)$ such that $B_G(z, 3d/2 - 2)$ intersects every X–Y path.*

Proof. Consider the graph H in Figure 10.1. If we let $X = \{x_1, x_2, x_3\}$ and $Y = \{y_1, y_2, y_3\}$, then there do not exist two disjoint X–Y paths P_1 and P_2 such that $d_H(P_1, P_2) \geqslant 2$. However, neither does H have a vertex z such that $B_H(z, 1)$ intersects every X–Y path. Let G be the graph obtained

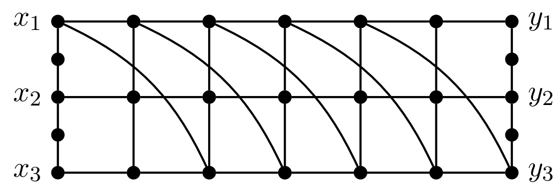

FIGURE 10.1: A graph without two disjoint anti-complete $\{x_1, x_2, x_3\} - \{y_1, y_2, y_3\}$ paths and no ball of radius one hitting all X–Y paths.

from H by subdividing each edge exactly $d-2$ times. We claim that G satisfies the lemma. To see this, let z be any vertex in G. Then there is a vertex $y \in H$ of distance at most $(d-1)/2$ to z in G. Every other vertex of H has distance at least $(d-1)/2$ to z in G. Since any two distinct vertices of H are at distance at least $d-1$ in G and $B_H(y, 1)$ does not intersect every X–Y path in H, it follows that $B_G(z, d - 3/2 + (d-1)/2)$ does not intersect every X–Y path in G. □

10.2 Interlaced systems of intervals

In this section we prove a weak characterisation of when a sequence of intervals contains an 'interlaced' subsequence. We begin with the definition of interlaced. See Figure 10.2 for a visualisation.

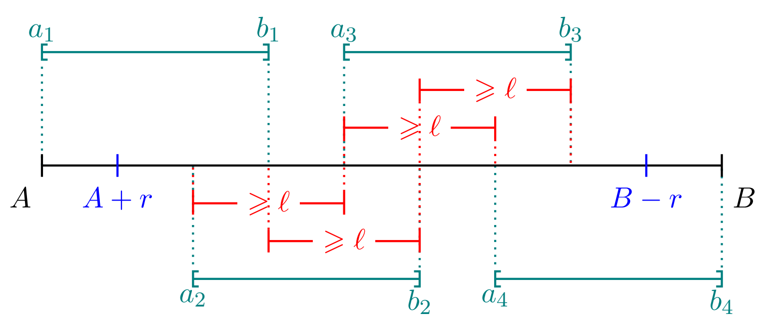

FIGURE 10.2: A clean (A, B, r)-system $([a_i, b_i])_{i \in [4]}$ indicated in teal which is interlaced with buffer ℓ.

Definition 10.2.1. Let $r \in \mathbb{N}$, and $A, B \in \mathbb{Z}$ with $A \leqslant B$. An (A, B, r)-*system* is a sequence $([a_i, b_i])_{i\in[t]}$ of closed intervals of the real line[2] such that for all $i \in [t]$,

- $a_i, b_i \in \mathbb{Z}$,
- $[a_i, b_i] \subseteq [A, B]$,
- if $a_i \neq A$, then $a_i \geqslant A + r$, and
- if $b_i \neq B$, then $b_i \leqslant B - r$.

Moreover, $([a_i, b_i])_{i\in[t]}$ is *interlaced with buffer* ℓ for an integer $\ell > 0$ if

- $a_1 = A, b_t = B$, and
- $a_i + \ell \leqslant a_{i+1} \leqslant b_i$ and $b_i + \ell \leqslant b_{i+1}$ for all $i \in [t - 1]$.

The system is *clean* if $[a_i, b_i] \cap [a_j, b_j] \neq \emptyset$ implies that $|i - j| \leqslant 1$.

Note that if $([a_i, b_i])_{i\in[t]}$ is interlaced with buffer ℓ, then the a_i and b_i values are both strictly increasing. An (A, B, r)-system is clean if and only if we can draw its intervals above and below the line without any overlaps (cf. Figure 10.3).

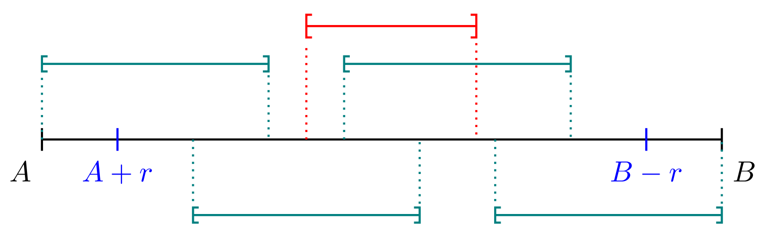

FIGURE 10.3: The (A, B, r)-system indicated in teal is clean. However, adding the red interval gives an (A, B, r)-system which is not clean.

Lemma 10.2.2. *Let $([a_i, b_i])_{i\in[t]}$ be an interlaced system with buffer ℓ. Then there exists a subsequence containing $[a_1, b_1]$ and $[a_t, b_t]$ which is a clean interlaced system with buffer ℓ.*

Proof. We proceed by induction on t. The lemma holds trivially if $t = 1$. Assume that $[a_1, b_1], \dots, [a_t, b_t]$ is not clean. Thus, there are integers $i, j \in [t]$ such that $[a_i, b_i] \cap [a_j, b_j] \neq \emptyset$ and $j \geqslant i + 2$. Hence, $a_j \leqslant b_i$. Moreover, $a_i + \ell \leqslant a_{i+1} \leqslant a_j$ and $b_i + \ell \leqslant b_{i+1} \leqslant b_j$. Thus, $[a_1, b_1], \dots, [a_i, b_i], [a_j, b_j], \dots, [a_t, b_t]$ is an interlaced system with buffer ℓ which is strictly shorter. Applying induction to this shorter sequence yields the desired subsequence of the original sequence, proving the claim. □

[2] We will in fact only consider the integers in such intervals.

An immediate consequence of the definition of a clean interlaced system is that every element $x \in [A, B]$ is contained in at most two intervals.

Lemma 10.2.3. *Let $([a_i, b_i])_{i\in[t]}$ be an (A, B, r)-system with $A \leqslant B-2r$ and let $\ell \leqslant r$ be a positive integer. Then either $([a_i, b_i])_{i\in[t]}$ contains an (A, B, r)-subsystem which is interlaced with buffer ℓ, or there exists an integer $z \in [A+r, B-r]$ such that, for $Z = [z-2\ell, z+2\ell] \cap (A, B)$, no interval in $([a_i, b_i])_{i\in[t]}$ intersects both segments of $[A, B] \setminus Z$.*

Proof. We say an integer $y \in [A + r, B - r]$ is *good* if $([a_i, b_i])_{i\in[t]}$ has an interlaced (A, B', ℓ)-subsystem $[a_{i_1}, b_{i_1}], \dots, [a_{i_s}, b_{i_s}]$ with buffer ℓ such that $y \in [a_{i_s} + \ell, b_{i_s} - \ell]$. If $B - r$ is good, then we find the desired interlaced subsystem. Indeed, then there exists an interlaced (A, B', ℓ)-subsystem $[a_{i_1}, b_{i_1}], \dots, [a_{i_s}, b_{i_s}]$ with buffer ℓ such that $B-r \in [a_{i_s}+\ell, b_{i_s}-\ell]$. Thus, $b_{i_s} > B-r$, which implies $b_{i_s} = B$ since $([a_i, b_i])_{i\in[t]}$ is an (A, B, r)-system. Hence, we have $B' = B$ and $[a_{i_1}, b_{i_1}], \dots, [a_{i_s}, b_{i_s}]$ is as desired.

Thus, we may assume that some integer in $[A + r, B - r]$ is not good, and we let z be the smallest such integer. We claim that z is as desired. So let $Z = [z - 2\ell, z + 2\ell] \cap (A, B)$ and assume for a contradiction that there exists an interval $[a_k, b_k]$ which intersects both segments of $[A, B] \setminus Z$. We fix such an interval $[a_k, b_k]$ for the remainder of the proof.

Suppose $a_k = A$. Since $b_k > z + 2\ell$, the interval $[a_k, b_k]$ by itself is an interlaced system with buffer ℓ certifying that z is good, a contradiction to our choice of z. Thus, $a_k \geqslant A + r$. Since $a_k < z$, it follows that a_k is good by the minimality of z.

Let $[a_{i_1}, b_{i_1}], \dots, [a_{i_s}, b_{i_s}]$ be an interlaced (A, B', ℓ)-subsystem with buffer ℓ such that $a_k \in [a_{i_s} + \ell, b_{i_s} - \ell]$. First, we observe that $b_{i_s} < z + \ell$. Otherwise, $z \in [a_{i_s} + \ell, b_{i_s} - \ell]$, contradicting the fact that z is not good. Since $a_k \in [a_{i_s} + \ell, b_{i_s} - \ell]$ and $b_k \geqslant z + 2\ell \geqslant b_{i_s} + \ell$, it follows that $[a_{i_1}, b_{i_1}], \dots, [a_{i_s}, b_{i_s}], [a_k, b_k]$ is an interlaced system with buffer ℓ. However, $z \in [a_k + \ell, b_k - \ell]$ implying that z is good. This again contradicts our choice of z and thus completes the proof. □

10.3 Construction of the two paths

We now prove our main theorem, which we restate for convenience.

Theorem 27. *Let $c = 129$. For all graphs G, all $X, Y \subseteq V(G)$, and all integers $d \geqslant 1$, either there exist two disjoint X–Y paths P_1, P_2 such that $d_G(P_1, P_2) \geqslant d$ or there exists $z \in V(G)$ such that $B_G(z, cd)$ intersects every X–Y path.*

Proof. Let $d_1 = 5d$ and $d_2 = 62d$. If G has no X–Y path, we may take z to be any vertex of G. Thus, let $P = v_0 \dots v_n$ be a shortest X–Y path in G. Let $\mathcal{H}$ be the family of components of $G - B_G(P, d_1)$. If some $H \in \mathcal{H}$ contains an X–Y path W, then we may take $P_1 = P$ and $P_2 = W$.

Note that $d_G(P_1, P_2) \geqslant d_1 \geqslant d$. Thus, we may assume that no $H \in \mathcal{H}$ contains vertices of both X and Y.

For each $H \in \mathcal{H}$, we define $a(H)$ to be $-d_2$ if H contains a vertex of X. Otherwise, $a(H)$ is the smallest index j such that $d_G(v_j, H) = d_1 + 1$. Similarly, we define $b(H)$ to be $n + d_2$ if H contains a vertex of Y. Otherwise, $b(H)$ is the largest index j such that $d_G(v_j, H) = d_1 + 1$. For each $H \in \mathcal{H}$, we define $I(H)$ to be the interval $[a(H), b(H)]$ (see Figure 10.4). Let $\mathcal{H}' \subseteq \mathcal{H}$ be some choice of components of $G - B_G(P, d_1)$ such that every component in $\mathcal{H}'$ corresponds to a maximal interval in $I(\mathcal{H}) := \{I(H) : H \in \mathcal{H}\}$ and, conversely, every maximal interval in $I(\mathcal{H})$ is of the form $[a(H), b(H)]$ for precisely one $H \in \mathcal{H}'$. Observe that $I(\mathcal{H}') := \{I(H) : H \in \mathcal{H}'\}$ is a $(-d_2, d_2 + n, d_2)$-system.

By Lemma 10.2.3, either there is an ordering $H_1, \dots, H_m$ of a subset of $\mathcal{H}'$ such that $I(H_1), \dots, I(H_m)$ is an interlaced system with buffer d_2, or there is an integer $j \in [0, n]$ such that no interval in $I(\mathcal{H}')$ intersects both segments of $[-d_2, d_2+n] \setminus Z$ for $Z = [j-2d_2, j+2d_2] \cap (-d_2, d_2+n)$.

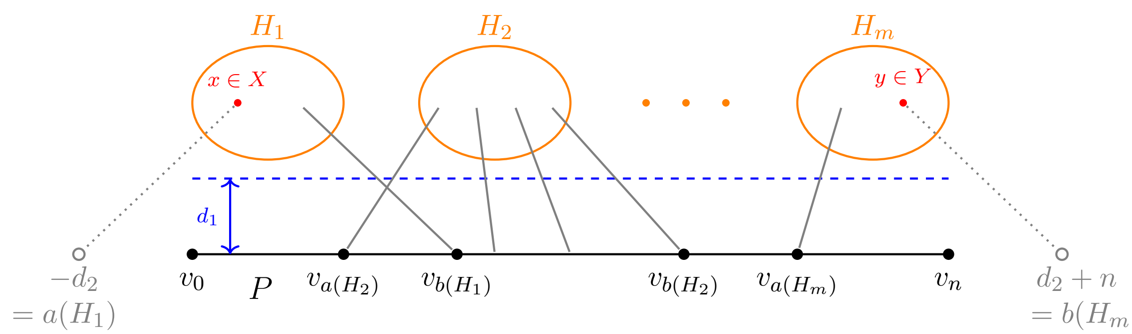

FIGURE 10.4: The components H_i along with their respective $a(H_i)$ and $b(H_i)$ represented by the corresponding vertex in P or a vertex in X or Y.

Suppose the latter holds. We claim that $B_G(z, cd)$ intersects all the X–Y paths for $z := v_j$. Suppose not and let Q be an X–Y path which avoids $B_G(z, cd)$. Let $R_1 := Pv_{j-2d_2-1}$ and $R_2 := v_{j+2d_2+1}P$. Since $cd \geqslant 2d_2 + d_1$, the path Q in particular avoids $B_G(v_{j-2d_2}Pv_{j+2d_2}, d_1) \supseteq B_G(P, d_1) \setminus (B_G(R_1, d_1) \cup B_G(R_2, d_1))$.

For $i \in \{1, 2\}$ let $\mathcal{H}_i$ be those components $H \in \mathcal{H}$ which have a neighbour in $B_G(R_i, d_1)$. By assumption, no interval in $I(\mathcal{H}')$ intersects both segments of $[-d_2, d_2 + n] \setminus Z$. Since $\mathcal{H}' \subseteq \mathcal{H}$ are the components which correspond to the maximal intervals in $I(\mathcal{H})$, we conclude that no interval in $I(\mathcal{H})$ intersects both segments of $[-d_2, d_2+n] \setminus Z$. For each $H \in \mathcal{H}$, the choice of $a(H)$ and $b(H)$ implies that if v_h is a vertex of P such that $d_G(v_h, H) = d_1 + 1$, then $h \in I(H) = [a(H), b(H)]$. Therefore, we conclude that $\mathcal{H}_1$ and $\mathcal{H}_2$ are disjoint, no $H \in \mathcal{H}_1$ meets Y, and no $H \in \mathcal{H}_2$ meets X.

Since P is a shortest path, $d_G(v_k, v_\ell) = |k - \ell|$ for all $0 \leqslant k, \ell \leqslant n$. Hence, $d_G(R_1, R_2) = 4d_2 + 2$. Since an edge in G from $B_G(R_1, d_1)$ to $B_G(R_2, d_1)$ would witness that $d_G(R_1, R_2) \leqslant 2d_1 + 1$ but $2d_1 + 1 < 4d_2 + 2$, it follows that no such edge exists.

Altogether, the X–Y path Q avoiding $B_G(z, cd)$ must be contained in $B_G(R_i, d_1) \cup \bigcup \mathcal{H}_i$ for some $i \in \{1, 2\}$. We consider the case $i = 1$ as the case $i = 2$ is analogous. Since no $H \in \mathcal{H}_1$ meets Y, the endvertex y of Q in Y is in $B_G(R_1, d_1)$. Thus, $v_0 \in X$ and $y \in Y$ have distance at most $j - 2d_2 - 1 + d_1 < n$ in G, which contradicts that P has length n and is a shortest X–Y path in G.

Therefore, we may assume there is an ordering $H_1, \dots, H_m$ of a subset of $\mathcal{H}'$ such that $I(H_1), \dots, I(H_m)$ is an interlaced system with buffer d_2. By Lemma 10.2.2, we may also assume that $I(H_1), \dots, I(H_m)$ is clean. Let $S \subseteq [m-1]$ be the set of indices i such that $b(H_i) - a(H_{i+1}) \leqslant 30d$. For each $i \in S$, there must be some component C_i of $G - B_G(P, d_1)$ such that $a(C_i) \leqslant a(H_{i+1}) - d_2$ and $b(C_i) \geqslant b(H_i) + d_2$. Otherwise, we may take z to be $v_{a(H_{i+1})}$, since $cd \geqslant d_2 + 30d + d_1$. Note that $C_i \notin \{H_1, \dots, H_m\}$ for all $i \in S$. Moreover, recall that for each $j \in [m]$, $I(H_j)$ is a maximal element of $I(\mathcal{H})$ under inclusion. Therefore, $a(H_i) < a(C_i)$ and $b(C_i) < b(H_{i+1})$ for all $i \in S$. Together, we have $a(H_i) < a(C_i) < a(H_{i+1})$ and $b(H_i) < b(C_i) < b(H_{i+1})$. Since the $a(H_i)$ and the $b(H_i)$ are strictly increasing, this also implies that $C_{i_1} \neq C_{i_2}$ for all distinct $i_1, i_2 \in S$. Suppose $|S| = s$. Re-index the elements in $\mathcal{M} := \{H_1, \dots, H_m\}$ and $\mathcal{S} := \{C_1, \dots, C_s\}$ so that all indices in $[m+s]$ appear exactly once, $a(H_{i-1}) \leqslant a(C_i)$ and $b(C_i) \leqslant b(H_{i+1})$ for all C_i, and the re-indexing preserves the original order of $H_1, \dots, H_m$ and of $C_1, \dots C_s$. Rename the corresponding sequence $H'_1, \dots H'_{m+s}$.

Claim 1. *$d_G(v_{b(H'_j)}, v_{a(H'_{j+1})}) \geqslant 30d$ for every $j \in [m+s-1]$.*

Proof. If $H'_j, H'_{j+1} \in \mathcal{M}$, then the claim follows directly from the choice of $\mathcal{S}$. Thus, we may assume that $H'_j \in \mathcal{M}$ and $H'_{j+1} \in \mathcal{S}$, since the case $H'_j \in \mathcal{S}$ and $H'_{j+1} \in \mathcal{M}$ is analogous. Then we have that $a(H'_{j+1}) + d_2 \leqslant a(H'_{j+2}) \leqslant b(H'_j)$. Thus, the claim follows from $d_2 \geqslant 30d$. ■

Claim 2. *$d_G(v_{a(H'_j)}, v_{a(H'_{j+2})}) \geqslant d_2$ for every $j \in [m+s-2]$.*

Proof. If at least two of $\{H'_j, H'_{j+1}, H'_{j+2}\}$ are in $\mathcal{M}$, then the claim follows from the choice of $\mathcal{M}$. Otherwise, $H'_j \in \mathcal{S}$ and $H'_{j+2} \in \mathcal{S}$. Then the claim follows from the choice of $\mathcal{S}$. ■

Claim 3. *If $q \geqslant p + 5$, then $d_G(H'_p, H'_q) \geqslant d_2 - 2d_1 + 2 - 30d$. Moreover, if $q \geqslant p + 7$, then $d_G(H'_p, H'_q) \geqslant d_2 - 2d_1 + 2$.*

Proof. Let $q \geqslant p + 5$ and let Q be a shortest $H'_p - H'_q$ path. Let h_p and h_q be the endvertices of Q in H'_p and H'_q, respectively, and let h'_p be the neighbour of h_p on Q and h'_q the neighbour of h_q on Q. Since both H'_p and H'_q are components of $G - B_G(P, d_1)$, we have $h'_p \in B_G(v_\ell, d_1)$ for some $\ell \leqslant b(H'_p)$ and $h'_q \in B_G(v_r, d_1)$ for some $r \geqslant a(H'_q)$. Then

$$d_G(v_\ell, v_r) \leqslant d_G(v_\ell, h'_p) + d_G(h'_p, h'_q) + d_G(h'_q, v_r) = 2d_1 + (||Q|| - 2).$$

Suppose there are at least four elements of $\mathcal{M}$ among $H'_p, \dots, H'_q$. Since $(I(H_i))_{i=1}^k$ is a clean interlaced system with buffer d_2, it follows that $r - \ell \geqslant d_2$. Since P is a shortest X–Y path, we thus obtain $d_G(v_\ell, v_r) = r - \ell \geqslant d_2$. Together with $d_G(v_\ell, v_r) \leqslant 2d_1 + (||Q|| - 2)$, this yields $||Q|| \geqslant d_2 - 2d_1 + 2$.

Note that if $q \geqslant p + 7$, then there are at least four elements of $\mathcal{M}$ among $H'_p, \dots, H'_q$. Thus, we may assume that $q \in \{p + 5, p + 6\}$ and that are at most three elements of $\mathcal{M}$ among $H'_p, \dots, H'_q$. Note that this can only happen if the sequence $H'_p, \dots, H'_q$ alternates between elements in $\mathcal{M}$ and $\mathcal{S}$. By symmetry, we may assume that $H'_p \in \mathcal{S}$. We have $\ell \leqslant b(H'_p) \leqslant b(H'_{p+1})$ and $a(H'_{p+5}) \leqslant a(H'_q) \leqslant r$. Since $H'_{p+4} \in \mathcal{S}$, $a(H'_{p+4}) \leqslant a(H'_{p+5}) - d_2$. Since $H'_{p+2} \in \mathcal{S}$, it follows that $d_G(v_{b(H'_{p+1})}, v_{a(H'_{p+3})}) \leqslant 30d$. Therefore, $d_G(v_\ell, v_r) \geqslant d_2 - 30d$. On the other hand, $d_G(v_\ell, v_r) \leqslant 2d_1 + (||Q|| - 2)$, as above. Thus, we have $||Q|| \geqslant d_2 - 2d_1 + 2 - 30d$, as required. ■

Let $H' = \bigcup_{i=1}^{m+s} H'_i$. A *fruit tree* is a subgraph $\mathcal{W} := \bigcup_{i \in I} H'_i \cup W_1 \cup \dots \cup W_k$ of G (see Figure 10.5) where

- I is a subset of $[m + s]$,
- for all $i \in [m + s]$ and $j \in [k]$, if $V(H'_i) \cap V(W_j) \neq \emptyset$, then $i \in I$,
- for each $i \in [k]$, W_i is a $\left(\bigcup_{i \in I} H'_i \cup \bigcup_{j=1}^{i-1} W_j\right)$-path with at least one edge, and
- the minor $M(\mathcal{W})$ of $\mathcal{W}$ obtained by contracting H'_i to a single vertex for each $i \in I$ is a tree.

We call the paths W_i the *composite* paths of $\mathcal{W}$. Since $M(\mathcal{W})$ is a tree, the endvertices of W_i are in distinct components of $\bigcup_{j \in I} H'_j \cup \bigcup_{j=1}^{i-1} W_j$; in particular, they are in distinct components of $\bigcup_{j=1}^{i} W_j \setminus E(W_i)$. For each $i \in [k]$, let n_i and m_i be the number of composite paths of the two components of $\bigcup_{j=1}^{i} W_j \setminus E(W_i)$ containing the endvertices of W_i. We say that $\mathcal{W}$ is *d-small* if W_i has length at most $(\max\{n_i, m_i\} + 2)d$, for all $i \in [k]$.

Claim 4. *Let $\mathcal{W}$ be a d-small fruit tree with composite paths $W_1, \dots, W_k$ and let C be a component of $\bigcup_{i=1}^{k} W_i$ consisting of ℓ composite paths. Then for every $v \in V(C)$, $d_{\mathcal{W}}(v, V(H')) \leqslant \ell d$.*

Proof. We proceed by induction on ℓ. Suppose $\ell = 1$. Then $C = W_i$ for some $i \in [k]$. Since $\mathcal{W}$ is d-small, $||W_i|| \leqslant 2d$, and both endvertices of W_i are in $V(H')$. Therefore, one of the two subpaths of W_i from v to $V(H')$ will have length at most d.

Now suppose that $\ell \geqslant 2$. Let $W_{i_1}, \dots, W_{i_\ell}$ with increasing $i_j \in [k]$ be all the composite paths of $\mathcal{W}$ contained in C. Since $M(\mathcal{W})$ is a tree, the endvertices of W_{i_ℓ} are in distinct components C_1 and C_2 of $\bigcup_{j=1}^{\ell} W_{i_j} \setminus E(W_{i_\ell})$. For each $i \in [2]$, let w_i be the endvertex of W_{i_ℓ} in C_i and let ℓ_i be the number of composite paths in C_i. Note that $\ell_1 + \ell_2 \leqslant \ell - 1$. By induction, we may assume that $v \in V(W_{i_\ell})$. Moreover, $d_{C_1}(w_1, V(H')) \leqslant \ell_1 d$ and $d_{C_2}(w_2, V(H')) \leqslant \ell_2 d$. Therefore, we are done unless $d_{W_{i_\ell}}(v, w_1) > (\ell - \ell_1)d$ and $d_{W_{i_\ell}}(v, w_2) > (\ell - \ell_2)d$. But this implies $||W_{i_\ell}|| > (\ell - \ell_1)d + (\ell - \ell_2)d \geqslant (\ell + 1)d \geqslant (\max\{\ell_1, \ell_2\} + 2)d$, which contradicts that $\mathcal{W}$ is d-small. ■

Claim 5. *Let $\mathcal{W}$ be a d-small fruit tree with composite paths $W_1, \ldots, W_k$. Then each component of $\bigcup_{i=1}^{k} W_i$ consists of at most 4 composite paths.*

Proof. Suppose that some component of $\bigcup_{i=1}^{k} W_i$ consists of composite paths $W_{i_1}, \ldots, W_{i_\ell}$ with $\ell \geqslant 5$ and increasing $i_j \in [k]$. Let $\ell' \in [\ell]$ be the smallest index such that $W_{i_1} \cup \cdots \cup W_{i_{\ell'}}$ has a component C with at least 5 composite paths.

Let c be the number of composite paths of C. Since $M(\mathcal{W})$ is a tree, $V(H') \cap V(C)$ contains $c+1$ vertices, no two of which are in the same H'_i. Thus, there exist $p, q \in [m+s]$ such that $q \geqslant p + c$ and C contains vertices $u_p \in V(H'_p)$ and $u_q \in V(H'_q)$.

Let C_1 and C_2 be the two components of $\bigcup_{j=1}^{i_{\ell'}} W_j \setminus E(W_{i_{\ell'}})$ containing the endvertices of $W_{i_{\ell'}}$. Let $i \in \{1, 2\}$. We denote by c_i the number of composite paths of C_i. The minimality of ℓ' implies that $c_i \leqslant 4$. Since $\mathcal{W}$ is d-small, the length of $W_{i_{\ell'}}$ is at most $(4+2)d$ and we can bound the number of edges in C_i from above by $\sum_{i=1}^{c_i} (i+1)d$. Thus, we have

$$d_C(u_p, u_q) \leqslant \sum_{i=1}^{c_1} (i+1)d + (4+2)d + \sum_{i=1}^{c_2} (i+1)d$$

by joining u_p to $W_{i_{\ell'}}$ in C_1, walking along $W_{i_{\ell'}}$ and then joining $W_{i_{\ell'}}$ to u_q in C_2.

If $c \geqslant 7$, then it follows from $c_1, c_2 \leqslant 4$ that

$$d_C(u_p, u_q) \leqslant \sum_{i=1}^{4} (i+1)d + (4+2)d + \sum_{i=1}^{4} (i+1)d = 34d < d_2 - 2d_1 + 2.$$

However, this contradicts Claim 3.

If $c \in \{5, 6\}$, then it follows from $c_1 + c_2 + 1 = c \leqslant 6$ that

$$d_C(u_p, u_q) \leqslant \sum_{i=1}^{4} (i+1)d + (4+2)d + \sum_{i=1}^{1} (i+1)d = 22d < d_2 - 2d_1 + 2 - 30d.$$

Again, this is a contradiction to Claim 3. ■

For each fruit tree $\mathcal{W} = \bigcup_{i \in I} H'_i \cup W_1 \cup \cdots \cup W_k$, we let $\min(\mathcal{W}) := \min\ I = \min\{i \in [m+s] : H'_i \subseteq \mathcal{W}\}$ and $\max(\mathcal{W}) := \max\ I = \max\{i \in [m+s] : H'_i \subseteq \mathcal{W}\}$.

An *orchard* is a sequence $\mathcal{W}^1, \ldots, \mathcal{W}^t$ of d-small fruit trees such that $\min(\mathcal{W}^1) = 1$, $\max(\mathcal{W}^t) = m+s$, and $\max(\mathcal{W}^i) + 1 = \min(\mathcal{W}^{i+1})$ for all $i \in [t-1]$. Note that each H'_i is itself a d-small fruit tree. Thus, $H'_1, \ldots, H'_{s+m}$ is an orchard.

Among all orchards, we choose $\mathcal{W}^1, \ldots, \mathcal{W}^t$ such that t is minimal.

For each $i \in [t]$, let $w(i) = \max(\mathcal{W}^i)$. For each $i \in [t-1]$ let B_i be a shortest $\mathcal{W}^i - \beta_i$ path, where $\beta_i := v_{b(H'_{w(i)})}$ and let A_i be a shortest $\alpha_i - \mathcal{W}^{i+1}$ path, where $\alpha_i := v_{a(H'_{w(i)+1})}$. For each $i \in [t-1]$, let r_i be the endvertex of B_i in $\mathcal{W}^i$ and ℓ_{i+1} be the endvertex of A_i in $\mathcal{W}^{i+1}$. Choose $\ell_1 \in X \cap V(H'_1)$ and $r_t \in Y \cap V(H'_{m+s})$ arbitrarily, and set $\alpha_t := v_n$ and $\beta_0 := v_0$. For each $i \in [t]$, let Q_i be an $\ell_i - r_i$ path in $\mathcal{W}^i$. See also Figure 10.5.

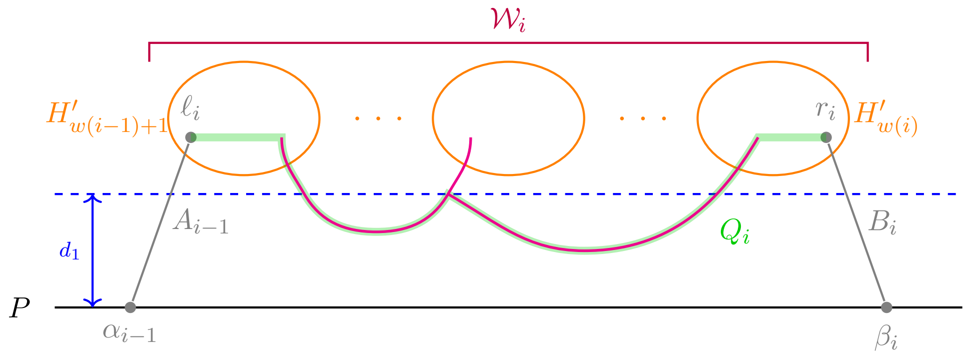

FIGURE 10.5: The fruit tree $\mathcal{W}_i$ of the orchard with its corresponding components (orange) and its composite paths (magenta).

We will now recursively define two X–Y walks P_1 and P_2 such that $d_G(P_1, P_2) \geqslant d$. These walks then contain the desired paths. In order to show that our construction indeed yields $d_G(P_1, P_2) \geqslant d$, we will need to ensure that $d_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

Initialize $P_1 := P\alpha_1$ and $P_2 := \ell_1 Q_1 r_1$. Now suppose that we have already defined P_1 and P_2 such that, for some $i \in [t]$,

(i) P_1 starts in v_0 and ends in α_i, and P_2 starts in ℓ_1 and ends in r_i; or P_1 starts in v_0 and ends in r_i, and P_2 starts in ℓ_1 and ends in α_i,

(ii) $P_1 \cup P_2$ does not contain any vertices of $\bigcup_{j=w(i)+1}^{m+s} H'_j$,

(iii) among all components of $(P_1 \cup P_2) \cap P$, the component P' containing α_i occurs last on P (note that α_i could be the left or right endvertex of this path),

(iv) the other endvertex of the component P' is some β_j with $j < i$,

(v) either α_i occurs after β_j along P, or $d_P(\alpha_i, \beta_j) \leqslant 20d$ and if additionally $d_P(\alpha_i, \beta_j) > 10d$, then α_{i+1} occurs between α_i and β_j along P, and

(vi) $d_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

We maintain these properties throughout the construction. Note that property (vi) suffices in order to prove that the final P_1 and P_2 have distance at least d in G. The other properties ensure that we can update the paths in the induction step without losing property (vi).

Now we describe how to extend P_1 and P_2. Suppose P_1 ends in α_i, and P_2 ends in r_i (the construction is analogous in the other case). Further, let β_j be the other endvertex of the component of $P_1 \cap P$ which contains α_i. If $i = t$, then P_1 and P_2 are our desired X–Y walks, since $r_t, \alpha_t \in Y$.

In order to ensure that $d_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$ we need to distinguish several cases. In each case we update

$$P_1 := P_1\alpha_i P\alpha_k A_k \ell_{k+1} Q_{k+1} r_{k+1} \text{ and } P_2 := P_2 r_i B_i \beta_i P \alpha_{k+1}$$

for some carefully chosen $k \in \{i, i+1, i+2, i+3\}$ where α_k occurs between α_i and β_i along P (see Figure 10.6).

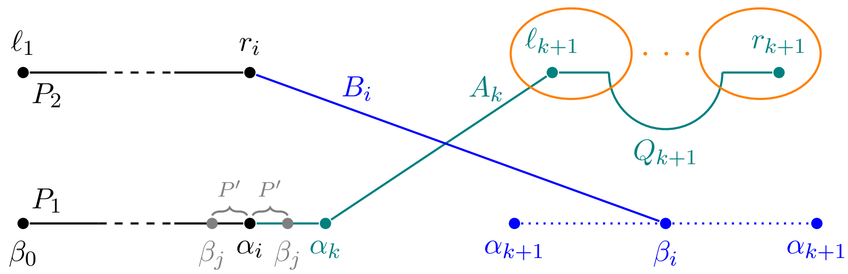

Figure 10.6: The old paths P_1 and P_2 in black together with their updates in teal and blue, respectively. Note that α_{k+1} may occur before or after β_i and that β_j may occur before or after α_i.

Then properties (i) and (ii) are clear from the definition of P_1 and P_2. By (iii), the component of $(P_1\alpha_i \cup P_2 r_i) \cap P$ containing α_i occurs last on P. Since α_{k+1} occurs after α_i and α_k along P and $\alpha_i P_1 r_{k+1} \cap P = \alpha_i P \alpha_k$, this implies that (iii) still holds for P_1 and P_2. To see that (iv) holds, we first note that once we have shown (vi), it follows that the component P' of $(P_1 \cup P_2) \cap P$ containing α_{k+1} is a subset of P_2. Since β_i occurs after α_i along P and $\alpha_i \in P_1$, property (iii) of $P_1\alpha_i$ and $P_2 r_i$ implies that P' has endvertices β_i and α_{k+1}. We are thus left to check (v) and (vi) in the individual cases. For this, we remark that since $d_P(\beta_i, \{\alpha_i, \beta_j\}) \geqslant 30d - 20d = 10d$ by assumption and Claim 1, it suffices to ensure that $d_P(\beta_i, \alpha_k) \geqslant 10d$ if $k \neq i$ and that $d_P(\alpha_{k+1}, \{\alpha_k, \beta_j\}) \geqslant 10d$ in order to obtain (vi).

To find the suitable $k \in \{i, i+1, i+2, i+3\}$ for the update of P_1 and P_2, we consider where $\alpha_i, \alpha_{i+1}, \alpha_{i+2}, \alpha_{i+3}$ occur relative to β_i on P. We depict the cases in Figures 10.7 to 10.10. In each of the figures, we indicate the corresponding assumptions in orange and their implications in green.

Suppose α_{i+1} occurs after β_i along P. Then $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant d_P(\beta_i, \{\alpha_i, \beta_j\}) \geqslant 10d$, and we set $k := i$ (see Figure 10.7). By assumption, (v) and (vi) are clear.

Now suppose that α_{i+1} occurs before β_i along P. Further, suppose α_{i+2} occurs after β_i along P (see Figure 10.8). In this case, we proceed as follows. If $d_P(\alpha_{i+1}, \beta_i) \geqslant 10d$, then we set $k := i+1$ (see Figure 10.8a). Again, (v) and (vi) are clear from the assumption. We remark that by assumption this in particular is the case if $d_P(\alpha_i, \beta_j) > 10d$ and β_j occurs after α_i along P,

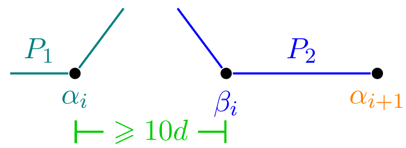

FIGURE 10.7: α_{i+1} occurs after β_i along P.

since then α_{i+1} occurs between α_i and β_j along P and, as shown above, $d_P(\beta_i, \beta_j) \geqslant 10d$.

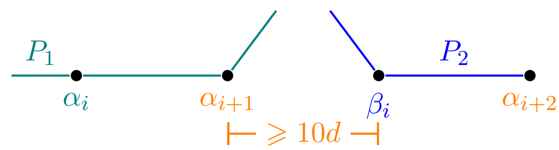

(a) $d_P(\alpha_{i+1}, \beta_i) \geqslant 10d$

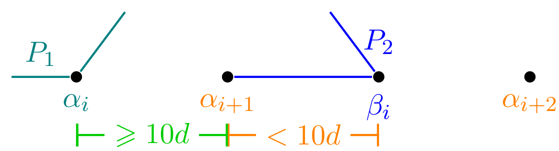

(b) $d_P(\alpha_{i+1}, \beta_i) < 10d$

FIGURE 10.8: α_{i+1} occurs before β_i along P and α_{i+2} occurs after β_i along P.

Otherwise, $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant 10d$. Indeed, either β_j occurs before α_i along P and thus we have $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) = d_P(\alpha_{i+1}, \alpha_i) = d_P(\alpha_i, \beta_i) - d_P(\alpha_{i+1}, \beta_i) \geqslant 30d - 10d \geqslant 10d$ by Claim 1, or β_j occurs after α_i along P and thus we have $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) = d_P(\alpha_{i+1}, \beta_j) = d_P(\alpha_i, \beta_i) - d_P(\alpha_{i+1}, \beta_i) - d_P(\alpha_i, \beta_j) \geqslant 30d - 10d - 10d \geqslant 10d$. In this case we set $k := i$ (see Figure 10.8b). Then (vi) holds by the previous calculation. Moreover, recall that $d_P(\alpha_{k+1}, \beta_i) < 10d$ by assumption, and thus also (v) holds.

We may thus assume that both α_{i+1} and α_{i+2} occur before β_i along P. Suppose further that α_{i+3} occurs after β_i along P (see Figure 10.9). In this case, we proceed as follows. If $d_P(\alpha_{i+2}, \beta_i) \geqslant 10d$, then we set $k := i + 2$ (see Figure 10.9a). Then (v) and (vi) are clear from the assumption.

Otherwise, if $d_P(\alpha_{i+2}, \beta_i) < 10d$, we have $d_P(\alpha_{i+2}, \{\alpha_i, \beta_j\}) \geqslant d_P(\alpha_{i+2}, \alpha_i) - 20d \geqslant d_2 - 20d \geqslant 10d$ by assumption and Claim 2. If $d_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant 10d$, then we set $k := i + 1$ (see Figure 10.9b). Then (vi) holds by assumption. Since also $d_P(\alpha_{k+1}, \beta_i) < 10d$ by assumption, (v) holds too.

Otherwise, $d_P(\alpha_{i+1}, \beta_i) = d_P(\alpha_{i+1}, \alpha_{i+2}) + d_P(\alpha_{i+2}, \beta_i) \leqslant 20d$ and $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant d_P(\alpha_i, \alpha_{i+2}) - d_P(\alpha_{i+1}, \alpha_{i+2}) - 20d \geqslant d_2 - 10d - 20d \geqslant 10d$ by assumption and Claim 2, and we then set $k := i$ (see Figure 10.9c). Then (vi) holds by the previous calculation. Moreover, (v)

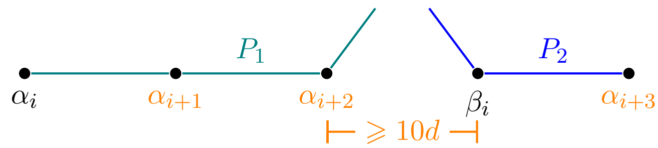

(a) $d_P(\alpha_{i+2}, \beta_i) \geqslant 10d$

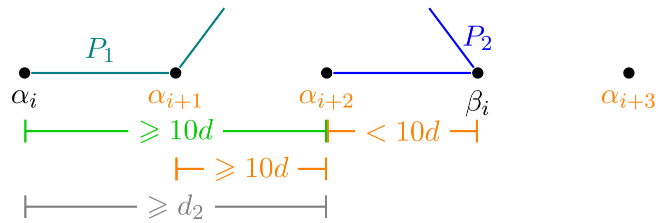

(b) $d_P(\alpha_{i+2}, \beta_i) < 10d$ and $d_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant 10d$

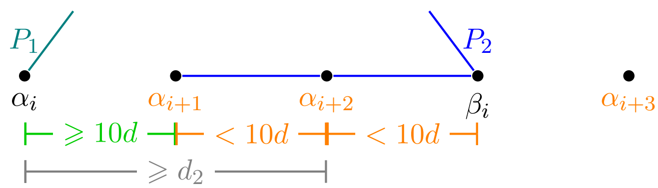

(c) $d_P(\alpha_{i+2}, \beta_i) < 10d$ and $d_P(\alpha_{i+1}, \alpha_{i+2}) < 10d$

FIGURE 10.9: α_{i+1} and α_{i+2} occur before β_i along P and α_{i+3} occurs after β_i along P.

holds since $\alpha_{k+2} = v_{a(H'_{w(k+2)+1})}$ occurs between $\alpha_{k+1} = v_{a(H'_{w(k+1)+1})}$ and β_i by assumption and because $a(H'_{w(k+1)+1}) < a(H'_{w(k+2)+1})$.

Thus, we may assume that α_{i+1}, α_{i+2} and α_{i+3} occur before β_i along P. We claim that α_{i+4} occurs after β_i along P and that $d_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant d_2$. Indeed, the fact that $\alpha_{i+3} = v_{a(H'_{w(i+3)+1})}$ occurs before $\beta_i = v_{b(H'_{w(i)})}$ implies that no three of $H'_{w(i)}, \ldots, H'_{w(i+3)+1}$ are in $\mathcal{M}$ since $\mathcal{M}$ is clean. Note that $w(i+3)+1 \geqslant w(i)+4$, since the $w(i)$ are strictly increasing by definition. By construction of $(H'_i)_{i \in [m+s]}$, we have that if $H'_i \in \mathcal{S}$ then $H'_{i+1} \in \mathcal{M}$. Hence, $w(i+3)+1 = w(i)+4$, and $H'_{w(i)}, H'_{w(i)+2}, H'_{w(i)+4} \in \mathcal{S}$, and $H'_{w(i)+1}, H'_{w(i)+3} \in \mathcal{M}$. In particular, $\mathcal{W}^{i+1} = H'_{w(i)+1}, \mathcal{W}^{i+2} = H'_{w(i)+2}$ and $\mathcal{W}^{i+3} = H'_{w(i)+3}$ consist of single components of H'. It follows that $d_P(\alpha_{i+1}, \alpha_{i+2}) = d_P(v_{a(H'_{w(i)+2})}, v_{a(H'_{w(i)+3})}) \geqslant d_2$ by the choice of $\mathcal{S}$. Moreover, as $H'_{w(i)+4} \in \mathcal{S}$, we have $H'_{w(i)+5} \in \mathcal{M}$. Together with $H'_{w(i)+1}, H'_{w(i)+3} \in \mathcal{M}$, this implies that $a(H'_{w(i+4)+1}) \geqslant a(H'_{w(i)+5}) \geqslant b(H'_{w(i)+1}) \geqslant b(H'_{w(i)})$ as $\mathcal{M}$ is clean. Thus, $\alpha_{i+4} = v_{a(H_{w(i+4)+1})}$ occurs after $\beta_i = v_{b(H'_{w(i)})}$ along P.

If $d_P(\alpha_{i+2}, \beta_i) \leqslant 20d$, then we set $k := i+1$ (see Figure 10.10a). Since $d_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant d_2 \geqslant 10d$ as shown above, (vi) holds. Moreover, $\alpha_{k+2} = v_{a(H'_{w(k+2)+1})}$ occurs between $\alpha_{k+1} = v_{a(H'_{w(k+1)+1})}$ and β_i by assumption and because $a(H'_{w(k+1)+1}) < a(H'_{w(k+2)+1})$. Thus (v) holds.

If $d_P(\alpha_{i+3}, \beta_i) \geqslant 10d$, then we set $k := i+3$ (see Figure 10.10b). Then (v) and (vi) are clear

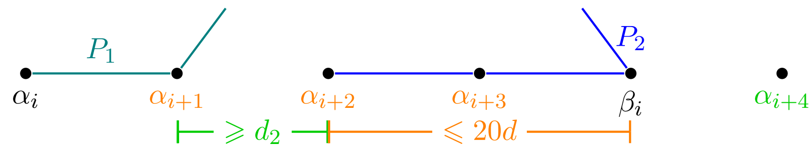

(a) $d_P(\alpha_{i+2}, \beta_i) \leqslant 20d$

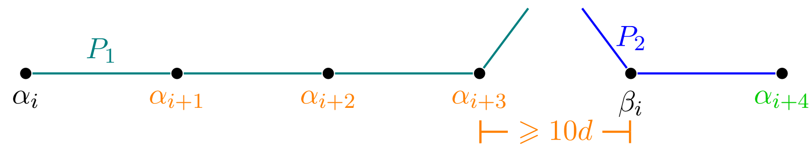

(b) $d_P(\alpha_{i+3}, \beta_i) \geqslant 10d$

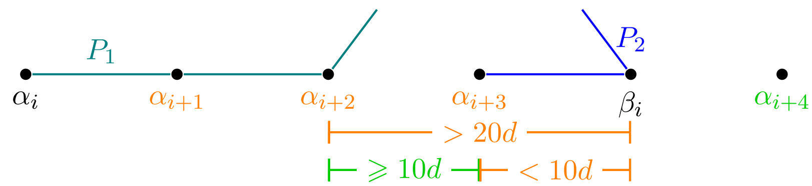

(c) $d_P(\alpha_{i+2}, \beta_i) > 20d$ and $d_P(\alpha_{i+3}, \beta_i) < 10d$

FIGURE 10.10: α_{i+1}, α_{i+2} and α_{i+3} occur before β_i along P.

by assumption.

Otherwise, $d_P(\alpha_{i+2}, \alpha_{i+3}) = d_P(\alpha_{i+2}, \beta_i) - d_P(\alpha_{i+3}, \beta_i) \geqslant 20d - 10d = 10d$. In this case we set $k := i + 2$ (see Figure 10.10c). Then (vi) holds by the previous calculation. Moreover, (v) holds since $d_P(\alpha_{i+3}, \beta_i) < 10d$ by assumption.

Claim 6. $d_G(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

Proof. By construction, we have $d_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$. Since P is a shortest X–Y path, the claim immediately follows. □

Claim 1. $d_G(P_1, P_2) \geqslant d$.

Proof. Suppose for a contradiction that $p_1 \in V(P_1)$, $p_2 \in V(P_2)$, and $d_G(p_1, p_2) \leqslant d - 1$. If $p_1 \in V(Q_i)$ and $p_2 \in V(Q_j)$ for some $i \neq j$, then $d_G(p_1, p_2) \geqslant d$ by the minimality of t. Suppose $p_1 \in V(P)$ and $p_2 \in V(Q_i)$ for some i. By Claim 4 and Claim 5, $d_G(p_2, H') \leqslant 4d$. Since $d_G(P, H') \geqslant d_1$, we have $d_G(p_1, p_2) \geqslant d_1 - 4d \geqslant d$. Suppose $p_1 \in V(P)$ and $p_2 \in V(A_i)$ for some i. By Claim 6, $d_G(p_1, \alpha_i) \geqslant 10d$. This contradicts $d_G(p_1, \alpha_i) \leqslant d_G(p_1, p_2) + d_G(p_2, \alpha_i) \leqslant d - 1 + d_1 < 10d$. The same argument handles the case $p_1 \in V(P)$ and $p_2 \in V(B_i)$ for some i.

Suppose $p_1 \in V(A_i)$ and $p_2 \in V(A_j)$ for some $i < j$. If $d_G(p_1, \ell_i) + d_G(p_2, \ell_j) \leqslant d + 1$, then

$$d_G(\ell_i, \ell_j) \leqslant d_G(\ell_i, p_1) + d_G(p_1, p_2) + d_G(p_2, \ell_j) \leqslant (d + 1) + (d - 1) = 2d,$$

which contradicts the minimality of t. Otherwise,

$$\begin{aligned} d_G(\alpha_i, \alpha_j) &\leqslant d_G(\alpha_i, p_1) + d_G(p_1, p_2) + d_G(p_2, \alpha_j) \\ &= (d_G(\alpha_i, \ell_i) - d_G(\ell_i, p_1)) + d_G(p_1, p_2) + (d_G(\alpha_j, \ell_j) - d_G(\ell_j, p_2)) \\ &\leqslant 2d_1 + 2 - (d+2) + (d-1) < 10d, \end{aligned}$$

which contradicts $d_G(\alpha_i, \alpha_j) \geqslant 10d$ by Claim 6. The same argument handles the case $p_1 \in V(B_i)$ and $p_2 \in V(B_j)$, and $p_1 \in V(A_i)$ and $p_2 \in V(B_j)$, for some $i \neq j$.

Suppose $p_1 \in V(Q_i)$ and $p_2 \in V(A_j)$ for some i, j. Since $d_G(H', P) = d_1 + 1$ and $||A_j|| \leqslant d_1 + 1$,

$$d_G(P, p_2) + d_G(p_2, \ell_j) = ||A_j|| \leqslant d_G(H', P) \leqslant d_G(H', p_1) + d_G(p_1, p_2) + d_G(p_2, P).$$

Moreover, assume that the component of $\mathcal{W}^i \setminus E(H')$ which contains p_1 consists of exactly c composite paths. By Claim 4, $d_{\mathcal{W}^i}(p_1, V(H')) \leqslant cd$. Altogether, we have $d_G(p_2, \ell_j) \leqslant d_G(H', p_1) + d_G(p_1, p_2) \leqslant cd + d - 1$. Therefore,

$$d_G(p_1, \ell_j) \leqslant d_G(p_1, p_2) + d_G(p_2, \ell_j) \leqslant cd + 2d - 2 \leqslant (c+2)d.$$

However, this contradicts the minimality of t, since there is $\mathcal{W}^i - \mathcal{W}^j$ path of length at most $(c+2)d$.

The same argument also handles the case $p_1 \in V(Q_i)$ and $p_2 \in V(B_j)$ for some i, j. By symmetry between p_1 and p_2, there are no more remaining cases. ■

This concludes the proof. □

11 A structural duality for path-decompositions into parts of small radius

It is an easy observation that if a graph G admits a path-decomposition whose parts have small radius, then G contains no large subdivision of $K_{1,3}$ or K_3 as a (quasi-) geodesic subgraph. We show that these are in fact the only obstructions to such path-decompositions of small radial width, and we prove analogous results for decompositions modelled on cycles and subdivided stars instead of paths.

With our results we confirm in a strong form a conjecture of Georgakopoulos and Papasoglu on fat-minor-characterisations of graphs quasi-isometric to paths, cycles and paths, and subdivided stars, respectively. For this, we present a novel view on quasi-isometries between graphs by graph-decompositions of bounded radial width and spread. This new perspective enables us to prove further results in coarse graph theory, and may thus be of independent interest.

This chapter is based on [8] and joint work with Reinhard Diestel, Ann-Kathrin Elm, Eva Fluck, Raphael W. Jacobs, Paul Knappe and Paul Wollan.

11.1 Introduction

All graphs in this chapter are finite.

11.1.1 Fat minors and quasi-isometries

Following Gromov's ideas on coarse geometry [80] into the realm of graphs, Georgakopoulos and Papasoglu [75] suggested a study of graphs from a coarse or metric perspective, which revolves around the concept of *quasi-isometry*. Roughly speaking, two metric spaces are quasi-isometric if their large-scale geometry coincides, and more formally, a quasi-isometry is a generalisation of bi-Lipschitz maps that allows for an additive error (see Section 9.2 for the definition).

As their favourite problem of metric graph-theoretic flavour, Georgakopoulos and Papasoglu proposed the following conjecture, whose qualitative converse is immediate:

Conjecture 11.1.1 (Fat minor conjecture [75])**.** *Let $\mathcal{X}$ be a finite set of finite graphs. Then there exists a function $f : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that, for all integers K, every graph with no K-fat X minor for any $X \in \mathcal{X}$ is $f(K)$-quasi-isometric to a graph with no X minor for any $X \in \mathcal{X}$.*

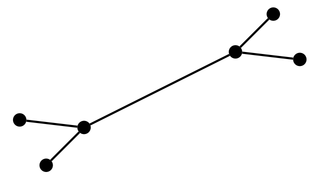

FIGURE 11.1: The ‘wrench’/‘spanner’ graph W.

Here, ‘K-fat minors’ are a metric variant of minors: Roughly speaking, a K-fat X minor is an X minor with additional distance constraints: its branch sets and branch paths are pairwise at least K apart, except for incident vertex-edge pairs (see Section 9.5 for the formal definition).

Georgakopoulos and Papasoglu [75] verified their conjecture for $\mathcal{X} = \{K_3\}$ (which describes graphs that are quasi-isometric to a forest), as well as for $\mathcal{X} = \{K_{1,m}\}$. An earlier result of Chepoi, Dragan, Newman, Rabinovich, and Vaxes [36] yields the case $\mathcal{X} = \{K_{2,3}\}$ (quasi-isometric to an outerplanar graph). Fujiwara and Papasoglu [70] solved the $\mathcal{X} = \{K_4^-\}$-case (quasi-isometric to a cactus); see also Chapter 12 for a simpler proof. Moreover, in Chapter 12, we solve the case $\mathcal{X} = \{K_4\}$. In contrast to these positive results, Davies, Hickingbotham, Illingworth and McCarty [38] showed that Conjecture 11.1.1 is false in general.

In this chapter, we establish Conjecture 11.1.1 for three further small cases but in a stronger form:

Theorem 29. *Conjecture 11.1.1 is true for $\mathcal{X} = \{K_3, K_{1,3}\}$ (quasi-isometric to a disjoint union of paths), $\mathcal{X} = \{K_{1,3}\}$ (cycles and paths)*[1] *and $\mathcal{X} = \{K_3, W\}$ (subdivided stars), where W is the graph depicted in Figure 11.1. In all these cases, Conjecture 11.1.1 even holds true when replacing ‘fat minors’ by ‘quasi-geodesic topological minors’.*

We refer the reader to Theorems 30 to 32 for the respective functions f.

In Theorem 29, ‘quasi-geodesic topological minors’ are a metric version of topological minors, whose model in G is ‘quasi-geodesically’ embedded (see Section 11.1.4 for the definition). Roughly speaking, the metric of the topological minor agrees (up to a multiplicative constant) with the metric induced by the graph G. In particular, quasi-geodesic topological minors yield fat minors but in general not the other way around (Lemma 11.3.3).

Our work on Theorem 29 was independent of Conjecture 11.1.1 and the involved concepts such as fat minors and quasi-isometries. In fact, we only discovered the connections of our results to quasi-isometries and fat minors, and in particular to Conjecture 11.1.1, through [75]. Our approach is not based on quasi-isometries, but on the notion of graph-decompositions, which we describe in what follows.

[1] Georgakopoulos and Papasoglu independently prove the more general case $\mathcal{X} = \{K_{1,m}\}$ in a second version of their paper [75].

11.1.2 Graph-decompositions

Graph-decompositions [48] are a natural extension of tree-decompositions which allow the bags V_h of decompositions $(H, \mathcal{V})$ to be arranged along general decomposition graphs H instead of just trees. We then say that $(H, \mathcal{V})$ is *modelled on* H and call it an H*-decomposition*. Moreover, given any graph class $\mathcal{H}$, we refer to H-decompositions with $H \in \mathcal{H}$ as $\mathcal{H}$*-decompositions*. Recent applications of graph-decompositions include a local-global decomposition theorem [48] as well as the study of local separations [32] and of locally chordal graphs [1].

In this chapter, we present a width-notion for graph-decompositions such that a graph G has an H-decomposition of small width if and only if H resembles the large-scale structure of G. The naive approach defines the 'width' of a graph-decomposition analogously to tree-width, that is, as the minimal cardinality of a bag of the decomposition (minus 1). However, the respective '$\mathcal{H}$-width', the minimal width of an $\mathcal{H}$-decomposition for a given graph class $\mathcal{H}$, does not yield a meaningful extension of tree-width: if a minor-closed class $\mathcal{H}$ of graphs has bounded tree-width, then every graph of small $\mathcal{H}$-width has small tree-width itself, and if $\mathcal{H}$ has unbounded tree-width, then the $\mathcal{H}$-width of every graph is at most 2 (minus 1) [51, 112].[2]

This inspired us to consider a metric perspective instead: To define the 'width' of a graph-decomposition, we measure the size of its bags not in terms of their cardinality, but by the radius of its *parts*, the induced subgraphs on the bags of the decomposition. More formally, recall that the *radius* of a graph G is the smallest $r \in \mathbb{N}$ such that some vertex of G has distance at most r to all vertices of G. We then let the *(inner-)radial width* of a decomposition be the largest radius among its parts and define the *radial* $\mathcal{H}$*-width* of G for a given class $\mathcal{H}$ of graphs to be the smallest radial width among all decompositions of G modelled on graphs in $\mathcal{H}$.

This notion of radial width is not new for tree-decompositions. Indeed, the radial $\mathcal{H}$-width for the class $\mathcal{H}$ of all trees, or *radial tree-width* for short, has been studied before, e.g. as the equivalent *tree-length* in [56] and *tree-breadth* in [57]. Similar to the classical tree-width notion [111], several computationally hard problems such as the computation of the metric dimension of a graph [17] can be efficiently solved on graphs of small radial tree-width (for a summary, see [56, Section 1] or [98]).

11.1.3 Interplay between graph-decompositions and quasi-isometries

As it turns out, graph-decompositions and quasi-isometries are closely related. In fact, these notions become qualitatively equivalent if we restrict to 'honest' graph-decompositions of bounded radial width and bounded 'radial spread'.

[2] By Diestel and Kühn [51, Proposition 3.7], every graph G has a grid-decomposition of width at most 2 (minus 1). Thus, the $\mathcal{H}$-width of G is at most 2 (minus 1) for any graph class $\mathcal{H}$ of unbounded tree-width by Robertson and Seymour's Grid Theorem [112].

An H-decomposition of G is *honest* if all its bags are non-empty and for every edge of its decomposition graph H, the bags corresponding to its endnodes intersect. Informally speaking, being honest ensures that all connectivity in the decomposition graph H also appears in the decomposed graph G. For each vertex $v \in G$ let H_v be the induced subgraph of H on the set of all nodes $h \in H$ whose corresponding bags contain v. The *(inner-)radial spread* of the H-decomposition is then defined as the largest radius of the H_v with $v \in V(G)$.

The equivalence between the existence of an honest H-decomposition with bounded radial width and spread and of a bounded quasi-isometry to H was observed for the case of trees H by Berger and Seymour [19, 4.1]. Here, we extend their observation to arbitrary graphs H.

Proposition 11.1.2. *There exist functions $g : \mathbb{N}^2 \to \mathbb{N}^2$ and $h_1, h_2 : \mathbb{N}^2 \to \mathbb{N}$ such that the following holds for all graphs G, H:*

(i) *If G admits an honest H-decomposition of radial width r_0 and radial spread r_1, then G is $g(r_0, r_1)$-quasi-isometric to H.*

(ii) *If G is (L, C)-quasi-isometric to H, then G admits an honest H-decomposition of radial width $h_1(L, C)$ and radial spread $h_2(L, C)$.*

We refer the reader to Section 11.2 for the detailed statement on the functions g and h_1, h_2.

The correspondence given by Proposition 11.1.2 paves the way for a new proof method towards Conjecture 11.1.1: one through the graph-theoretic construction of a suitable graph-decomposition. We follow this approach in the present chapter and further demonstrate its power and versatility in Chapter 12.

Proposition 11.1.2 in particular implies that we reached our initial goal from the beginning of Section 11.1.2 with the notions of 'radial width' and 'radial spread': a graph H resembles the large-scale geometry of G (in terms of bounded quasi-isometries) if and only if G admits an honest H-decomposition of bounded radial width and spread.

11.1.4 Our results

With a suitable width measure at hand, our next goal was to identify obstructions to small radial $\mathcal{H}$-width and to characterise which graphs have small radial $\mathcal{H}$-width for given classes $\mathcal{H}$ of graphs. Similar to Conjecture 11.1.1, we considered metric versions of minors as candidates for suitable obstructions. For the graph classes $\mathcal{H}$ that we study in this chapter, our obstructions are 'quasi-geodesic topological minors', which can be seen as a special case of fat minors. Let us make this precise in what follows.

A $(\geqslant k)$-*subdivision* of a graph X is a graph X' which arises from X by subdividing each edge at least k times, i.e. replacing every edge with a path of at least $k + 1$ edges. Further, a subgraph X of a graph G is *c-quasi-geodesic*[3] for some $c \in \mathbb{N}$ if the distance of any two vertices x and y in X is at most c times their distance in G. Here, the parameter c describes how well the distances in X approximate the distances in G; in particular, a subgraph is geodesic if and only if it is 1-quasi-geodesic.

Quasi-geodesic topological minors are indeed an obstruction to small radial width: if a $(\geqslant k)$-subdivision of a graph X is a c-quasi-geodesic subgraph of G, then X is a $\lfloor \frac{k-2}{2c} \rfloor$-fat minor of G (Lemma 11.3.3). Hence, as fat minors form an obstruction to small radial width (Lemma 11.3.4), so do quasi-geodesic topological minors (Proposition 11.3.5).

Our main result, Theorem 29, asserts that these are in fact the only obstructions to small radial $\mathcal{H}$-width for the three cases where $\mathcal{H}$ consists of paths, cycles and paths, and subdivided stars, respectively. We give the detailed statements in the following Theorems 30 to 32, from which we then deduce Theorem 29.

Theorem 30 (Radial path-width)**.** *Let $k \in \mathbb{N}$. If a connected graph G contains no $(\geqslant k)$-subdivision of K_3 as a geodesic subgraph and no $(\geqslant 3k)$-subdivision of $K_{1,3}$ as a 3-quasi-geodesic subgraph, then G admits an honest decomposition modelled on a path P of radial width at most $18k+2$ and radial spread at most $18k+1$.*

Moreover, P is $(1, 18k+2)$-quasi-isometric to G.

Theorem 31 (Radial cycle-width)**.** *Let $k \in \mathbb{N}$. If a connected graph G contains no $(\geqslant 3k)$-subdivision of $K_{1,3}$ as a 3-quasi-geodesic subgraph, then G admits an honest decomposition modelled on a cycle or path C of radial width at most $18k+2$ and radial spread at most $36k+2$.*

Moreover, C is $(1, 18k+2)$-quasi-isometric to G.

Theorem 32 (Radial star-width)**.** *Let $k \in \mathbb{N}$. If a connected graph G contains no $(\geqslant k)$-subdivision of K_3 as a geodesic subgraph and no $(\geqslant 3k)$-subdivision of the wrench graph W as a 3-quasi-geodesic subgraph, then G admits an honest decomposition modelled on a subdivided star S of radial width at most $58k+9$ and radial spread at most $58k+9$.*

Moreover, there exists some $C_k \in \mathbb{N}$ such that some subdivided star is $(1, C_k)$-quasi-isometric to G.[4]

Proof of Theorem 29 given Theorems 30 to 32. Let $K \in \mathbb{N}$, and let G be a graph with no K-fat X minor for any X in the respective $\mathcal{X}$. By Lemma 11.3.3 there is an integer k depending on K and $\mathcal{X}$ only, such that, for every $c \in \mathbb{N}$, G contains no $(\geqslant c \cdot k)$-subdivision of any $X \in \mathcal{X}$ as c-quasi-geodesic subgraph. Now we deduce Theorem 29 from Theorems 30 to 32 by applying the respective

[3]Note that in metric spaces this property is often called $(c, 0)$-quasi-geodesic. In [18] it is called c-multiplicative.

[4]While we will only formally prove that some subdivided star S' and some $C_k \in \mathbb{N}$ exists, one can check by carefully reading the proof that we may choose $S' = S$ and $C_k = 60k + 14$ (see the paragraph after the proof of Theorem 32 in Section 11.6 for details).

theorem to the components of G and combine the obtained $(1, C_k)$-quasi-isometries to one from the disjoint union H of their domains to G. Lemma 9.2.1 yields the desired $f(K)$-quasi-isometry from G to H. □

11.1.5 Sketch of the proofs

For the proof of Theorem 30 (radial path-width), we start with a longest geodesic path P in the connected graph G. We then show that either balls of small radius around $V(P)$ cover all of G or we can find a geodesic ($\geqslant k$)-subdivision of the triangle K_3 or a 3-quasi-geodesic ($\geqslant 3k$)-subdivision of the claw $K_{1,3}$. In the former case we construct a P-decomposition of G by letting the bag to some node $p \in P$ be the union of all those small radius-balls, whose centre vertex has small distance to p in P.

The proof technique for Theorem 30 immediately generalizes to Theorem 31 (radial cycle-width). The proof of Theorem 32 (radial star-width) is more involved. As before we start with a longest geodesic path P in G and consider the subgraph G' of G which is covered by balls of small radius around $V(P)$. Unlike in the path-case, G' will in general not be equal to G. But if G does not contain a large subdivision of K_3 or W as a (3-quasi-)geodesic subgraph, then every component of $G - G'$ will have its neighbours in G' only close to the same vertex p_s of P. We then identify a geodesic path within each component of $G - G'$ such that all vertices in the component have small distance to this path. All these paths can then be combined with P into a subdivided star S, by adding edges between their last vertices and p. We then assign to each node s of S a suitable bag V_s of vertices of bounded distance to s.

11.1.6 An open conjecture about quasi-isometries to trees

For the special case of $\mathcal{X} = \{K_3\}$, Georgakopoulos and Papasoglu answered Conjecture 11.1.1 in the affirmative, proving that the absence of a K-fat K_3 minor implies the existence of an $f(K)$-quasi-isometry to a forest. Berger and Seymour [19] characterised the graphs quasi-isometric to a forest using a similar kind of obstruction. Our results Theorems 30 to 32 yield another natural guess for such an obstruction – long quasi-geodesic cycles:

Conjecture 11.1.3. *There is a function $f : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that if a connected graph G does not contain a c-quasi-geodesic cycle of length at least $3ck$ for some $c, k \in \mathbb{N}$, then G is $f(k)$-quasi-isometric to a tree.*

If true, this statement would strengthen the respective results of Georgakopoulos and Papasoglu and of Berger and Seymour. Note that by Proposition 11.1.2 and [19, 4.1] we can equivalently ask about the existence of a function f which guarantees that G admits an honest tree-decomposition of radial width at most $f(k)$.

11.1.7 How this chapter is organised

In Section 11.2 we show how graph-decompositions are related to quasi-isometries, proving Proposition 11.1.2. In Section 11.3 we prove that quasi-geodesic topological minors yield fat minors and are an obstruction to small radial width. The following three Sections 11.4–11.6 contain the proofs of Theorems 30 to 32, respectively.

11.1.8 Some notation

We remark that the 'outer' versions of radial width and spread are only used in Lemmas 11.2.1 and 11.2.2. Hence, we will usually omit the 'inner' when talking about inner-radial width and spread.

For classes $\mathcal{H}$ of graphs that we frequently use in this chapter, we name the (inner)-radial $\mathcal{H}$-width as in Table 11.1.

TABLE 11.1: Nomenclature for frequently used classes $\mathcal{H}$ of graphs. The second column gives the list of graphs which defines the minor-closure of $\mathcal{H}$ via forbidden (topological) minors.

$\mathcal{H}$ consists of disjoint unions of	**Forbidden (topological) minors**	**$\mathcal{H}$-decomposition**	**radial $\mathcal{H}$-width**
paths	K_3, $K_{1,3}$	path-decomposition	radial path-width
subdivided stars	K_3, W	star-decomposition	radial star-width
trees	K_3	tree-decomposition	radial tree-width
cycles	$K_{1,3}$	cycle-decomposition	radial cycle-width

11.2 Interplay between graph-decompositions and quasi-isometries

In this section we describe the interplay between quasi-isometries and honest graph-decomposition of bounded radial width and spread, i.e. we prove Proposition 11.1.2. To make this relationship more precise, we will use the following equivalent definition of quasi-isometry, which describes the parameters in (Q1) and (Q2) more precisely:

For given integers $m, M \geqslant 1$ and $a, A, r \geqslant 0$, an (m, a, M, A, r)-*quasi-isometry* from a graph H to a graph G is a map $\varphi\colon V(H) \to V(G)$ such that

(Q1'a) $d_H(h, h') \leqslant m \cdot d_G(\varphi(h), \varphi(h')) + a$ for every $h, h' \in V(H)$,

(Q1'b) $d_G(\varphi(h), \varphi(h')) \leqslant M \cdot d_H(h, h') + A$ for every $h, h' \in V(H)$, and

(Q2') for every vertex $v \in G$, there exists a node $h \in H$ with $d_G(v, \varphi(h)) \leqslant r$.

Note that every (m, a, M, A, r)-quasi-isometry is a $(\max\{m, M\}, \max\{\frac{a}{m}, A, r\})$-quasi-isometry. Conversely, an (L, C)-quasi-isometry is a (L, LC, L, C, C)-quasi-isometry.

We split the proof of Proposition 11.1.2 into the following two lemmas.

Lemma 11.2.1. *Let $(H, \mathcal{V})$ be an honest graph-decomposition of a graph G of outer-radial width r_0 and outer-radial spread r_1. Then there is a $(2r_1, 2r_1, 2r_0, 0, r_0)$-quasi-isometry from H to G.*

Proof. Since $(H, (V_h)_{h \in H})$ has outer-radial width r_0, there exists for every bag V_h a vertex $\varphi(h) \in G$ with $V_h \subseteq B_G(\varphi(h), r_0)$. We show that the respective map $\varphi : V(H) \to V(G)$ is the desired quasi-isometry from H to G. Note that the definition of φ immediately implies (Q2') by (H1).

For (Q1'b) we have to show that for every $h, h' \in V(H)$, we have

$$d_G(\varphi(h), \varphi(h')) \leqslant 2r_0 \cdot d_H(h, h').$$

So let h and h' be arbitrary vertices of H, and let $h_0 h_1 \dots h_\ell$ be a shortest h–h' path in H. We now aim to build from this path a $\varphi(h)$–$\varphi(h')$ walk in G of length at most $2r_0\ell$. For every $i \in \{1, \dots, \ell\}$, we fix a vertex $v_i \in V_{h_{i-1}} \cap V_{h_i}$; such v_i exist because the considered H-decomposition of G is honest by assumption. Furthermore, we set $v_0 := \varphi(h)$ and $v_{\ell+1} := \varphi(h')$. Now for every $i \in \{0, \dots, \ell\}$, let P_i be a shortest v_i–v_{i+1} path in G. Since our H-decomposition of G has outer-radial width r_0, the paths P_0 and P_ℓ have length at most r_0 and all other P_i have length at most $2r_0$. Thus, $v_0 P_0 v_1 \dots v_\ell P_\ell v_{\ell+1}$ is a $\varphi(h)$–$\varphi(h')$ walk in G of length at most $r_0 + 2r_0(\ell - 1) + r_0$, as desired.

For (Q1'a), we have to check that for every $h, h' \in V(H)$,

$$d_H(h, h') \leqslant 2r_1 \cdot d_G(\varphi(h), \varphi(h')) + 2r_1.$$

So let h and h' be arbitrary vertices of H and let $v_0 v_1 \dots v_\ell$ be a shortest $\varphi(h)$–$\varphi(h')$ path in G. Similar as before, we build a h–h' walk in H of length at most $2r_1\ell + 2r_1$. For every $i \in \{1, \dots, \ell\}$ we fix a node $h_i \in V(H)$ satisfying $v_{i-1}v_i \in V_{h_i}$; such h_i exist by (H1). Furthermore, we set $h_0 \coloneqq h$ and $h_{\ell+1} \coloneqq h'$. For every $i \in \{0, \dots, \ell\}$, let P_i be a shortest h_i–h_{i+1} path in H. As $h_i, h_{i+1} \in V(H_{v_i})$ for all $i \in \{0, \dots, \ell\}$ and $(H, \mathcal{V})$ has outer-radial spread r_1, all P_i have length at most $2r_1$. Thus, $h_0 P_0 h_1 \dots h_\ell P_\ell h_{\ell+1}$ is an h–h' walk in H of length at most $2r_1(\ell + 1)$, as desired. □

Lemma 11.2.2. *Let φ be an (m, a, M, A, r)-quasi-isometry from a graph H to a graph G, and let $r' \geqslant r$ be an integer. Write $B_h := B_G(\varphi(h), r')$ and set*

$$V_h := \bigcup_{h' \in B_H(h, mr' + \lceil (m+a)/2 \rceil)} B_{h'}.$$

Then $(H,(V_h)_{h\in H})$ is an honest graph-decomposition of G of outer-radial width at most $r'+M(mr'+\lceil(m+a)/2\rceil)+A$ and (inner-)radial spread at most $4mr'+m+2a+1$. Moreover, if $r'\geqslant M+A$, then even the (inner-)radial width is at most $4mr'+m+2a+1$.

Proof. We first show that the pair $(H,(V_h)_{h\in H})$ is indeed an H-decomposition of G'. It then follows immediately from the definition of the V_h that $(H,\mathcal{V})$ is honest (as $\varphi(h)\in V_h$ for all $h\in V(H)$ and $\varphi(h),\varphi(h')\in V_h\cap V_{h'}$ for all $hh'\in E(H)$ because $m\geqslant 1$).

(H1): By (Q2'), each vertex of G is contained in some B_h. Now consider an edge $e=v_0v_1$ of G. There are nodes h_0 and h_1 of H such that $v_0\in B_{h_0}$ and $v_1\in B_{h_1}$. Hence, $d_G(\varphi(h_0),\varphi(h_1))$ is at most $2r'+1$. By (Q1'a), there exists a path P in H of length at most $m(2r'+1)+a$. There exists a 'middle' vertex h on this path P such that both h_0 and h_1 are contained in $B_H(h,mr'+\lceil(m+a)/2\rceil)$, so both v_0 and v_1 are contained in $B_h\subseteq V_h$. In particular, $e=v_0v_1\in G[V_h]$, as desired.

(H2): Let v be any vertex of G, and let h_0 and h_1 be nodes of H such that $v\in V_{h_0}\cap V_{h_1}$. We want to find an h_0–h_1 path P in H such that V_h contains v for every node $h\in P$. By definition of the bags V_h, there are $h_0'\in B_H(h_0,mr'+\lceil(m+a)/2\rceil)$ and $h_1'\in B_H(h_1,mr'+\lceil(m+a)/2\rceil)$ such that $v\in B_{h_0'}\cap B_{h_1'}$. Hence, $d_G(\varphi(h_0'),\varphi(h_1'))$ is at most $2r'$. So by (Q1'a), there is an h_0'–h_1' path P' in H of length at most $2mr'+a$. Let W be the walk $P_0P'P_1$ in H joining h_0 and h_1 where P_0 is a shortest h_0–h_0' path in H and P_1 a shortest h_1'–h_1 path in H. It follows directly from the construction of W that h_0' or h_1' is contained in $B_H(h,mr'+\lceil(m+a)/2\rceil)$ for every node h visited by W. Since $v\in B_{h_0'}\cap B_{h_1'}$, this implies $v\in V_h$ for every node h visited by W by the definition of V_h. In particular, W contains a h_0–h_1 path P which is as desired.

Secondly, let us verify that $(H,(V_h)_{h\in H})$ has the desired radial spread and outer-radial width. For the radial spread, observe that the above constructed walk W has length at most $2(mr'+\lceil(m+a)/2\rceil)+2mr'+a\leqslant 2mr'+m+a+1+2mr'+a=4mr'+m+2a+1$; so the radial spread of $(H,(V_h)_{h\in H})$ is as desired.

For the outer-radial width, consider any node $h\in H$ and vertex $v\in V_h$. By definition, there is $h'\in B_H(h,mr'+\lceil(m+a)/2\rceil)$ such that $v\in B_{h'}$. By (Q1'b), we obtain

$$d_G(v,\varphi(h))\leqslant d_G(v,\varphi(h'))+d_G(\varphi(h),\varphi(h'))\leqslant r'+M(mr'+\lceil(m+a)/2\rceil)+A.$$

Thus, every $v\in V_h$ has distance at most $r'+M(mr'+\lceil(m+a)/2\rceil)+A$ from $\varphi(h)$. This yields the desired outer-radial width.

To obtain the moreover-part, let us investigate the above equation in more detail. Any shortest v–$\varphi(h')$ in G is contained in $B_h'\subseteq V_h$. Fix a shortest h'–h path Q and replace each edge h_0h_1 of Q by a shortest $\varphi(h_0)$–$\varphi(h_1)$ path $Q_{h_0h_1}$ in G to obtain a $\varphi(h)$–$\varphi(h')$ walk Q' in G. If $r'\geqslant M+A$, the path $Q_{h_0h_1}$ is contained in $B_{h_0}'\subseteq V_h$. Thus, the (inner-)radial width is already at most $r'+M(mr'+\lceil(m+a)/2\rceil)+A$. □

Proof of Proposition 11.1.2. Use Lemma 11.2.1 for (i) and Lemma 11.2.2 for (ii). □

To conclude this section, let us look at a possible path-way towards omitting the condition on the radial spread. Berger and Seymour proved that a graph has bounded radial tree-width if and only if it is quasi-isometric to a tree [19, 4.1]. More precisely, they show that if G has a T-decomposition of low radial-width for some tree T, then G is quasi-isometric to some tree T' that is obtained from a subtree of T by contracting and subdividing edges.

We ask whether such an argument transfers to arbitrary decomposition graphs:

Question 11.2.3. Given an integer $r \geqslant 1$, does there exist an integer R such that if a graph G has a decomposition modelled on a graph H of radial width at most r, then there exists an honest decomposition of G modelled on a graph H' obtained from a subgraph of H by subdividing and contracting edges such that both its radial width and radial spread are at most R?

An affirmative answer to Question 11.2.3 would in particular imply the equivalence of small radial $\mathcal{H}$-width and quasi-isometry to an element of $\mathcal{H}$ for graph classes $\mathcal{H}$ closed under takings subgraphs, and contracting and subdividing edges.

11.3 Obstructions to small radial width

In this section we study quasi-geodesic topological minors as obstructions to small radial width. We show in Lemma 11.3.3 that quasi-geodesic topological minors are also an instance of the more general fat minors used in Conjecture 11.1.1. Fat minors are obstructions to small radial width (Lemma 11.3.4), which then implies that quasi-geodesic topological minors are also such obstructions (Proposition 11.3.5).

Let us recall the definition of ($\geqslant k$)-subdivisions T_kX and quasi-geodesity from the introduction.

Definition 11.3.1 (T_kX)**.** A ($\geqslant k$)*-subdivision* of a graph X, which we denote with T_kX, is a graph which arises from X by subdividing every edge at least k times, i.e. replacing every edge in X with a new path of length at least $k+1$ such that no new path has an inner vertex in $V(X)$ or on any other new path.

The original vertices of X are the *branch vertices of the* T_kX. Note that the well-known topological minor relation can be phrased in terms of ($\geqslant 0$)-subdivisions in that a graph X is a *topological minor* of a graph G if G contains a T_0X as a subgraph.

Definition 11.3.2 (Quasi-geodesic)**.** A subgraph X of a graph G is *c-quasi-geodesic (in G)* for some $c \in \mathbb{N}$ if for every two vertices $u, v \in V(X)$ we have $d_G(u,v) \leqslant c \cdot d_X(u,v)$. We call X *geodesic* if it is 1-quasi-geodesic.

Let us now turn to the question how quasi-geodesic topological minors relate to fat minors, the obstruction investigated by Georgakopoulos and Papasoglu in Conjecture 11.1.1.

Lemma 11.3.3. *If a graph G contains a T_kX as a c-quasi-geodesic subgraph for some graph X and $c, k \in \mathbb{N}$ with $k \geqslant 2$, then X is a K-fat minor in G for $K = \lfloor \frac{k-2}{2c} \rfloor$.*

Proof. Denote the c-quasi-geodesic subgraph of G which is a T_kX by X_G. By definition, any two branch vertices of X_G have distance at least $k+1$ in X_G. For $x \in V(X)$, we then choose $V_x = B_{X_G}(x, \lceil (k+1)/4 \rceil)$, and for $e = xx' \in E(X)$, we let E_e be the (unique) V_x-V_x' path in X_G not meeting any other branch vertices; in particular, each E_e has length at least $(k+1) - 2\lceil (k+1)/4 \rceil \geqslant k/2 - 1$. We then let $\mathcal{V}$ be the set of all V_x and $\mathcal{E}$ be the set of all E_e. Since X_G is a T_kX, this construction directly yields $d_{X_G}(Y, Z) \geqslant k/2 - 1$ for every two distinct $Y, Z \in \mathcal{V} \cup \mathcal{E}$ unless $Y = E_e$ and $Z = V_x$ for some vertex $x \in V(X)$ incident to $e \in E(X)$, or vice versa. Since X_G is a c-quasi-geodesic subgraph of X, we then obtain $d_G(Y, Z) \geqslant d_{X_G}(Y, Z)/c \geqslant \frac{k-2}{2c}$; so X_G is indeed a K-fat model of X in G. □

Lemma 11.3.4. *Suppose X is a K-fat minor in a graph G for some $K \in \mathbb{N}$. If $(H, \mathcal{V})$ is a graph-decomposition of G of radial width less than $K/2$, then X is a minor of H.*

Proof. Fix a K-fat model $(\mathcal{U}, \mathcal{E})$ of X in G. We aim to define a model $(\mathcal{U}', \mathcal{E}')$ of X in H. For every vertex $x \in X$, let U_x' be the set of nodes $h \in H$ whose corresponding bag V_h meets U_x. Note that the U_x' are pairwise disjoint, as the radial width of $(H, \mathcal{V})$ is $< K/2$ and $(\mathcal{U}, \mathcal{E})$ is K-fat. Since the $G[U_x]$ are connected, the $H[U_x']$ are connected by (H2').

Analogously, for each edge xy of X, the subgraph induced on the nodes $h \in H$ whose V_h meets E_{xy} is connected by (H2'), meets U_x' and U_y', and all these subgraphs are pairwise disjoint. Thus, we may pick a U_x'–U_y' path E_{xy}' in each of these subgraphs to obtain a model of X in H. □

Combining the previous two lemmas, we obtain that quasi-geodesic topological minors are indeed obstructions to small radial width:

Proposition 11.3.5. *If a graph G contains a $T_{4ck+2}X$ as a c-quasi-geodesic subgraph for some graph X and $c, k \in \mathbb{N}$ with $k \geqslant 2$, then G admits no H-decomposition of radial width less than k modelled on a graph H with no X minor.* □

11.4 Radial path-width

In this section we prove Theorem 30, which we restate here for convenience.

Theorem 30. *Let $k \in \mathbb{N}$. If a connected graph G contains no $T_{4k+1}K_3$ as a geodesic subgraph and no $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph, then G admits an honest decomposition modelled on a path P of radial width at most $18k+2$ and radial spread at most $18k+1$.*

Moreover, P is $(1, 18k+2)$-quasi-isometric to G.

We remark that we did not optimise the bounds on the radial width and radial spread.

In fact, we show the following stronger statement, which immediately implies Theorem 30.

Theorem 11.4.1. *Let $k \in \mathbb{N}$, and let P be a longest geodesic path in a connected graph G. If G contains no $T_{4k+1}K_3$ as a geodesic subgraph and no $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph, then either P has length at most $18k+2$ or every vertex of G has distance at most $9k$ from P.*

Let us first show that Theorem 11.4.1 indeed implies Theorem 30. For this, we need the following auxiliary lemma, which asserts that taking balls of radius $r \in \mathbb{N}$ around a quasi-geodesic subgraph H of a graph G yields a 'partial decomposition' of G modelled on H (see Figure 11.2).

Lemma 11.4.2. *Let G be a graph, and let H be a c-quasi-geodesic subgraph of G for some $c \in \mathbb{N}$. Given $r \in \mathbb{N}$ we write $B_h := B_G(h, r)$ for $h \in H$ and set*

$$V_h := \bigcup_{h' \in B_H(h, cr+\lceil c/2 \rceil)} B_{h'}.$$

Then $(H, (V_h)_{h \in H})$ is an honest H-decomposition of $G' := G[\bigcup_{h \in H} B_h]$ of radial width at most $(c+1)r + \lceil c/2 \rceil$ and radial spread at most $4cr + c + 1$.

Moreover, if H is a path, then $(H, (V_h)_{h \in H})$ has radial spread at most $(c+1)r + \lceil c/2 \rceil$.

Proof. The pair $(H, (V_h)_{h \in H})$ is an honest graph-decomposition of G by Lemma 11.2.2, since the embedding φ of a c-quasi-geodesic subgraph H of G into $G[\bigcup_{h \in H} B_G(h, r)]$ is a $(c, 0, 1, 0, r)$-quasi-isometry. By Lemma 11.2.2 and its moreover-part, this graph-decomposition has (inner-)radial width $r + (cr + \lceil c/2 \rceil) = (c+1)r + \lceil c/2 \rceil$ and radial spread $4rc + c + 1$.

For the 'moreover'-part, assume that $H = h_0 \dots h_n$ is a path, and let v be any vertex of G. Let $i, j \in [n]$ be the smallest and largest integer such that $v \in B_{h_k} = B_G(h_k, r)$ for $k = i, j$. Then $d_G(h_i, h_j) \leqslant 2r$, and hence $j - i \leqslant 2cr$ since H is c-quasi-geodesic in G. By the definition of the V_h, it follows that only V_{h_k} with $i - cr + \lceil c/2 \rceil \leqslant k \leqslant j + cr + \lceil c/2 \rceil$ contain v, and hence H_v has radius at most $(c+1)r + \lceil c/2 \rceil$ as it is contained in $B_H(h_k, (c+1)r + \lceil c/2 \rceil)$ for $k := \lfloor (j-i)/2 \rfloor$. □

Proof of Theorem 30 given Theorem 11.4.1. Let G be a connected graph that contains neither $T_{3k}K_{1,3}$ as a 3-geodesic subgraph nor $T_{4k+1}K_3$ as a geodesic subgraph. Let P' be a longest geodesic path in G. By Theorem 11.4.1, either P' has length at most $18k+2$ or every vertex of G has

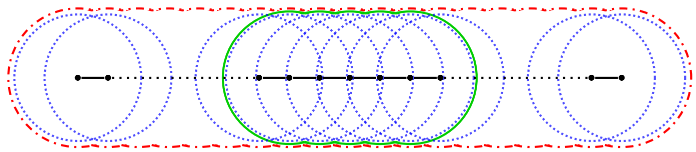

FIGURE 11.2: A decomposition modelled on the black path using the blue balls as given by Lemma 11.4.2 for $r = 2$, $c = 1$. The bag corresponding to the centre vertex is green.

distance at most $9k$ from P'. In the former case, it follows that G has radius at most $18k + 2$. Let P be the trivial path on a single vertex p. Then P is $(1, 18k + 2)$-quasi-isometric to G, and $(P, (V_p)_{p \in P})$ with $V_p := V(G)$ is an honest decomposition of G of radial width at most $18k + 2$ and radial spread 0.

So we may assume the latter case. Set $P := P'$, and apply Lemma 11.4.2 to $H := P$ and $r := 9k$. This yields an honest P-decomposition of $G' := [B_G(P, 9k)]$ of radial width at most $18k + 1$ and radial spread at most $18k + 1$. Since every vertex of G has distance at most $9k$ from P, we have $G' = G$, and hence this is the desired decomposition of G. $\square$

The remainder of this section is devoted to the proof of Theorem 11.4.1. Let us first give a brief sketch of the proof. For this, let G be a connected graph and, let P be a longest geodesic path in G. For some suitably chosen $r = r(k) < 9k$, we let $G_P := G[B_G(P, r)]$. We then analyse how the components of $G - G_P$ attach to G_P, and show that either all vertices in a component have distance at most $9k$ from P or we can use the component to find a $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph of G or a $T_{3k}K_3$ as a geodesic subgraph of G.

The analysis of the components will be done in three lemmas, Lemmas 11.4.3, 11.4.4 and 11.4.6 below. They are stated in a slightly more general form than needed for the proof of Theorem 11.4.1, which enables us to use them later also in the proofs of Theorems 31 and 32.

The first lemma, Lemma 11.4.3 shows that enlarging a c-quasi-geodesic subgraph with a shortest path to it yields a subgraph which is $(2c + 1)$-quasi-geodesic. This lets us find a 3-quasi-geodesic $T_{3k}K_{1,3}$ in G if some component of $G - G_P$ attaches 'to the middle' of G_P, that is, to some ball $B_G(p, r)$ where p lies 'in the middle' of P. The second lemma, Lemma 11.4.4, demonstrates that we can find a long geodesic cycle in G if some component of $G - G_P$ attaches to G' close to the start and the end of P, but nowhere in between (see Figure 11.3). The third lemma, Lemma 11.4.6, shows that if a component of $G - G_P$ attaches to G_P only towards one end of P, then its vertices all have distance at most $9k$ from P.

Lemma 11.4.3. *Let G be a graph, and let X be a c-quasi-geodesic subgraph of G for some $c \in \mathbb{N}$. If P is a shortest v–X path in G for some vertex $v \in G$, then $X \cup P$ is $(2c + 1)$-quasi-geodesic* *in G.*

Proof. We have to show that, for every two vertices u and w of $X \cup P$, the distance of u and w in $X \cup P$ is at most $(2c+1)$ times their distance in G. Since X is c-quasi-geodesic in G and P is geodesic in G, it is (by symmetry) enough to consider the case where u is a vertex of P and w is a vertex of X. Let x be the endvertex of P in X. Since P is a shortest v–X path in G, uPx is a shortest u–X path in G and hence $d_P(u,x) = d_G(u,x) \leqslant d_G(u,w)$ as $w \in X$. We then have $d_X(x,w) \leqslant c \cdot d_G(x,w) \leqslant c \cdot (d_G(x,u) + d_G(u,w))$, where the first inequality follows from X being c-quasi-geodesic in G while the second one applies the triangle inequality. Again using the triangle inequality, we can then combine these inequalities to

$$\begin{aligned} d_{X\cup P}(u,w) &\leqslant d_{X\cup P}(u,x) + d_{X\cup P}(x,w) = d_P(u,x) + d_X(x,w) \\ &\leqslant d_G(u,x) + c \cdot (d_G(x,u) + d_G(u,w)) \leqslant (2c+1) \cdot d_G(u,w), \end{aligned}$$

which shows the claim. □

Lemma 11.4.4. *Let $r, n, m_0, m_1 \in \mathbb{N}$ such that $m := n - m_0 - m_1 - 2r > 0$. Let G be a graph containing a geodesic path $P = p_0 \dots p_n$ of length n, and write $B_i := B_G(p_i, r)$ for every $p_i \in P$. Suppose that a component C of $G - \bigcup_{i=0}^{n} B_i$ has at least one neighbour in both $\bigcup_{i=0}^{m_0} B_i$ and $\bigcup_{i=n-m_1}^{n} B_i$, but no neighbours in $\bigcup_{i=m_0+1}^{n-m_1-1} B_i$. Then G contains a geodesic cycle of length at least $2m$.*

See Figure 11.3 for a sketch of the situation in Lemma 11.4.4.

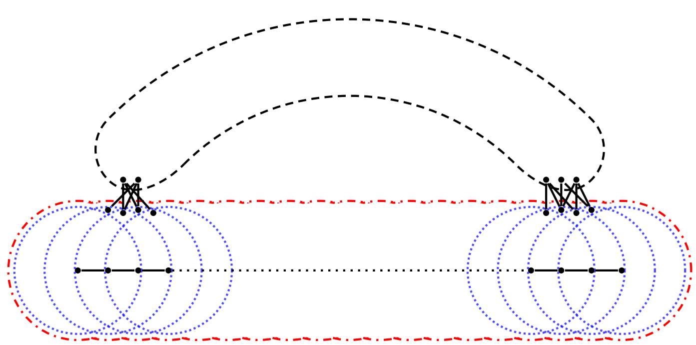

FIGURE 11.3: The setting of Lemma 11.4.4.

The proof of Lemma 11.4.4 builds on the study of how cycles interact with a given separation of the graph. More formally, we have the following preparatory lemma.

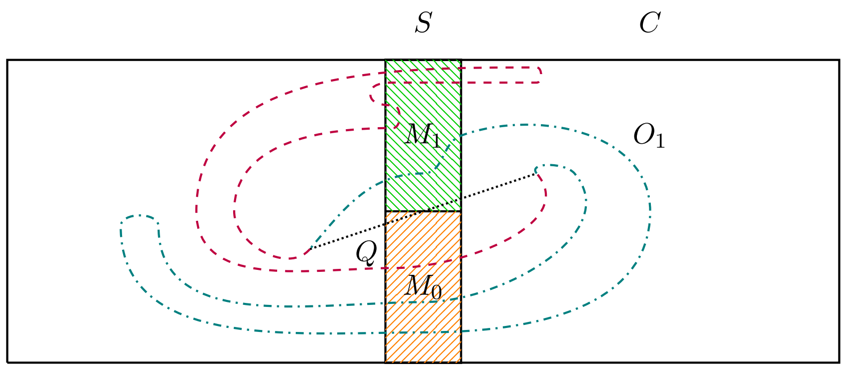

FIGURE 11.4: The setting of Lemma 11.4.5. Here, O_2 is the cycle containing Q (black, dotted line) and the purple, dashed part of O_1, and O_3 is the cycle containing Q and the teal, dash dotted part of O_1.

Lemma 11.4.5. *Let G be a graph, S a set of vertices and C a component of $G - S$, and $\{M_0, M_1\}$ be a bipartition of S. Given a cycle O in G, let $W(O)$ be the number of M_0–M_1 paths in O meeting C.*

Suppose that Q is an O_1-path of a cycle O_1 in G, and let O_2 and O_3 denote the two cycles in $O_1 \cup Q$ containing Q. Then $W(O_1) + W(O_2) + W(O_3)$ is even.

Proof. The reader may look at Figure 11.4 to follow the proof more easily. Let u and v be the endvertices of Q on O_1. Write $P_1 := Q$, and let P_2 and P_3 be the two u–v paths in O_1. Then the cycle O_1 consists of P_2 and P_3, and without loss of generality, we may assume that O_2 consists of P_1 and P_2 while O_3 consists of P_1 and P_3.

Let $\mathcal{P}$ be the set of M_0–M_1 paths in G meeting C that are paths in at least one O_i. We remark that every path in $\mathcal{P}$ has all its inner vertices in C since $S = M_0 \cup M_1$ separates C from the rest of the graph. Moreover, let us emphasise that, in the definition of $W(O)$, we do not consider the direction in which O traverses the M_0–M_1 paths it contains.

If $T \in \mathcal{P}$ is a subpath of some P_i, then T is also a subpath of exactly two of the three cycles O_1, O_2 and O_3 and thus contributes exactly 2 to the sum $W(O_1) + W(O_2) + W(O_3)$. So when we check that this sum is even, all the elements of $\mathcal{P}$ that are a subpath of some P_i contribute an even amount to the sum. In particular, if all elements of $\mathcal{P}$ are paths in some P_i, then the claim holds.

So let $\mathcal{P}' \subseteq \mathcal{P}$ consist of precisely those elements of $\mathcal{P}$ that are not a path in any P_i, and suppose that $\mathcal{P}'$ is non-empty. This implies that at least one P_i meets $M_0 \cup M_1$. Every path in $\mathcal{P}$ is a subpath of one of the cycles O_i, which in turn is disjoint from the interior of some P_i. Hence every path in $\mathcal{P}'$ meets the interior of at most, and thus precisely, 2 of the P_i. We remark that a path in an O_i that is also a path in some O_j with $i \neq j$ is already a path in some P_i. Thus, every element of $\mathcal{P}'$ contributes exactly 1 to the sum $W(O_1) + W(O_2) + W(O_3)$. In order to prove that this sum is even, it thus suffices to show that there are evenly many paths in $\mathcal{P}'$. We will check this by a case distinction. Note that a path in some O_j is also a path in some P_i if and only if it contains neither of u and v as inner vertex. Hence, all paths in $\mathcal{P}'$ contain at least one of u and v as an inner vertex. Because $\mathcal{P}'$ is non-empty and all inner vertices of its elements are contained in C, at least one of u and v is contained in C.

First, consider the case that all three paths P_1, P_2 and P_3 meet $M_0 \cup M_1$. In particular, this implies that no path in $\mathcal{P}'$ contains both u and v as an inner vertex. If all paths P_1, P_2 and P_3 have their first vertex in $M_0 \cup M_1$ contained in the same element of $\{M_0, M_1\}$, then no element of $\mathcal{P}'$ contains u as an inner vertex. Otherwise, precisely two paths P_i have their first vertex in $M_0 \cup M_1$ contained in the same element of $\{M_0, M_1\}$. In this case, there are exactly two elements of $\mathcal{P}'$ that contain u as an inner vertex. By symmetry, there are also exactly 2 or 0 elements of $\mathcal{P}'$ that contain v as an inner vertex. Hence, $\mathcal{P}'$ contains precisely $0, 2$ or 4 paths, and thus in this case $W(O_1) + W(O_2) + W(O_3)$ is even.

Secondly, consider the case that precisely one P_i, say P_1, meets $M_0 \cup M_1$. Since P_2 does not meet $M_0 \cup M_1$ and at least one of u and v is contained in C, this implies that both u and v are contained in C. Because $\mathcal{P}'$ is non-empty, the first and last vertex of P_1 in $M_0 \cup M_1$ are contained in distinct elements of $\{M_0, M_1\}$. Thus, $\mathcal{P}'$ contains precisely two paths, one containing P_2 and one containing P_3. Hence, $W(O_1) + W(O_2) + W(O_3)$ is even.

Lastly, consider the case that precisely two P_i meet $M_0 \cup M_1$. Say P_3 does not meet $M_0 \cup M_1$. Thus, again both u and v are contained in C. Here, we need to distinguish some more cases.

The first case which we consider is that both P_1 and P_2 have their first and last vertex in $M_0 \cup M_1$ contained in distinct elements of $\{M_0, M_1\}$ and that the first vertex of P_1 in $M_0 \cup M_1$ and the first vertex of P_2 in $M_0 \cup M_1$ are contained in distinct elements of $\{M_0, M_1\}$. Then $\mathcal{P}'$ contains precisely four elements: two containing P_3 and two not containing any edges of P_3. So again $W(O_1) + W(O_2) + W(O_3)$ is even.

Next, we assume that both P_1 and P_2 have their first and last vertex in $M_0 \cup M_1$ contained in distinct elements of $\{M_0, M_1\}$ but that the first vertex of P_1 in $M_0 \cup M_1$ and the first vertex of P_2 in $M_0 \cup M_1$ are contained in the same element of $\{M_0, M_1\}$. Then $\mathcal{P}'$ contains precisely two elements: both contain P_3. So again $W(O_1) + W(O_2) + W(O_3)$ is even.

Now we assume that at least one of P_1 and P_2, say P_1, has their first and last vertex in $M_0 \cup M_1$ contained in the same element of $\{M_0, M_1\}$, say in M_0. Since $\mathcal{P}'$ is non-empty and every path in $\mathcal{P}'$ contains u or v as inner vertex, this implies that either the first or the last vertex of P_2 in $M_0 \cup M_1$ is contained in M_1. If only the first or only the last vertex of P_2 in $M_0 \cup M_1$ is contained in M_1, then $\mathcal{P}'$ contains precisely two elements: one of these contains P_3 and the other does not contain edges of P_3. If both the first and the last vertex of P_2 in $M_0 \cup M_1$ are contained in M_1, then $\mathcal{P}'$ contains also precisely two elements: neither of them contains edges of P_3. Again in both cases, $W(O_1) + W(O_2) + W(O_3)$ is even. □

Proof of Lemma 11.4.4. Note that, as $n > m_0 + m_1$, we in particular have $n - m_1 > m_0$. We set

$$M_0 := N_G(C) \cap \bigcup_{i=0}^{m_0} B_i \text{ and } M_1 := N_G(C) \cap \bigcup_{i=n-m_1}^{n} B_i.$$

As C has no neighbours in $\bigcup_{i=m_0+1}^{n-m_1-1} B_i$, the neighbourhood of C is equal to $M_0 \cup M_1$ and C is a component of $G - M_0 - M_1$.

First, we note that M_0 and M_1 are disjoint. Indeed, if $u \in B_{j_0}$ for some $j_0 \leqslant m_0$ and $v \in B_{j_1}$ for some $j_1 \geqslant n - m_1$, then using the triangle inequality we obtain

$$n-m_0-m_1 \leqslant d_P(p_{j_0}, p_{j_1}) = d_G(p_{j_0}, p_{j_1}) \leqslant d_G(p_{j_0}, u)+d_G(u, v)+d_G(v, p_{j_1}) = d_G(u, v)+2r. \tag{11.1}$$

So $d_G(u, v) \geqslant n - m_0 - m_1 - 2r = m > 0$ by assumption. In particular, u and v are distinct vertices, and thus M_0 and M_1 are disjoint.

We now define, for a cycle O in G, the number $W(O)$ as the number of M_0–M_1 paths in O which have an inner vertex in C.

The remainder of the proof now consists of three steps. First, we show that a cycle O with $W(O) \neq 0$ has length at least $2m$. Second, we find a cycle O with odd $W(O)$, and lastly we show that a shortest such cycle is geodesic in G.

Claim 1. *Every cycle O in G with $W(O) \neq 0$ has length at least $2m$.*

Proof. Since $W(O) \neq 0$, the cycle O contains at least one M_0–M_1 path. As the number of M_0–M_1 paths is even for every cycle in G, the cycle O contains at least two M_0–M_1 paths, which then have to be internally disjoint. By (11.1), every such path has length at least m, and thus O has length at least $2m$. ■

Claim 2. *There exists a cycle O in G with odd $W(O)$.*

Proof. Among all M_0–M_1 paths that contain an inner vertex in C, let Q be a shortest such path, and denote by v_0 and v_1 its end vertices in M_0 and M_1, respectively. Let R_0 be a shortest v_0–P path in G and R_1 a shortest v_1–P path. Since $v_i \in M_0 \cup M_1 = N_G(C)$, both R_i have length exactly r. If R_0 and R_1 share a vertex x, say such that R_0x is at least as long as R_1x, then R_1xR_0 contains a path of length at most r from v_1 to some p_s with $s \leqslant m_0$, contradicting $v_1 \in M_1$ and $M_0 \cap M_1 = \emptyset$. Hence, R_0 and R_1 are disjoint. Now the concatenation R_0QR_1 is a path that starts and ends in distinct vertices of P and is internally disjoint from P. Hence, it can be closed by a (unique) subpath of P to a cycle O. By construction of O, there is precisely one M_0–M_1 path in O that has an inner vertex in C, and that is Q. In particular, $W(O)$ is odd. ■

Claim 3. *Let O be a cycle in G with odd $W(O)$. If O is not geodesic, then there is a cycle O' that is shorter than O such that $W(O')$ is odd.*

Proof. Since $O_1 := O$ is not geodesic, there is an O_1-path Q in G with endvertices u and v in O which is shorter than the distance of u and v in O_1. Let O_2 and O_3 be the two cycles in $O_1 \cup Q$ which contain Q. Since both O_2 and O_3 are shorter than O_1, it suffices to show that at least one of $W(O_2)$ and $W(O_3)$ is odd. Now $W(O_1)$ is odd by assumption, so it is enough to show that the sum $W(O_1) + W(O_2) + W(O_3)$ is even. This however follows from Lemma 11.4.5 applied to the component C together with the bipartition $\{M_0, M_1\}$ of $S = N_G(C)$. ■

To formally complete the proof, pick a cycle O in G with $W(O)$ odd such that O is as short as possible among these. Such O exists by Claim 2 and is geodesic in G by Claim 3. Moreover, O has length at least $2m$ by Claim 1, as desired. □

Lemma 11.4.6. *Let U be a set of vertices of a connected graph G, and let $p_0 \in V(G) \setminus U$ have maximal distance from U in G. Let $P = p_0 \dots p_n$ be a shortest p_0–U path in G. Given $r \in \mathbb{N}$, we write $B_i := B_G(p_i, r)$ for every $p_i \in P$. Suppose that some component C of $G - \bigcup_{i=0}^{n} B_i$ is disjoint from U and satisfies $N_G(C) \subseteq \bigcup_{i=0}^{\ell} B_i$ for some $\ell \leqslant n$. Then $d_G(v, P) \leqslant 2r + \ell$ for every $v \in C$.*

Proof. Let v be an arbitrary vertex of C, and let Q be a shortest v–U path in G with endvertex $u \in U$. As $V(C) \cap U = \emptyset$, the path Q intersects $\bigcup_{i=0}^{n} B_i$. Let q be the first vertex on Q in $\bigcup_{i=0}^{n} B_i$, and let j be the smallest index of a ball B_j with $q \in B_j$. Note that $j \leqslant \ell$ since $N_G(C) \subseteq \bigcup_{i=0}^{\ell} B_i$.

Since P is a longest geodesic path in G which ends in U, we have that $n = ||P|| \geqslant ||Q|| = d_G(v, u) = d_G(v, q) + d_G(q, u)$, where the last equality follows from Q being geodesic in G. Additionally, since $p_n \in U$, we have $d_G(p_j, p_n) \leqslant d_G(p_j, u) \leqslant d_G(p_j, q) + d_G(q, u) = r + d_G(q, u)$. These two inequalities combine to

$$d_G(v, q) \leqslant n - d_G(q, u) \leqslant n - (d_G(p_j, p_n) - r).$$

All in all, we get

$$d_G(v, P) \leqslant d_G(v, p_j) \leqslant d_G(v, q) + r \leqslant n - (d_G(p_j, p_n) - r) + r = n - (n - j) + 2r \leqslant 2r + \ell,$$

where the last inequality follows from $j \leqslant \ell$. This concludes the proof of the claim. □

With all the previous lemmas at hand, we are now ready to prove Theorem 11.4.1.

Proof of Theorem 11.4.1. Let G be a connected graph that contains neither $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph nor $T_{4k+1}K_3$ as a geodesic subgraph.

Let $P = p_0 \dots p_n$ be a longest geodesic path in G. We may assume $n \geqslant 18k + 3$; otherwise we are done. For every $0 \leqslant i \leqslant n$, define $B_i := B_G(p_i, 3k)$ as the ball in G of radius $3k$ around p_i.

If $G = G_P$, then we are done. Otherwise, we consider the components of $G - G_P$. Since G is connected by assumption, each component C of $G - G_P$ has a neighbour in at least one B_i. The following two claims show that in fact $N_G(C) \subseteq \bigcup_{i=0}^{3k} B_i$ or $N_G(C) \subseteq \bigcup_{i=n-3k}^{n} B_i$.

Claim 1. *No component C of $G - G_P$ has a neighbour in B_i with $3k < i < n - 3k$.*

Proof. Suppose for a contradiction that there is such a component C. We show that G then contains a $T_{3k}K_{1,3}$ as 3-quasi-geodesic subgraph, which contradicts our assumption on G.

Let $v \in C$ be a vertex with a neighbour in B_i with $3k < i < n - 3k$, and let Q be a shortest v–p_i path in G. The choice of v guarantees that the path Q has length exactly $3k + 1$, and thus Q is a shortest v–P path in G as every vertex in $G - G_P$ and hence in C has distance at least $3k + 1$ to P. By Lemma 11.4.3, $P \cup Q$ is a 3-quasi-geodesic subgraph of G. It follows from $3k < i < n - 3k$ that $P \cup Q$ is a $T_{3k}K_{1,3}$, as desired. ■

Claim 2. *No component C of $G - G_P$ has at least one neighbour in $\bigcup_{i=0}^{3k} B_i$, at least one neighbour in $\bigcup_{i=n-3k}^{n} B_i$ and no neighbours in any B_i with $3k < i < n - 3k$.*

Proof. Suppose towards a contradiction that there is such a component C. Applying Lemma 11.4.4 to C with $m_0 = m_1 = r = 3k$, we find that G contains a geodesic cycle O of length at least $2(n - 12k)$. But $n \geqslant 18k + 3$ as we assumed the graph to have a large radius, so O has length at least $2(18k + 3 - 12k) = 12k + 6$. Thus, O is a geodesic cycle in G of length at least $12k + 6$, a contradiction to our assumptions on G. ■

Thus, each component of $G - G_P$ attaches either only to balls B_i with $i \leqslant 3k$ or to balls B_i with $i \geqslant n - 3k$. But the vertices in such components have bounded distance from P in G, as the following claim shows.

Claim 3. *Let C be a component of $G - G_P$ such that $N_G(C) \subseteq \bigcup_{i=0}^{3k} B_i$ or $N_G(C) \subseteq \bigcup_{i=n-3k}^{n} B_i$. Then every vertex $v \in C$ has distance at most $9k$ from P in G.*

Proof. If $N_G(C) \subseteq \bigcup_{i=0}^{3k} B_i$, then the claim follows by applying Lemma 11.4.6 with $U = \{p_n\}$. Otherwise, $N_G(C) \subseteq \bigcup_{i=n-3k}^{n} B_i$, and the claim follows by applying Lemma 11.4.6 with $U = \{p_0\}$. ■

By Claims 1 to 3, every vertex of G has distance at most $9k$ from P, as desired. □

11.5 Radial cycle-width

In this section we build on our result on radial path-width, Theorem 30, and use a small adaptation of its proof to address radial cycle-width by proving Theorem 31. More precisely, Theorem 31 follows from the lemmas that we have already shown in Section 11.4. Let us restate Theorem 31 here for convenience.

Theorem 31. *Let $k \in \mathbb{N}$. If a connected graph G contains no $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph, then G admits an honest decomposition modelled on a path or cycle C of radial width at most $18k+2$ and radial spread at most $36k+2$.*

Moreover, C is $(1, 18k+2)$-quasi-isometric to G.

We remark that we did not optimise the bound on the radial width and radial spread.

In fact, we show the following stronger statement, which immediately implies Theorem 31.

Theorem 11.5.1. *Let $k \in \mathbb{N}$, and let G be a connected graph. If G contains no $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph, then there exists a geodesic cycle or path C in G such that C is either a path of length at most $18k+2$ or every vertex of G has distance at most $9k$ from C.*

Let us first show that Theorem 11.5.1 implies Theorem 31.

Proof of Theorem 31 given Theorem 11.5.1. Let G be a connected graph that contains no $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph, and let C' be given by applying Theorem 11.5.1 to G. If C' is a path of length at most $18k+2$, then G has radius at most $18k+2$. Let C be the trivial path on a single vertex c. Then C is $(1, 18k+2)$-quasi-isometric to G, and $(C, (V_c))$ with $V_c := V(G)$ is an honest decomposition of G of radial width at most $18k+2$ and radial spread 0.

So we may assume that every vertex of G has distance at most $9k$ from C. Set $C := C'$, and apply Lemma 11.4.2 to $H := C$ and $r := 9k$. This yields an honest C-decomposition of $G' := [B_G(C, 9k)]$ of radial width at most $18k+1$ and radial spread at most $36k+2$. Since every vertex of G has distance at most $9k$ from C, we have $G' = G$, and hence this is the desired decomposition of G. □

Let us now prove Theorem 11.5.1.

Proof of Theorem 11.5.1. Let G be a connected graph that contains no $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph. Applying Theorem 11.4.1 to G yields that G either contains a geodesic path P such that $C := P$ is as desired, or G contains a $T_{4k+1}K_3$ as a geodesic subgraph. Since we are done in the first case, we may thus assume that G contains a $T_{4k+1}K_3$ as a geodesic subgraph, which means that there exists a geodesic cycle C in G of length at least $3 \cdot (4k+2) = 12k+6$. Set $G_C := G[B_G(C, 3k)]$. If $G = G_C$, then we are done.

Otherwise, there exists a neighbour v of $V(G_C)$ in G since G is connected. Let P be a shortest v–C path in G, and let u be its endvertex in C. Since $v \in N_G(G_C)$, it has distance $3k+1$ from C, and thus P has length $3k+1$. Let $Q := B_C(u, 3k+1)$ be the subpath of C of length $6k+2$ which contains u as its 'middle' vertex. Note that Q is actually a path as C has length at least $12k+6 > 6k+2$. Since P is a path from v to C which ends in u, the paths P and Q only meet in u. Hence, $P \cup Q$ is a $T_{3k}K_{1,3}$ in G.

To conclude the proof and obtain the desired contradiction to our assumptions on G, it remains to show that $P \cup Q$ is 3-quasi-geodesic in G. Since Q is a subpath of C with $||Q|| = 6k+2 \leqslant \frac{1}{2}(12k+6) \leqslant \frac{1}{2}||C||$, we have $d_Q(u,v) = d_C(u,v)$ for every two vertices $u, v \in Q$. Thus, Q is geodesic in G, since C is geodesic in G. By Lemma 11.4.3, $P \cup Q$ then is a 3-quasi-geodesic subgraph of G, as desired. □

11.6 Radial star-width

In this section we prove Theorem 32, which we restate here for convenience.

Theorem 32. *Let $k \in \mathbb{N}$. If a connected graph G contains no T_kK_3 as a geodesic subgraph and no $T_{3k}W$ as a 3-quasi-geodesic subgraph, then G admits an honest decomposition modelled on a subdivided star of radial width at most $58k+9$ and radial spread at most $58k+9$.*

Moreover, there exists $C_k \in \mathbb{N}$ such that some subdivided star is $(1, C_k)$-quasi-isometric to G.

Recall that W is the wrench graph depicted in Figure 11.1. As we already mentioned in the introduction, one can check by carefully reading the proof that the subdivided star which we construct for the first part of the statement already satisfies the 'moreover'-part, and that we may choose $C_k = 60k+14$ (see the paragraph after the proof of Theorem 32 for details).

Before we start with the proof of Theorem 32, we first show an auxiliary lemma. Recall that in the proof of Theorem 30 we used Lemma 11.4.4 as a tool to show that a graph contains a geodesic T_kK_3. Similarly, we will use the following lemma in the proof of Theorem 32 to show that a graph contains a 3-quasi-geodesic $T_{3k}W$.

Lemma 11.6.1. *Let $P = p_0 \dots p_n$ be a geodesic path in a graph G. Given two vertices $u, v \in G$, let Q be a shortest u–P path in G with endvertex $q \in P$, and let R be a shortest v–$(P \cup Q)$ path in G with endvertex $r \in P \cup Q$. Suppose that at least one of the following two conditions hold:*

(i) *$r \in P$ and $d_G(q,r) \geqslant 4\max\{d_G(u,P), d_G(v,P)\}$;*

(ii) *$r \in Q$ and $d_G(q,r) \geqslant 4\max\{d_G(v,Q), d_G(p_0,q), d_G(p_n,q)\}$.*

Then $P \cup Q \cup R$ is a 3-quasi-geodesic subgraph of G.

Proof. We first consider the case that $r \in P$. By Lemma 11.4.3, we have that $P \cup Q$ and $P \cup R$ are each 3-quasi-geodesic subgraphs of G. Hence, in order to show that $X := P \cup Q \cup R$ is 3-quasi-geodesic, it suffices to consider vertices $x, y \in X$ with $x \in Q$ and $y \in R$.

Since P is geodesic, we have that $d_X(q,r) = d_G(q,r) \leqslant d_G(q,x) + d_G(x,y) + d_G(y,r)$ and thus

$$\begin{aligned} d_G(x,y) &\geqslant d_G(q,r) - d_G(x,q) - d_G(y,r) \geqslant 4\max\{d_G(u,P), d_G(v,P)\} - d_G(u,P) - d_G(v,P) \\ &\geqslant 2\max\{d_G(u,P), d_G(v,P)\} \geqslant 2\max\{d_G(x,q), d_G(y,r)\}. \end{aligned}$$

Therefore, making again use of $d_G(q,r) \leqslant d_G(q,x) + d_G(x,y) + d_G(y,r)$, we obtain

$$\begin{aligned} d_X(x,y) &= d_X(x,q) + d_X(q,r) + d_X(r,y) = d_G(x,q) + d_G(q,r) + d_G(r,y) \\ &\leqslant 2d_G(x,q) + 2d_G(y,r) + d_G(x,y) \leqslant 3d_G(x,y). \end{aligned}$$

For the second case, assume that $r \in Q$. Similar as before, we find that $P \cup Q$ and $Q \cup R$ are each 3-quasi-geodesic subgraphs of G. Hence, we only need to consider vertices $x, y \in X$ with $x \in P$ and $y \in R$. As before, we find that

$$d_G(x,y) \geqslant d_G(q,r) - d_G(x,q) - d_G(y,r) \geqslant 2\max\{d_G(v,Q), d_G(p_0,q), d_G(p_n,q)\},$$

and thus

$$d_X(x,y) = d_G(x,q) + d_G(q,r) + d_G(r,y) \leqslant 2d_G(x,q) + 2d_G(y,r) + d_G(x,y) \leqslant 3d_G(x,y). \qquad \square$$

Now we prove Theorem 32. We remark that we did not optimise the bounds on the radial width and radial spread.

Proof of Theorem 32. Let G be a graph that contains neither T_kK_3 as a geodesic subgraph nor $T_{3k}W$ as a 3-quasi-geodesic subgraph. We will construct a subdivided star S and an S-decomposition of G of radial width at most $58k+9$ and radial spread at most $58k+9$. The 'moreover' part of the statement then follows: By Lemmas 9.2.1 and 11.2.1, there exists an (L,C)-quasi-isometry from G to S such that L and C depend on k only. Applying [102, 1.6] (to $L, C, 2$)[5] then yields a constant C'_k such that G is $(1, C'_k)$-quasi-isometric to some subdivided star S'. By Lemma 9.2.1, it then follows that S' is $(1, C_k)$-quasi-isometric to G for $C_k := 3C'_k$.

Let $P = p_0 \dots p_n$ be a longest geodesic path in G. Observe that we may assume $n \geqslant 58k+10$; otherwise, G has a trivial $\mathcal{H}$-decomposition into a single ball of radius at most $58k+9$. For every

[5]Note that we use here that subdivided stars admit path-decompositions into bags of size 3 and thus have path-width 2.

$0 \leqslant i \leqslant n$ define B_i as the ball in G of radius $3k$ around p_i, and let $G_P := G[\bigcup_{i=0}^{n} B_i]$.

In the following claims we will analyse how the components of $G - G_P$ attach to G_P. Note that, as we assumed G to be connected, every component of $G - G_P$ has some neighbour in some B_i. First we show that we may assume that some component of $G - G_P$ attaches to G_P somewhere in the middle of P.

Claim 1. *Either G has an honest P-decomposition of radial width at most $58k+9$ and radial spread at most $58k+9$ or some vertex of $G-G_P$ has a neighbour in some B_s with $23k+5 \leqslant s \leqslant n-23k-5$.*

Proof. We assume for a contradiction that whenever a component of $G - G_P$ has a neighbour in some B_s then $s \leqslant 23k+4$ or $n - 23k - 4 \leqslant s$.

First we consider the case that some component C of $G - G_P$ has at least one neighbour in both $\bigcup_{i=0}^{23k+4} B_i$ and $\bigcup_{i=n-23k-4}^{n} B_i$. Let $m_0 = m_1 = 23k+4$ and $r = 3k$. As

$$n - m_0 - m_1 - 2r = n - (23k+4) - (23k+4) - 6k = n - 52k - 8 \geqslant 2k + 2 \geqslant 2,$$

we can apply Lemma 11.4.4 to obtain a geodesic cycle in G of length at least $2 \cdot (2k+2) > 3k+3$ which contradicts our assumption that G does not contain a geodesic $T_k K_3$.

Thus, for every component C of $G - G_P$, either $N_G(C) \subseteq \bigcup_{i=0}^{23k+4} B_i$ or $N_G(C) \subseteq \bigcup_{i=n-23k-4}^{n} B_i$ (this includes in particular the case that $G = G_P$). Then by Lemma 11.4.6 applied to $U = \{p_0\}$ or $U = \{p_n\}$, every vertex in $G - G_P$ has distance at most $2 \cdot 3k + 23k + 4 = 29k + 4$ to P. Hence, by defining $B_i' = B_G(p_i, 29k+4)$, every vertex of G is contained in some B_i'. So by Lemma 11.4.2 applied to the geodesic path P and to $r = 29k+4$ we obtain an honest P-decomposition of G of radial width at most $2 \cdot (29k+4)+1 = 58k+9$ and radial spread at most $2 \cdot (29k+4)+1 = 58k+9$. ■

By Claim 1 (and since P is a path and hence a subdivided star) we may assume that there is some s with $23k+5 \leqslant s \leqslant n-23k-5$ such that some vertex of $G-G_P$ has a neighbour in B_s. We fix s for the rest of the proof. The following two claims now show that every component of $G - G_P$ can only attach to G_P either close to the start of P or close to the end of P or close to p_s.

Claim 2. *If a vertex of $G - G_P$ has a neighbour in a ball B_t, then $t \leqslant 3k$ or $s - 12k - 3 \leqslant t \leqslant s + 12k + 3$ or $n - 3k \leqslant t$.*

Proof. We assume for a contradiction that there is some vertex v of $G-G_P$ that has a neighbour in some B_t with $3k < t < s - 12k - 3$, the other case is symmetric. Let $R = r_0 \dots r_{3k+1}$ be a shortest path from $v = r_0$ to $p_t = r_{3k+1}$ in G. Similarly, let u be a vertex of $G - G_P$ that has a neighbour in B_s and let $Q = q_0 q_1 \dots q_{3k+1}$ be a shortest path from $u = q_0$ to $p_s = q_{3k+1}$ in G. Then $p_t P p_s$ has length at least $12k+4 = 4 \cdot (3k+1)$, so by applying Lemma 11.6.1 to $p_{t-3k-1} P p_{s+3k+1}$, Q and R we obtain a 3-quasi-geodesic $T_{3k} W$ in G, a contradiction to our assumption on G. ■

Claim 3. *For every component C of $G - G_P$, its neighbourhood $N_G(C)$ is contained either in $\bigcup_{i=0}^{3k} B_i$, in $\bigcup_{i=s-12k-3}^{s+12k+3} B_i =: D_s$, or in $\bigcup_{i=n-3k}^{n} B_i$.*

Proof. First, we assume for a contradiction that C has neighbours in both $\bigcup_{i=0}^{3k} B_i$ and D_s. Let $r = 3k$, $m_0 = 3k$ and $m_1 = n - (s - 12k - 3)$. Then, since $s \geqslant 23k + 5$,

$$n - m_0 - m_1 - 2r = n - 3k - (n - (s - 12k - 3)) - 6k = s - 21k - 3 \geqslant 2k + 2 \geqslant 2,$$

so we can apply Lemma 11.4.4 to obtain a geodesic cycle in G of length at least $2 \cdot (2k+2) > 3k+3$. This contradicts our assumption that G does not contain a geodesic $T_k K_3$.

Hence, if C has neighbours in D_s, then it does not have neighbours in $\bigcup_{i=0}^{3k} B_i$, and by symmetry it neither has neighbours in $\bigcup_{i=n-3k}^{n} B_i$. Furthermore, if C has neighbours in both $\bigcup_{i=0}^{3k} B_i$ and $\bigcup_{i=n-3k}^{n} B_i$, then it follows again from Lemma 11.4.4 that G contains a geodesic $T_k K_3$, which yields the same contradiction. ■

Let us first consider all the components of $G - G_P$ that attach to G_P either close to the start or close to the end of P. By Lemma 11.4.6 these components can only contain vertices which are close to P. This fact allows us to find a path-decomposition of $G' := G[V']$ of low radial width and spread where V' is the union of the B_i and the vertices of components C of $G - G_P$ with $N_G(C) \subseteq \bigcup_{i=0}^{3k} B_i$ or $N_G(C) \subseteq \bigcup_{i=n-3k}^{n} B_i$, as follows.

Claim 4. *All vertices of G' have distance at most $9k$ from P. Also, there is an honest decomposition $(P, \mathcal{V}')$ of G' of radial width at most $24k+5$ and radial spread at most $48k+8$ such that*

- *all vertices in the bag V'_s of the node p_s of P have distance at most $24k+5$ from p_s in $G[V'_s]$, and*
- *for every $v \in N_G(G')$ the set $N_G(v) \cap V'$ is contained in the bag V'_s.*

Proof. First, we show that vertices of $G' - G_P$ are close to p_0 or p_n. Let C be a component of $G' - G_P$ with $N_G(C) \subseteq \bigcup_{i=0}^{3k} B_i$. By Lemma 11.4.6 applied to $U = \{p_n\}$ and $\ell = 3k$, every vertex $v \in C$ has distance at most $2 \cdot (3k) + 3k = 9k$ to P, and thus has distance at most $12k$ to p_0. By symmetry, every component C of $G' - G_P$ with $N_G(C) \subseteq \bigcup_{i=n-3k}^{n} B_i$ only contains vertices whose distance from p_n is at most $12k$.

Now we construct the P-decomposition. For $i \in [n]$ define

$$V'_i := \bigcup_{p_j \in B_P(p_i, 12k+3)} B_{G'}(p_j, 12k+2).$$

By Lemma 11.4.2, the V'_i are the bags of a P-decomposition of G' of radial spread at most $4 \cdot (12k+2) + 2 = 48k+8$, and every element of V'_i has distance at most $24k+5$ from p_i in $G[V'_i]$

by construction. Furthermore, every component of $G' - G_P$ is contained in $B_{G'}(p_0, 12k+2) \subseteq V_0'$ or in $B_{G'}(p_n, 12k+2) \subseteq V_n'$.

Also, for every B_i that contains a neighbour of $G - G' \subseteq G - G_P$ we have $s - 12k - 3 \leqslant i \leqslant s + 12k + 3$ by Claim 2 and thus $B_i \subseteq B_{G'}(p_i, 12k+2) \subseteq V_s'$. As V_s' is the bag corresponding to p_s, this completes the proof. ■

We now construct path-decompositions of low radial width and spread of the remaining components, that is of the components of $G - V'$. We then combine these path-decompositions with the path-decomposition $(P, \mathcal{V}')$ of G' to a star-decomposition of G, whose central node will be p_s. For this, we need to enlarge the bag V_s' assigned to p_s a little. Indeed, at the moment, the components of $G - G'$ need not have low radial path-width; in fact, they can still be star-like. But if we delete larger balls around the nodes in P that are close to p_s, we indeed end up with components that are path-like. For this, we define V'' to be the set of vertices of G that have distance at most $18k+4$ from P, and let $G'' := G[V'']$ (so $G' \subseteq G''$). In particular, by Claim 2, every component of $G - G''$ 'attaches in the middle of P', i.e. every neighbour of G'' in a component of $G - G''$ has distance $18k+5$ to a p_i with $s - 12k - 3 \leqslant i \leqslant s + 12k + 3$. Let further $\mathcal{C}$ be the set of components that contain a vertex whose distance from P is at least $24k+5$, and let $\mathcal{C}'$ be the set of all other components of $G - V''$. We now add all vertices in $V'' \setminus V'$ and all components from $\mathcal{C}'$ to the bag V_s' that is indexed by p_s. More precisely, we define

$$V_s := V_s' \cup (V'' \setminus V') \cup \{V(C) \colon C \in \mathcal{C}'\}.$$

Further, we set $V_i := V_i'$ for all $i \neq s \in [n]$. Let also $V''' := V'' \cup V_s$ and $G''' := G[V''']$.

Claim 5. *$(P, \mathcal{V})$ is a decomposition of G''' of radial width at most $36k+7$ and radial spread at most $48k+8$ and every vertex in V_s has distance at most $36k+7$ from p_s in $G[V_s]$.*

Proof. Since $N_G(v) \cap V'$ is contained in $V_s' \subseteq V_s$ for every vertex $v \in N_G(G')$ by Claim 4, and because $(P, \mathcal{V}')$ is a decomposition of G', it follows that $(P, \mathcal{V})$ is a decomposition of G'''. Its radial spread is at most $48k+8$ by Claim 4 and because each vertex in $G''' - V'$ is only contained in V_s. Moreover, every part $G[V_i]$ with $i \neq s$ has radius at most $24k+5$ by Claim 4, so it remains to consider $G[V_s]$.

Let $v \in V_s$ be given. If $v \in V_s'$, then $d_{G[V_s]}(v, p_s) \leqslant 24k+5$ by Claim 4. Otherwise, v has distance at most $24k+4$ from some vertex p_i of P, where $s - 12k - 3 \leqslant i \leqslant s + 12k + 3$ by Claim 2, and hence

$$d_{G[V_s]}(v, p_s) \leqslant d_{G[V_s]}(v, p_i) + d_{G[V_s]}(p_i, p_s) \leqslant (24k+4) + (12k+3) = 36k+7$$

where we used that $p_i P p_s \subseteq G[V_s'] \subseteq G[V_s]$ by the definition of V_s'. ■

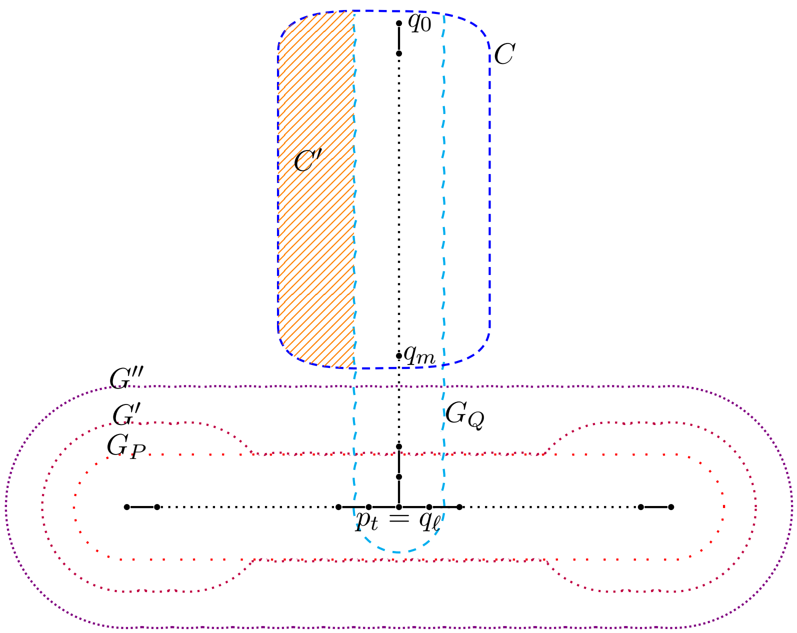

FIGURE 11.5: The setting of the proof of Claim 6.

We now show that all components of $G - V'''$, i.e. those in $\mathcal{C}$, are path-like.

Claim 6. *Let C be a component in $\mathcal{C}$. Then there is an honest decomposition of $G[V(C) \cup V_s]$ modelled on a path Q of radial width at most $54k + 8$ and radial spread at most $12k + 2$ such that the bag corresponding to the last node of Q contains V_s, and V_s only meets bags assigned to nodes of distance at most $6k$ to the last node of Q.*

Proof. The reader may look at Figure 11.5 to follow the proof more easily. Let q_0 be a vertex of C of maximal distance from P. Since $C \in \mathcal{C}$ the vertex q_0 has distance at least $24k + 5$ from P. Let $Q = q_0 \dots q_\ell$ be a shortest path in G from q_0 to P (so Q is geodesic in G). Note that C is contained in a component of $G - G'$, and thus Q ends in a vertex $p_t = q_\ell$ with $s - 12k - 3 \leqslant t \leqslant s + 12k + 3$ by Claim 2. In particular $3k < t < n - 3k$. Let m be the last index such that $q_m \notin G''$. Then Qq_m is a shortest q_0–$N_G(G'')$-path, and $m = \ell - 18k - 4$.

For every $0 \leqslant i \leqslant \ell$ we define B_i^Q to be the ball in G of radius $3k$ around q_i, and let $G_Q = G[\bigcup_{i=0}^{\ell} B_i^Q]$. We claim that every vertex of $C - G_Q$ has distance at most $12k$ from q_0. To see this, let C' be any component of $C - G_Q$, and let C'' be the component of $G - G_Q$ containing C'. We first show that C'' has no neighbour in $\bigcup_{i=3k+1}^{\ell} B_i^Q$ and that this in particular implies that $C' = C''$.

Towards a contradiction, we first assume that C'' has a neighbour in some B_j^Q with $3k < j < \ell - 12k - 3$. Then q_jQ has length at least $12k + 4 = 4 \cdot (3k + 1)$. So Lemma 11.6.1 (ii) applied to $p_{t-3k-1}Pp_{t+3k+1}$, $q_{j-3k-1}Q$ and a shortest path from C'' to q_j yields that G contains a 3-quasi-geodesic $T_{3k}W$, a contradiction.

Second, suppose for a contradiction that C'' has a neighbour in $\bigcup_{i=\ell-12k-3}^{\ell} B_i^Q$. In addition, C'' has a neighbour in $\bigcup_{i=0}^{3k} B_i^Q$. Indeed, since vertices in $\bigcup_{i=\ell-12k-3}^{\ell} B_i^Q$ have distance at most $15k+3 < 18k+4$ from P, they are neither contained in C' nor do they have neighbours in C'. Since $C' \subseteq C''$, this implies that C' has no neighbours in $\bigcup_{i=3k+1}^{\ell} B_i^Q$. As C is connected and contains $q_0 \in G_Q$, C' has a neighbour in G_Q which then has to be contained in $\bigcup_{i=0}^{3k} B_i^Q$. Hence, C'' also has a neighbour in $\bigcup_{i=0}^{3k} B_i^Q$.

For $r = 3k$, $m_0 = 3k$ and $m_1 = 12k+3$, we have that

$$\ell - m_0 - m_1 - 2r = \ell - 3k - 12k - 3 - 6k = \ell - 21k - 3 \geqslant 2k+2,$$

so we can apply Lemma 11.4.4 to Q and C'' to obtain a geodesic cycle of length at least $2(2k+2) > 3k+3$, which contradicts our assumption on G. Thus, $N_G(C'') \subseteq \bigcup_{i=0}^{3k} B_i^Q$. As G'' is connected and contains B_ℓ^Q, it follows that $C'' \cap G'' = \emptyset$ and thus $C' = C''$, if G'' and $\bigcup_{i=0}^{3k} B_i^Q$ do not meet.

Indeed, suppose for a contradiction that G'' and $\bigcup_{i=0}^{3k} B_i^Q$ meet in a vertex v. Then

$$d_G(P, q_0) \leqslant d_G(P, v) + d_G(v, q_i) + d_G(q_i, q_0) \leqslant 18k + 4 + 3k + 3k = 24k + 4,$$

which contradicts the choice of q_0.

We can now show that all vertices in C' have distance at most $12k$ from q_0. For this, consider the induced subgraph of G on the vertex set $V(G'') \cup V(C)$. In this graph, q_0 has maximal distance from P and thus we can apply Lemma 11.4.6 to the shortest $q_0 - P$ path Q, $r = 3k$, $\ell = 3k$ and the component C'. So every vertex in C' has distance at most $9k$ from Q. As every shortest C'–Q path ends in some q_i with $i \leqslant 3k$, this implies that every vertex in C' has distance at most $12k$ from q_0.

Now we obtain a path-decomposition of $G[V(C) \cup V_s]$ as follows. In a first step, for $i \leqslant \ell$ we define

$$U_i' := \bigcup_{q_j \in B_Q(q_i, 3k+1)} B_j^Q.$$

By Lemma 11.4.2, the U_i' are the bags of a Q-decomposition of G_Q of radial spread at most $4 \cdot (3k) + 2 = 12k + 2$; and every vertex in any U_i' has distance at most $6k+1$ from q_i in $G[U_i']$.

In a second step, for $Q^* = Qq_m$, we define a Q^*-decomposition of $G_Q^* \cup C$ where $G_Q^* := G[\bigcup_{i=0}^{m+3k+1} B_i^Q]$. For that, we let $U_i'' := U_i'$ for $1 \leqslant i \leqslant m$. The neighbourhood of every component of $C - G_Q$ is contained in U_0', so we can just add its vertices to U_0': let U_0'' be the union of U_0' and all vertices of components of $C - G_Q$. By construction U_0'' has radius at most $12k$. Then the U_i'' are the bags of a Q^*-decomposition of $G_Q^* \cup C$. Indeed, every vertex of $C - G_Q$ is contained in U_0''

by definition, and every vertex of $C \cap G_Q$ is contained in G_Q^* since all vertices in $\bigcup_{i=m+3k+2}^{\ell} B_i^Q$ have distance at most $18k+3$ from P and are thus not contained in C. Moreover, $(Q^*,\mathcal{U}'')$ has radial width at most $12k$ and radial spread at most $12k+2$.

We now restrict the bags U_i'' to C and add V_s to all U_i'' with $i \geqslant m-6k$, i.e. we set $U_i := (U_i'' \cap V(C))$ for $i < m-6k$ and $U_i := (U_i'' \cap V(C)) \cup V_s$ for all $i \geqslant m-6k$. To verify that $(Q^*,\mathcal{U})$ is a decomposition of $G[V(C) \cup V_s]$, it suffices to show that $N_G(G-C) \subseteq U_m$.

For this, let $v \in N_G(G-C) \subseteq V(C)$ be given, and let w be a neighbour of v in $G-C$. As C is a component of $G - G''$, this means that $w \in G''$. As we have already seen above, no component of $C - G_Q$ has neighbours in G'', so $v \in G_Q$, which in particular implies that $v \in G_Q^*$. If $v \in \bigcup_{i=m-3k-1}^{m+3k+1} B_i^Q$, then $v \in U_m$ by construction. So we may assume that $v \in B_j^Q$ for some $j < m-3k-1$. But this implies that q_0 has distance at most

$$d_G(q_0,q_j)+d_G(q_j,v)+d_G(v,w)+d_G(w,P) \leqslant (m-3k-2)+3k+1+(18k+4) = m+18k+3 < \ell$$

from P, which contradicts that Q is a shortest q_0–P path.

Hence, $(Q^*,\mathcal{U})$ is a path-decomposition of $G[V(C) \cup V_s]$, and it has radial spread at most $12k+2$ by construction. Moreover, it has radial width at most $54k+8$: Let $i \leqslant m$ and let $v \in U_i$. If $i \leqslant m-6k$, then, since $Q^* \subseteq C$ is a shortest q_0–$N_G(G-C)$ path in G and $G[U_i']$ has radius at most $12k$, we have $U_i = U_i'$ and hence $G[U_i]$ has radius at most $12k$. Now suppose that $i > m-6k$. If $v \in V_s$, then v has distance at most $36k+7$ from p_s by Claim 5. Otherwise, if $v \in U_i \setminus V_s$ then

$$d_{G[U_i]}(v,p_s) < d_{G[U_i]}(v,q_i) + d_{G[U_i]}(q_i,q_m) + d_{G[U_i]}(q_m,p_s) \leqslant 12k+6k+(36k+8) = 54k+8,$$

where we used that $q_iQq_m \subseteq G[U_i]$ by definition, that $q_m \in N_G(V_s') \subseteq V_s$ by Claim 4, and that every vertex in V_s has distance at most $36k+7$ from p_s by Claim 5. ■

We now combine these path-decompositions of the components of $G - V'''$ with the path-decomposition $(P,\mathcal{V})$ of G''' to a star-decomposition of G. For this, recall that by Claim 5, $(P,\mathcal{V})$ is a path-decomposition of G''' of radial width at most $36k+7$ and radial spread at most $48k+8$. Moreover, the components in $\mathcal{C}$ are precisely the components of $G - V'''$. For every component $C \in \mathcal{C}$, Claim 6 guarantees the existence of a decomposition $(Q_C,\mathcal{U}_C)$ of C modelled on a path Q_C of radial width at most $54k+8$ and radial spread at most $12k+2$ such that the bag in $\mathcal{U}_C$ corresponding to the last vertex of Q_C contains V_s.

We obtain a subdivided star S from the disjoint union of P and the Q_C by adding edges from the last vertices of the Q_C to p_s. Now every vertex h of S already has a bag, which we denote by V_h, in exactly one of the path-decompositions $(P,\mathcal{V})$ or $(Q_C,\mathcal{U}_C)$. It is straightforward to check that $(S,\mathcal{V})$ is a star-decomposition of G and that is radial width is at most $54k+8$. Moreover, its radial spread is at most $(48k + 8) + 1 + 6k = 52k + 9$, as only vertices in V_s may lie in more than one decomposition of the form $(P, \mathcal{V})$ or $(Q_C, \mathcal{V})$. This completes the proof. $\square$

We remark that the proof actually yields that the subdivided star S is $(1, 60k + 14)$-quasi-isometric to G (where we may choose $\varphi(h) = h$ for all vertices h of S as $V(S) \subseteq V(G)$ by construction). Here is a hint for the proof: The proof of Theorem 32 shows that all vertices in G have distance at most $58k + 9 \leqslant 60k + 14$ from some vertex of S. Moreover, since P and all paths Q_C are geodesic in G, it remains to check the distances of vertices h, h' of S that lie on distinct such paths. Assume $h \in V(Q_C)$ and $h' \in V(Q_{C'})$ ($h \in V(Q_C)$ and $h' \in V(P)$ is similar). Then $d_S(h, h') \leqslant d_G(h, h')$ as $Q_C \subseteq C$ is a shortest path in G between its first vertex and $N_G(G-C)$. Also, $d_G(h, h') \leqslant d_S(h, p_s) + d_G(N_G(G-C), N_G(G-C')) + d_S(p_s, h') = d_S(h, h') + 60k + 14$ since p_s is the centre of the subdivided star S and because $d_G(N_G(G-C), N_G(G-C')) \leqslant 2 \cdot (18k+4) + 2 \cdot (12k+3)$ follows from the fact that C, C' are components of $G - V'' = G - B_G(P, 18k + 4)$ and from Claim 2.

12 A characterisation of the graphs quasi-isometric to K_4-minor-free graphs

We prove that there is a function f such that every graph with no K-fat K_4 minor is $f(K)$-quasi-isometric to a graph with no K_4 minor. This solves the K_4-case of a general conjecture of Georgakopoulos and Papasoglu. Our proof technique also yields a new short proof of the respective K_4^--case, which was first established by Fujiwara and Papasoglu.

This chapter is based on [13] and joint for with Raphael W. Jacobs, Paul Knappe and Paul Wollan.

12.1 Introduction

All graphs in this chapter may be infinite, unless otherwise stated.

12.1.1 Quasi-isometry and fat minors

Gromov's [80] coarse geometry viewpoint had a profound influence on the field of geometric group theory and has resonated in neighbouring areas of study. At its core, this perspective revolves around the concept of *quasi-isometry*, a generalisation of bi-Lipschitz maps which allows for an additive error. Roughly speaking, two metric spaces are quasi-isometric whenever their large scale geometry coincides (see Section 9.2 for the definition).

Following Gromov's idea into the realm of graphs, Georgakopoulos and Papasoglu [75] recently presented results and questions regarding the interplay of geometry and graphs, which they hope evolves into a coherent *Coarse Graph Theory* or *Graph-Theoretic Geometry.* At the heart of their paper, they proposed a conjecture [75, Conjecture 1.1] which would offer, for some prescribed finite graph X, a characterisation of the graphs whose large scale geometry is that of a graph with no X minor. This characterisation is in terms of 'fat minors', a coarse variant of graph minors. Roughly speaking, a *fat minor* is a minor with additional distance constraints, which in particular ensures its branch sets to be pairwise far apart; especially, 0-fat minors are the usual minors (see Section 9.5 for the definition).

Conjecture 11.1.1 ([75])**.** *For every finite graph X there exists a function $f\colon \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that every graph with no K-fat X minor is $f(K)$-quasi-isometric to a graph with no X minor.*

Note that the case $K = 0$ of Conjecture 11.1.1 holds trivially (with $f(0) := (1, 0)$), since 0-fat X minors are precisely the X minors.

We remark that it is easy to see that the (qualitative) converse of Conjecture 11.1.1 holds for any (not even necessarily finite) graph X (cf. Lemmas 11.2.2 and 11.3.4).

Recently, Conjecture 11.1.1 has been disproved by Davies, Hickingbotham, Illingworth and McCarty for general graphs X [38, Theorem 1]. However, Conjecture 11.1.1 is known to be true for particular graphs X: The case $X = K_3$, which characterises graphs quasi-isometric to a forest, follows from a result of Manning [100] (see [75] or [19] for a graph-theoretic proof). The case $X = K_{2,3}$ was proved by Chepoi, Dragan, Newman, Rabinovich, and Vaxes [37], which characterises graphs quasi-isometric to an outerplanar graph. Further, Fujiwara and Papasoglu [70] proved the case $X = K_4^-$, a K_4 with an edge removed, which characterises graphs quasi-isometric to a cactus. The case $X = K_{1,m}$ was proved by Georgakopoulos and Papasoglu [75]. Additionally, we showed in Chapter 11 a similar, but even stronger characterisations of graphs quasi-isometric to a disjoint union of paths or to a disjoint union of subdivided stars, respectively.

12.1.2 Our results

The main contribution of this chapter is the resolution of Conjecture 11.1.1 for $X = K_4$:

Theorem 33. *There exists a function $f\colon \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that every graph with no K-fat K_4 minor is $f(K)$-quasi-isometric to a graph with no K_4 minor.*

It follows that every connected finite graph with no K-fat K_4 minor is either (M, A)-quasi-isometric to a 2-connected series-parallel graph where $(M, A) := f(K)$ for the function f from Theorem 33, or it can be disconnected by removing a ball of radius A.

In general, the construction of a quasi-isometry and the argument for its verification may be quite lengthy. For example, the proof of Conjecture 11.1.1 for $X = K_4^-$ given by Fujiwara and Papasoglu [70] is more than 20 pages long.

There is an alternative framework in which quasi-isometries between graphs can be cast, which enables us to look at Conjecture 11.1.1 from a different angle: Graph-decompositions [48] are a natural extension of tree-decompositions which allow the bags V_h of decompositions $(H, \mathcal{V})$ to be arranged along general decomposition graphs H instead of just trees. Recent applications of graph-decompositions include a local-global decomposition theorem [48] as well as the study of local separations [32] and of locally chordal graphs [1]. In [8, § 3.4] it was shown that two graphs are quasi-isometric if and only if each has a decomposition modelled on the other, subject to bounds on width parameters that correspond to the constants of the quasi-isometry (see Section 11.2 for details). To prove Theorem 33, we construct such a decomposition of G modelled on a graph H

with no K_4 minor instead of defining a quasi-isometry from G to H. This in particular yields a more graph-theoretic flavour of the arguments used in our proof of Theorem 33.

Our techniques additionally yield a new proof of Conjecture 11.1.1 for $X = K_4^-$, which is significantly shorter than the original one by Fujiwara and Papasoglu [70]:

Theorem 34. *There exists a function $f\colon \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that every graph with no K-fat K_4^- minor is $f(K)$-quasi-isometric to a graph with no K_4^- minor.*

Recall that the connected graphs with no K_4^- minor are precisely the cacti, i.e. those graphs in which any two cycles meet in at most one vertex. In particular, every connected graph with no K-fat K_4^- minor is $f(K)$-quasi-isometric to a cactus where f is the function from Theorem 34. We also remark that our proof yields a function f which is quite different from the one given by Fujiwara and Papasoglu.[1]

Let us finally emphasise that we prove Theorems 33 and 34 for arbitrary (finite or infinite) graphs.[2] Thus, Theorems 33 and 34 transfer verbatim to length spaces, a generalisation of geodesic metric spaces [75, Observation 2.1]

12.1.3 Applications

Bonamy, Davies, Esperet and Wesolek [24] brought the following consequences of Theorem 33 to our attention.

Gromov [80] introduced *asymptotic dimension* of metric spaces in the context of geometric group theory; see [15] for a survey on asymptotic dimension and its group-theoretic applications.[3] Bonamy, Bousquet, Esperet, Groenland, Liu, Pirot and Scott proved that for any finite graph X, the class of graphs with no X minor has asymptotic dimension at most 2 [23, Theorem 2], and they asked whether this can be generalised to classes of graphs excluding a graph X as a K-fat minor for some $K \in \mathbb{N}$ [23, Question 5]. It now follows from Theorem 33 that this is true for $X = K_4$.

Theorem 12.1.1. *For every $K \in \mathbb{N}$, the class of graphs with no K-fat K_4 minor has asymptotic dimension at most 1.*

Proof. Let $\mathcal{H}$ be the class of all graphs with no K_4 minor. Since K_4 is planar, $\mathcal{H}$ has asymptotic dimension at most 1 by [23, Theorem 1.2]. Let $K \in \mathbb{N}$. By Theorem 33, every graph with no K-fat

[1] For a given graph G with no K-fat K_4^- minor, we construct a $(84K + 2, 84K + 2)$-quasi-isometry from some cactus H to G (which then by Lemma 9.2.1 yields that G is quasi-isometric to H). In contrast to that Fujiwara and Papasoglu construct a $(2, 6000 \cdot 10^{100}K)$-quasi-isometry from some cactus H' to G; their constants can be found in [70, Proofs of Lemma 1.9 and Theorem 2.1]. We remark that neither of us optimised their constants.

[2] This made no difference in the proofs except for Lemma 12.4.8.

[3] For the various equivalent definitions of 'asymptotic dimension', we refer the reader to [23, § 1.1]. For example, the *asymptotic dimension* of a graph class $\mathcal{G}$ is the least $n \in \mathbb{N}$ such that there exists a function f such that for all $G \in \mathcal{G}$ and for all $r \in \mathbb{N}$, $V(G)$ can be coloured with $n+1$ colours so that for all $v, u \in V(G)$, if they are connected by a monochromatic path in G^r, then the distance between v, u in G is $\leqslant f(r)$ [23, Proposition 1.17].

K_4 minor is $f(K)$-quasi-isometric to a graph in $\mathcal{H}$ where f is the function from Theorem 33. Since asymptotic dimension is preserved under quasi-isometry [15, Proposition 22], the class of graphs with no K-fat K_4 minor thus has asymptotic dimension at most 1, too. $\square$

A class of graphs is *hereditary* if it is closed under taking induced subgraphs. Such a class of graphs is *χ-bounded* if there exists a function $f\colon \mathbb{N} \to \mathbb{N}$ such that $\chi(G) \leqslant f(\omega(G))$ for every graph G in the class, where χ and ω denote the *chromatic number* and the *clique number*, respectively.

Burling [25] constructed a sequence of graphs with increasing chromatic number that contain no K_3 as a subgraph; for a modern study of these graphs, see [108]. In the following, we will refer to the induced subgraphs of graphs in the Burling sequence as *Burling graphs*. From Theorem 12.1.1 it now follows that the class of Burling graphs provides an example for a graph class which has bounded asymptotic dimension but is not χ-bounded.

Example 12.1.2. *There exists a hereditary class of graphs which has asymptotic dimension at most* 1 *and is not χ-bounded.*

Proof. The hereditary class of Burling graphs is not χ-bounded, as Burling graphs do not contain the complete graph K_3, but have unbounded chromatic number. Hence, it remains to prove that the class of Burling graphs has asymptotic dimension at most 1. For this, by Theorem 12.1.1, it suffices to show that Burling graphs do not contain 2-fat K_4 minors. As Burling graphs do not contain K_3 as a subgraph, every 2-fat K_4 minor in a Burling graph G yields a subdivision of K_4 as an induced subgraph of G such that its branch vertices have distance at least 2. By [109, Theorem 7.4], every subdivision of a K_4 which is an induced subgraph of a Burling graph contains at least one non-subdivided edge. Hence, G cannot contain a 2-fat K_4 minor. $\square$

12.1.4 An open conjecture

Even though Conjecture 11.1.1 is disproved for general graphs X [38], we would like to draw the reader's attention to the following special case[4] of Conjecture 11.1.1, which is still open. It is a coarse variant of Kuratowski's theorem and was first conjectured by Georgakopoulos and Papasoglu [75]:

Conjecture 12.1.3. *There exists a function $f : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that every graph with no K-fat K_5 and $K_{3,3}$ minor is $f(K)$-quasi-isometric to a planar graph.*

[4]Note that one may replace the graph X in Conjecture 11.1.1 with any finite set of finite graphs.

12.1.5 How this chapter is organised

Section 12.2 contains some preparatory lemmas, which simplify the proofs of Theorems 33 and 34. We then prove Theorem 34 in Section 12.3 and Theorem 33 in Section 12.4.

12.2 Preparatory work: sequences of graph-decompositions

As we have mentioned already in the introduction, we will prove Theorems 33 and 34 by showing equivalent statements about graph-decompositions, that is, we will prove the cases $X = K_4^-$ and $X = K_4$ of the following conjecture, which is qualitatively equivalent to Conjecture 11.1.1 by Lemmas 11.2.1 and 11.2.2.

Conjecture 11.1.1'. *Let X be a finite graph. Then there exist a function $f\colon \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that every graph G with no K-fat X minor admits an honest $f(K)$-radial decomposition $(H, \mathcal{V})$ modelled on a graph H with no X minor.*

We will later refer to the coordinates of f as f_0 and f_1 in the sense of $f(K) = (f_0(K), f_1(K))$.

For this we construct for a given graph G the desired graph-decomposition of G recursively by 'extending' partial graph-decompositions of G step-by-step to one of G. More precisely, given a (suitable) partial graph-decomposition $(H^d, \mathcal{V}^d)$ of G with support Y^d we will extend it to a (R_0, R_1)-radial partial graph-decomposition $(H^{d+1}, \mathcal{V}^{d+1})$ of G with support $Y^{d+1} \supseteq Y^d$. Here, we will ensure that $N_G(Y^d) \subseteq V(Y^{d+1})$, so that the desired graph-decomposition may arise as the 'limit' of the $(H^d, \mathcal{V}^d)$.

In this section we will formalise this procedure and establish a theorem (Theorem 12.2.5) that will yield the graph-decomposition at once if we can do one simplified 'extension step'; see $(*)$ below.

The *restriction* $(H, \mathcal{V})$ of a graph-decomposition $(H, \mathcal{V}')$ of G to a subgraph Y of G is given by $V_h := V'_h \cap V(Y)$. Conversely, a partial graph-decomposition $(H, \mathcal{V})$ of G with support Y *extends* to a partial graph-decomposition $(H', \mathcal{V}')$ of G with support Y' if $H \subseteq H'$, $Y \subseteq Y'$, and $V_h \subseteq V'_h$ for every node h of H. Note that if $(H, \mathcal{V})$ and $(H, \mathcal{V}')$ are partial graph-decompositions of G with supports Y and Y', respectively, such that $Y \subseteq Y'$, and with the same decomposition graph H, then $(H, \mathcal{V})$ extends to $(H, \mathcal{V}')$ if and only if $(H, \mathcal{V})$ is the restriction of $(H, \mathcal{V}')$ to Y.

Let $(\mathcal{H}^d)_{d\in\mathbb{N}}$ be a sequence of partial graph-decompositions $\mathcal{H}^d = (H^d, \mathcal{V}^d)$ of G with supports Y^d such that $\mathcal{H}^d$ extends to $\mathcal{H}^{d+1}$ for every $d \in \mathbb{N}$. Then we set $Y := \bigcup_{n\in\mathbb{N}} Y^d$, $H := \bigcup_{d\in\mathbb{N}} H^d$ and $V_h := \bigcup_{d\in\mathbb{N}: h\in H^d} V_h^d$ for every node h of H. We call the pair $\mathcal{H} = (H, \mathcal{V})$ the *limit* of $(\mathcal{H}^d)_{d\in\mathbb{N}}$.

The following observation is immediate from the definitions:

Observation 12.2.1. *Let $(\mathcal{H}^d)_{d\in\mathbb{N}}$ be a sequence of partial graph-decompositions of a graph G such that $\mathcal{H}^d$ extends to $\mathcal{H}^{d+1}$ for every $d \in \mathbb{N}$. Then the limit $\mathcal{H}$ is a partial graph-decomposition of G with support $Y = \bigcup_{d\in\mathbb{N}} Y^d$. Moreover,*

(i) *if, for a finite graph X, $H^d \in \mathrm{Forb}_{\preccurlyeq}(X)$ for every $d \in \mathbb{N}$, then $H \in \mathrm{Forb}_{\preccurlyeq}(X)$,*

(ii) *if $\mathcal{H}^d$ is honest for every $d \in \mathbb{N}$, then $\mathcal{H}$ is honest,*

(iii) $\mathrm{rad}_G(V_h) \leqslant \sup_{d:h\in H^d} \mathrm{rad}_G(V_h^d)$ *for every node h of H, and*

(iv) $\mathrm{irads}(\mathcal{H}, v) \leqslant \sup_{d:v\in Y^d} \mathrm{irads}(\mathcal{H}^d, v)$ *for every vertex v of Y.* □

First, we describe how to extend a partial graph-decomposition $(H^d, \mathcal{V}^d)$ with support Y_d if it is possible to extend it in the direction of each remaining component of $G - Y_d$ individually, that is if we are given for each $C \in \mathcal{C}(G - Y^d)$ a partial graph-decomposition $\mathcal{H}^C = (H^C, \mathcal{V}^C)$ of G with support Y^C such that $N_G(C) \subseteq V(Y^C) \subseteq B_G(C, 1)$.

Let $\mathcal{H} = (H, \mathcal{V})$ be a partial graph-decomposition of G with support Y. For $C \in \mathcal{C}(G - Y)$, a partial graph-decomposition $\mathcal{H}^C = (H^C, \mathcal{V}^C)$ of G with support $Y^C \subseteq G[C, 1]$ is *feasible for $\mathcal{H}$ and C*, if $\partial_G C \subseteq V(Y^C)$ and there exist nodes h^C of H^C and g^C of H with $N_G(C) = V^C_{h^C}$ and $N_G(C) \subseteq V_{g^C}$. A family $\{\mathcal{H}^C \mid C \in \mathcal{C}(G - Y)\}$ of partial graph-decompositions of G is *feasible for $\mathcal{H}$* if each $\mathcal{H}^C$ is feasible for $\mathcal{H}$ and C.

Construction 12.2.2. Assume that $\mathcal{H} = (H, \mathcal{V})$ is a partial graph-decomposition of a graph G with support Y. Let $\{\mathcal{H}^C \mid C \in \mathcal{C}(G - Y)\}$ be a family of partial graph-decompositions $\mathcal{H}^C = (H^C, \mathcal{V}^C)$ of G with support Y^C which is feasible for $\mathcal{H}$.

We define the pair $\mathcal{H}' := (H', \mathcal{V}')$ as follows. Let H' be the graph obtained from the disjoint union of H and the H^C by identifying $h^C \in V(H^C)$ with $g^C \in V(H)$ for every $C \in \mathcal{C}(G - Y)$. For each node h of H', we set $V'_h := V_h$ if $h \in V(H)$ and $V'_h := V^C_h$ if $h \in V(H^C) \setminus \{h^C\}$ for some $C \in \mathcal{C}(G - Y)$. Set $Y' := Y \cup \bigcup_{C\in\mathcal{C}(G-Y)} Y^C$.

Note that the graph H' from Construction 12.2.2 is obtained from H and the H^C by 1-sums. Here, the 1-*sum* of two graphs G_0 and G_1 with respect to vertices v_0 of G_0 and v_1 of G_1 is obtained from the disjoint union of G_0 and G_1 by identifying v_0 with v_1. We remark that $\mathrm{Forb}_{\preccurlyeq}(X)$ is closed under 1-sums if and only if X is 2-connected. Hence, if H and the H^C have no X-minor for $X = K_4$ or $X = K_4^-$, then also the arising graph H' has no X minor[5].

The following observation about $\mathcal{H}'$ from Construction 12.2.2 is immediate from its construction.

Observation 12.2.3. *Let $G, \mathcal{H}$ and the $\mathcal{H}^C$ be as in Construction 12.2.2. Then $\mathcal{H}'$ constructed as in Construction 12.2.2 from $\mathcal{H}$ and the $\mathcal{H}^C$ is a partial graph-decomposition of G with support Y', which extends $\mathcal{H}$. Moreover,*

[5]One could, more generally, obtain a partial graph-decomposition $(H', \mathcal{V}')$ by identifying more than one pair of vertices of H and an H^C if the respective bags of H and H^C are suitable for this. However, since we need to ensure for the proof of Theorems 33 and 34 that the arising graph H' has no K_4^-- or K_4 minor, we will not do this here.

(i) *if* $H, H^C \in \mathrm{Forb}_{\preccurlyeq}(X)$ *for all* $C \in \mathcal{C}(G-Y)$ *and some* 2*-connected graph* X*, then* $H' \in \mathrm{Forb}_{\preccurlyeq}(X)$*,*

(ii) *if* $\mathcal{H}$ *and* $\mathcal{H}^C$ *are honest for every* $C \in \mathcal{C}(G-Y)$*, then* $\mathcal{H}'$ *is honest,*

(iii) $\mathrm{rad}_G(V'_h) = \mathrm{rad}_G(V_h)$*, if* $h \in V(H)$*,*

(iv) $\mathrm{rad}_G(V'_h) = \mathrm{rad}_G(V_h^C)$*, if* $h \in V(H^C) \setminus \{h^C\}$ *for some* $C \in \mathcal{C}(G-Y)$*,*

(v) $\mathrm{irads}(\mathcal{H}', v) = \mathrm{irads}(\mathcal{H}, v)$ *for* $v \in V(Y) \setminus \partial_G Y$*,*

(vi) $\mathrm{irads}(\mathcal{H}', v) = \mathrm{irads}(\mathcal{H}^C, v)$ *for* $v \in V(Y^C) \setminus \partial_G Y = V(Y^C \cap C)$ *for* $C \in \mathcal{C}(G-Y)$*, and*

(vii) $\mathrm{irads}(\mathcal{H}', v) \leqslant \mathrm{irads}(\mathcal{H}, v) + 2\sup_{C \in \mathcal{C}(G-Y)} \mathrm{irads}(\mathcal{H}^C, v)$ *for* $v \in \partial_G Y$*.* □

A partial graph-decomposition $\mathcal{H}$ of G with support Y is *R-component-feasible* for $R \in \mathbb{N}$ if for every component C of $G - Y$ there is $h \in V(H)$ with $N_G(C) \subseteq V_h$ and $\mathrm{rad}_G(V_h) \leqslant R$. A subgraph Y of G is *R-ball-componental* if for every component C of $G - Y$ there exists a vertex w of G such that C is a component of $G - B_G(w, R)$.

The following lemma shows that we may assume without loss of generality that every R-component-feasible partial graph-decomposition has an R-ball-componental support.

Lemma 12.2.4. *For every* $R \in \mathbb{N}$*, every honest* R*-component-feasible partial decomposition* $\mathcal{H}$ *of a graph* G *with support* Y *extends to an honest* R*-component-feasible partial graph-decomposition* $\mathcal{H}'$ *of* G *modelled on the same decomposition graph with* R*-ball-componental support* Y' *and* $\mathrm{oradw}(\mathcal{H}') \leqslant \max\{\mathrm{oradw}(\mathcal{H}), R\}$*,* $\mathrm{irads}(\mathcal{H}', v) = \mathrm{irads}(\mathcal{H}, v)$ *for* $v \in V(Y)$*, and* $\mathrm{irads}(\mathcal{H}', v) = 0$ *for* $v \in V(Y') \setminus V(Y)$*.*

Proof. For each node h of H with $\mathrm{rad}_G(V_h) \leqslant R$, we fix a vertex v_h of G such that $V_h \subseteq B_G(v_h, R)$. Since $(H, \mathcal{V})$ is R-component-feasible, we may further fix for each $C \in \mathcal{C}(G-Y)$, a node h_C of H such that $N_G(C) \subseteq V_{h_C}$ and $\mathrm{rad}_G(V_{h_C}) \leqslant R$. Set $U := V(Y) \cup \bigcup_{C \in \mathcal{C}(G-Y)} (B_G(v_{h_C}, R) \cap V(C))$. It is immediate from the construction that every component C' of $G - U$ is also a component of $G - B_G(v_{h_C}, R)$ where C is the unique component of $G - Y$ with $C' \subseteq C$. Hence, the induced subgraph $Y' := G[U]$ is R-ball-componental.

It remains to extend $(H, \mathcal{V})$ to the desired graph-decomposition $(H, \mathcal{V}')$ of Y'. Set $V'_h := V_h \cup (B_G(v_h, R) \cap (\bigcup_{C \in \mathcal{C}(G-Y): h_C = h} V(C)))$ for every $h \in V(H)$. Note that, for every $h \in V(H)$ every vertex v of Y is contained in V'_h if and only if it is contained in V_h, and every vertex $v \in U \setminus V(Y)$ is contained in exactly one V'_h. Thus, $(H, \mathcal{V}')$ is indeed a graph-decomposition of Y', and its outer-radial width and radial spread are as claimed. Moreover, by construction, $N_G(C') \subseteq V'_{h_C}$ where C is the unique component of $G - Y$ with $C' \subseteq C$. Hence, $(H, \mathcal{V}')$ is R-component-feasible. □

Recall that our overall goal is to construct, for a given graph X and a graph G without K-fat X minor, a graph-decomposition of G along a graph H without X minor. For this, we recursively construct $f(K)$-radial partial decompositions $(H^d, \mathcal{V}^d)$ with support Y^d, where each $(H^d, \mathcal{V}^d)$ extends $(H^{d-1}, \mathcal{V}^{d-1})$ 'into the direction' of each component. Combining Observation 12.2.1 and Lemma 12.2.4, it suffices to consider the case $Y^d = G[B]$ where B is a ball of bounded radius and to 'make progress' into the direction of one component of $G - Y^d$. The following result, Theorem 12.2.5, and property $(*)$ formalise this.

We say that a graph X has *property* $(*)$ if for every $K \in \mathbb{N}$ there are $R(K)$, $f_0'(K)$, $f_1'(K)$, $f_1''(K) \in \mathbb{N}$ with $R(K) \leqslant f_0'(K)$ such that, for every ball B of radius $\leqslant R(K)$ in a graph G with no K-fat X minor and for every component C of $G - B$, there exists an honest $R(K)$-component-feasible $(f_0'(K), f_1'(K))$-radial partial graph-decomposition $\mathcal{H}^C$ of G with support Y^C modelled on a graph $H^C \in \mathrm{Forb}_{\preccurlyeq}(X)$ such that $N_G(C) \cup \partial_G C \subseteq V(Y^C) \subseteq B_G(C, 1)$, the neighbourhood $N_G(C)$ is a bag of $\mathcal{H}^C$, and $\mathrm{irads}(\mathcal{H}^C, v) \leqslant f_1''(K)$ for all $v \in N_G(C)$.

Theorem 12.2.5. *If a* 2*-connected finite graph* X *has property* $(*)$ *(with respect to functions* R, f_0', f_1', f_1''*), then Conjecture 11.1.1' holds for* X *(with respect to* $f := (f_0, f_1) := (f_0', f_1' + 2 \cdot f_1'' + 1)$*).*

A partial graph-decomposition $\mathcal{H}$ of G with support Y is (R_0, R_1, R_1')*-radial* in G if it is (R_0, R_1)-radial and $\mathrm{irads}(\mathcal{H}, v) \leqslant R_1'$ for every $v \in \partial_G Y$.

Proof. For $K \in \mathbb{N}$, let $R(K), f_0'(K), f_1'(K), f_1''(K)$ be given by property $(*)$ of X, and let G be any graph with no K-fat X minor. By Observation 9.4.3, we may assume that G is connected.

We define a sequence of partial graph-decompositions $(H^d, \mathcal{V}^d)$ of G with supports Y^d as follows. Pick an arbitrary vertex v of G and initialise $Y^0 := G[v, R_0(K)]$. Let $(H^0, \mathcal{V}^0)$ be the partial graph-decomposition of G with support Y^0 whose decomposition graph H^0 is the graph on a single vertex h and whose single bag is $V_h^0 := V(Y^0)$. Clearly, $(H^0, \mathcal{V}^0)$ is honest, $R(K)$-component-feasible and $(f_0(K), 0, 0)$-radial, and $H^0 \in \mathrm{Forb}_{\preccurlyeq}(X)$.

Now let $d \in \mathbb{N}$ be given, and assume that for $i \leqslant d$ we have already constructed a sequence of honest, $R(K)$-component-feasible, $(f_0(K), f_1(K), f_1'(K))$-radial partial graph-decompositions $\mathcal{H}^i := (H^i, \mathcal{V}^i)$ of G with $R(K)$-ball-componental support Y^i and with $H^i \in \mathrm{Forb}_{\preccurlyeq}(X)$ such that each $\mathcal{H}^{i+1}$ extends $\mathcal{H}^i$. Let further $(H^C, \mathcal{V}^C)$, for every $C \in \mathcal{C}(G - Y^d)$, be given by property $(*)$ of X.

Then applying Construction 12.2.2 to $\mathcal{H}^d$ and the $(H^C, \mathcal{V}^C)$ yields a partial graph-decomposition $(H', \mathcal{V}')$ of G, which extends $\mathcal{H}^d$, is honest and $(f_0(K), f_1(K), f_1'(K))$-radial by Observation 12.2.3, and such that $H' \in \mathrm{Forb}_{\preccurlyeq}(X)$, since X is 2-connected. Moreover, $(H', \mathcal{V}')$ is clearly still $R(K)$-component-feasible, so by Lemma 12.2.4, $(H', \mathcal{V}')$ extends to an honest $R(K)$-component-feasible $(f_0(K), f_1(K), f_1'(K))$-radial partial graph-decomposition $(H^{d+1}, \mathcal{V}^{d+1})$ of G with $R(K)$-ball-componental support Y^{d+1} and with $H^{d+1} = H' \in \mathrm{Forb}_{\preccurlyeq}(X)$.

Now by Observation 12.2.1, the limit $(H, \mathcal{V})$ of the $(H^d, \mathcal{V}^d)$ is an honest $(f_0(K), f_1(K))$-radial partial graph-decomposition of G with support $Y := \bigcup_{d \in \mathbb{N}} Y^d$ and with $H \in \mathrm{Forb}_{\preccurlyeq}(X)$. Since Y^{d+1}

contains $G[Y^d, 1]$ by construction, and because G is connected, we have $Y = G$, and hence $(H, \mathcal{V})$ is a graph-decomposition of G. Thus, $(H, \mathcal{V})$ is as desired. $\square$

12.3 Proof of the fat minor conjecture for $X = K_4^-$

In this section we prove Theorem 34, which we restate here in the terminology of graph-decompositions.

Theorem 34'. *For every $K \in \mathbb{N}$, every graph with no K-fat K_4^- minor admits an honest $(42K+1, 28K+3)$-radial decomposition $(H, \mathcal{V})$ modelled on a graph H with no K_4^- minor.*

Note that by Lemmas 9.2.1 and 11.2.1, Theorem 34' immediately implies Theorem 34.

To prove Theorem 34', it suffices by Theorem 12.2.5 to show that for every ball B in G of radius $\leqslant 42K + 1$ and every component C of $G - B$ there exists a $(42K + 1, 28K + 1)$-radial, $(42K+1)$-component-feasible partial decomposition $\mathcal{H}^C$ of G modelled on a cactus H^d with support $Y^C \subseteq G[C, 1]$ such that $\partial_G C \cup N_G(C) \subseteq V(Y^C)$ and $\text{irads}(\mathcal{H}^C, v) \leqslant 1$ for all $v \in N_G(C)$.

For this, we first show that if G has no K-fat K_4^- minor, then no component C of $G - B$ contains three vertices in its boundary $\partial_G C$ that are pairwise far apart in G (see Lemma 12.3.1 and Figure 12.1). It follows that the boundary of every component C of $G - B$ can be partitioned into two sets N_1, N_2 of small radius. If N_1 and N_2 are close in G, then the boundary of C has small radius, and we may obtain the desired partial graph-decomposition $\mathcal{H}^C$ of G by letting H^C be a K_2 on vertices h, h' and setting $V_h^C := N_G(C), V_{h'}^C := N_G(C) \cup \partial_G C$. Otherwise, the sets N_1 and N_2 are far apart in G. We then show that there is a partial path-decomposition $\mathcal{H}^C$ with support $Y^C \subseteq G[C, 1]$ as desired 'connecting' N_1 and N_2 such that its first and last bag are its only bags which intersect B (see Lemma 12.3.3 and Figure 12.2).

Lemma 12.3.1. *Let B be a ball in a graph G around a vertex w, and let C be a component of $G - B$. If $\partial_G C$ contains three vertices which are pairwise at least $7K$ apart in G for some $K \in \mathbb{N}_{\geqslant 1}$, then G has a K-fat K_4^- minor.*

Proof. Let r be the radius of the ball B around w, and let u_1, u_2, u_3 be three vertices in $\partial_G C$ which are pairwise at least $7K$ apart in G. Then also $d_G(u_i, u_j) \leqslant 2r$ for $i \neq j$ since B has radius r, and thus

$$r \geqslant d_G(u_i, u_j)/2 \geqslant 7K/2 > 3K.$$

Hence, the ball $B_G(w, r - (3K - 1))$ is non-empty, so we may pick for every $i \in [3]$ a shortest v_i–$B_G(w, r - (3K - 1))$ path $Q^i = q_0^i \dots q_{3K}^i$ in G. By defining the branch sets as

$$V_1 := V(C),\ V_2 := B_G(w, r - (3K - 1)),\ V_3 := V(q_K^1 Q^1 q_{2K}^1),\ V_4 := V(q_K^2 Q^2 q_{2K}^2)$$

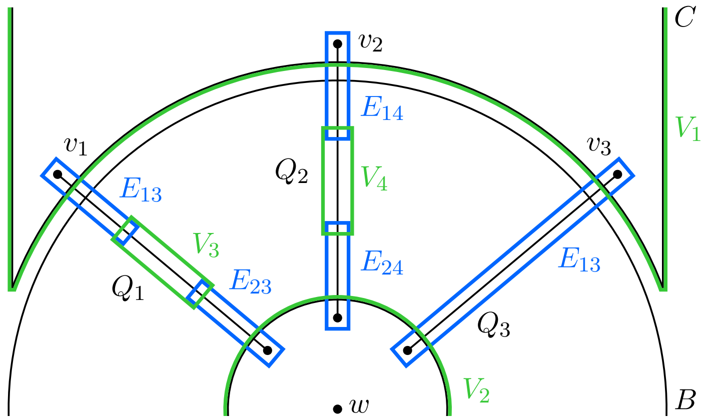

FIGURE 12.1: The K-fat K_4^- minor in the proof of Lemma 12.3.1.

and the branch paths as

$$E_{13} := q^1_{2K} Q^1 q^1_{3K},\ E_{14} := q^2_{2K} Q^2 q^2_{3K},\ E_{23} := q^1_0 Q^1 q^1_K,\ E_{24} := q^2_0 Q^2 q^2_K,\ E_{12} := Q^3$$

we obtain a K-fat K_4^--model (see Figure 12.1). Indeed, the sets V_i and the paths E_{ij} obviously form a model of K_4^-, so it remains to show that this K_4^--model is K-fat.

Since $G[V_1] = C$ is a component of $G - B_G(w, r)$ and $V_2 = B_G(w, r - (3K - 1))$, we have $d_G(V_1, V_2) \geqslant 3K \geqslant K$. In particular, as the Q^i are shortest paths, it follows that $d_G(V_i, E_{(3-i)j}) \geqslant 2K \geqslant K$ as well as $d_G(E_{1j}, E_{2j}) \geqslant K$ for $i = 1, 2$ and $j = 3, 4$. By the choice of V_3, V_4 and the E_{ij} as subsets of the Q^i, it remains to show $d_G(Q^i, Q^j) \geqslant K$ for $i \neq j$. Indeed, we find

$$d_G(Q^i, Q^j) \geqslant d_G(u_i, u_j) - ||Q_i|| - ||Q_j|| \geqslant 7K - 2 \cdot 3K = K,$$

where the first inequality holds because any Q^i–Q^j path extends via Q^i and Q^j to a u_i–u_j path. $\square$

For the next lemma, we need Lemma 11.4.2 from Chapter 11, which we recall here for convenience (in a slightly weaker form which suffices in this case).

Lemma 12.3.2. *Let G be a graph, and let P be a geodesic path in G. Given $r \in \mathbb{N}$ we write $B_p := B_G(p, r)$ for $p \in V(P)$ and set*

$$V_p := \bigcup_{p' \in B_P(p, r+1)} B_{p'}.$$

Then $\mathcal{H} = (P, (V_p)_{p \in P})$ is a partial graph-decomposition of G with support $G' = G[\bigcup_{p \in P} B_p]$ of outer-radial width at most $2r + 1$ and radial spread at most $2r + 1$. Moreover, $V_p \subseteq B_G(p, 2r + 1)$ for every $p \in V(P)$.

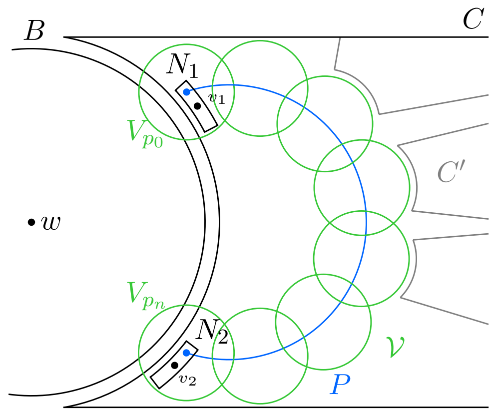

FIGURE 12.2: Depicted in green is the path-decomposition $(P, \mathcal{V})$ from Lemma 12.3.3 along the path P which is depicted in blue. By (ii), the neighbourhood of every component C' of $C - Y$ is contained in some bag in $\mathcal{V}$.

Lemma 12.3.3. *Let B be a ball in a graph G around a vertex w, and let C be a component of $G - B$ whose boundary $\partial_G C$ contains two vertices that are at least $42K + 1$ apart for $K \in \mathbb{N}_{\geqslant 1}$.*

If G has no K-fat K_4^- minor, then there is an induced subgraph Y of $G[C, 1]$ which admits an honest $(28K + 1, 28K + 1)$-radial decomposition $(P, \mathcal{V})$ modelled on some path $P = p_0 \dots p_n$ such that

(i) *$N_G(C) \cup \partial_G C \subseteq V_{p_0} \cup V_{p_n}$, and $N_G(C)$ is disjoint from every V_p with $p \neq p_0, p_n$, and*

(ii) *for every component C' of $G - Y$ which meets, or equivalently is contained in, C, its neighbourhood $N_G(C')$ is contained in some bag V_p of $(P, \mathcal{V})$.*

Proof. Let us assume that G has no K-fat K_4^- minor and show that G then has a partial path-decomposition as claimed. Let v_1, v_2 be two vertices in $\partial_G C$ which are at least $42K + 1$ apart. Then the balls B_1 and B_2 in G of radius $21K - 1 \leqslant \lfloor d_G(v_1, v_2)/2 \rfloor - 1$ around v_1 and v_2, respectively, are disjoint and also there is no B_1–B_2 edge in G. In particular, the sets $N_i := B_G(v_i, 7K - 1) \cap \partial_G C \subseteq B_i$ for $i \in [2]$ are at least $28K + 3$ apart because for all $n_1 \in N_1, n_2 \in N_2$ we have

$$d_G(n_1, n_2) \geqslant d_G(v_1, v_2) - d_G(n_1, v_1) - d_G(v_2, n_2) \geqslant (42K + 1) - 2 \cdot (7K - 1) = 28K + 3. \quad (12.1)$$

Further, since G has no K-fat K_4^- minor, Lemma 12.3.1 yields $N_1 \cup N_2 = \partial_G C$.

Let $P = p_0 \dots p_n$ be a shortest N_1–N_2 path in $G' := G[B_1 \cup V(C) \cup B_2]$. Since B_1 and B_2 are disjoint and not joined by an edge, N_i, for $i \in [2]$, separates $B_i \setminus V(C)$ and $N_{3-i} \subseteq V(C)$ in G'.

Hence, every N_1–N_2 path in G', in particular P, is contained in C. Moreover, for $r_P := 14K$, we have

$$B_G(P, r_P) \subseteq B_G(C, r_P) = B_G(N_1, r_P) \cup V(C) \cup B_G(N_2, r_P) \subseteq B_1 \cup V(C) \cup B_2 = V(G'), \quad (12.2)$$

where we used that $N_1 \cup N_2 = \partial_G C$, $N_i \subseteq B_G(v_i, 7K-1)$ and $B_i = B_G(v_i, 21K-1)$. In particular,

$$B_{G'}(p, r_P) = B_G(p, r_P) \quad (12.3)$$

for all $p \in P$. So by Lemma 12.3.2, the pair $(P, \mathcal{V}')$ given by

$$V'_p := \bigcup_{p' \in B_P(p, r_P+1)} B_G(p', r_P)$$

is an honest graph-decomposition of $G'[P, r_P] = G[P, r_P]$ of radial width at most $2r_P + 1 = 28K+1$ and radial spread at most $2r_P + 1 = 28K + 1$. Moreover, $d_G(p, v) \leqslant 28K + 1$ for all $p \in V(P)$ and $v \in V'_p$.

Now set $Y := G[P, r_P] \cap G[C, 1]$. As we will see in a moment, the restriction of $(P, \mathcal{V}')$ to Y has already all desired properties, except that it does not satisfy the second part of (i). To solve this problem, we adapt $(P, \mathcal{V}')$ by deleting $N_G(C)$ from all bags V'_p of internal nodes $p \neq p_0, p_n$ of P. Formally, define the pair $(P, \mathcal{V})$ by $V_p := (V'_p \cap V(Y)) \setminus N_G(C)$ for $p \neq p_0, p_n$ and $V_p := V'_p \cap V(Y)$ for $p = p_0, p_n$. We claim that $(P, \mathcal{V})$ and Y are as desired.

We first prove that $(P, \mathcal{V})$ is indeed a graph-decomposition of Y. Since $(P, \mathcal{V}')$ is a graph-decomposition of $G[P, r_P]$, its restriction to $Y \subseteq G[P, r_P]$ is a graph-decomposition of Y. Thus, it suffices to show that every vertex $v \in N_G(C)$ is contained in precisely one of the bags V'_{p_0}, V'_{p_n} and that this bag also contains all neighbours of v.

Since $p_0 \in N_1$, we have $N_1 \subseteq B_G(p_0, 7K-1)$. Otherwise, there is some $u \in N_1$ with $d_G(u, p_0) \geqslant 7K$. Since also $d_G(u, v_2), d_G(p_0, v_2) \geqslant d_G(N_1, N_2) \geqslant 7K$ by (12.1), Lemma 12.3.1 applied to $u, p_0, v_2 \in \partial_G C$ yields a K-fat K_4^- minor in G, which is a contradiction. Analogously, we find $N_2 \subseteq B_G(p_n, 7K-1)$. Since $d_G(p_0, N_2), d_G(p_n, N_1) \geqslant 28K+3$ by (12.1) but $d_G(p, v) \leqslant 28K+1$ for all $p \in V(P)$ and $v \in V'_p$, the set $B_G(N_1, 1)$ is disjoint from V'_{p_n}, and $B_G(N_2, 1)$ is disjoint from V'_{p_0}.

Further, since $(P, \mathcal{V}')$ has outer-radial width and radial spread at most $28K+1$, so does $(P, \mathcal{V})$. It remains to show that $(P, \mathcal{V})$ satisfies (ii). Suppose for a contradiction that (ii) does not hold, that is, there is a component C' of $G - Y$ which is contained in C and which contains $z_1 \in N_G(B_G(p_{j_1}, r_P))$ and $z_2 \in N_G(B_G(p_{j_2}, r_P))$ with $0 \leqslant j_1 \leqslant j_2 \leqslant n$ such that $j_2 - j_1 > 2 \cdot (r_P + 1)$.

To derive a contradiction, we find a K-fat K_4^- minor in G as follows. Let $Q^i = q^i_0 \dots q^i_{r_P+1}$ be

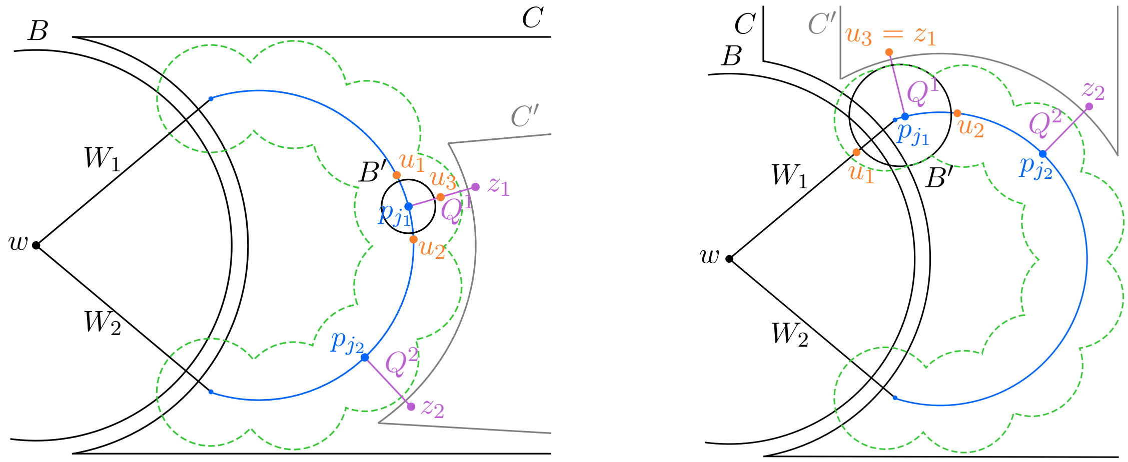

(a) Case 1: $j_1 \geqslant 7K$ and thus $u_1 \in V(P)$

(b) Case 2: $j_1 < 7K$ and thus $u_1 \in V(W_1)$

FIGURE 12.3: The situation in the proof of Lemma 12.3.3 if $(P, \mathcal{V})$ does not satisfy (ii).

a shortest p_{j_i}–z_i path in G and let $W^i = w_0^i \dots w_{r+1}^i$ be a shortest w–v_i' path in G for $v_1' = p_0$ and $v_2' = p_n$. Note that by definition of z_i and p_{j_i}, the Q^i have length $r_P + 1$ and are hence shortest P–C' paths in G. Further, recall that the radius r of the ball B is at least $d_G(v_1, v_2)/2 \geqslant 14K + 1$.

Now first assume that $j_1 \geqslant 7K$ (see Figure 12.3a). Set $B' := B_G(p_{j_1}, 7K)$, and consider $u_1 := p_{j_1 - 7K - 1}$, $u_2 := p_{j_1 + 7K + 1}$ and $u_3 := q^1_{7K+1}$. Since P is a shortest path in G' and Q^1 is also a shortest P–C' path in G, it follows from (12.3) for p_{j_1} that $u_1, u_2, u_3 \in N_G(B')$. A similar reasoning applied to u_1 and u_2 yields that the u_i have pairwise distance at least $7K$. We claim that u_1, u_2, u_3 lie in the same component of $G - B'$, which by Lemma 12.3.1 yields a K-fat K_4^- minor in G – a contradiction. Indeed, $Z := p_0 P u_1 \cup W^1 \cup W^2 \cup u_2 P p_n \cup Q^2 \cup C' \cup u_3 Q^1 z_1$ is connected and contains u_1, u_2, u_3. Moreover, it follows from a similar reasoning as above together with $j_2 - j_1 > 2(r_P + 1)$, also Q^2 being a shortest P–C' path in G that Z is disjoint from B', and $W^1, W^2 \subseteq G[B, 1]$.

Second, assume that $j_1 < 7K$ (see Figure 12.3b). Set $B' := B_G(p_{j_1}, 14K)$, and consider $u_1 := w^1_{r - 14K + j_1}$, $u_2 := p_{j_1 + 14K + 1}$ and $u_3 := z_1$. As above, $u_1, u_2, u_3 \in N_G(B')$ and the u_i have pairwise distance at least $7K$. Again, we obtain a contradiction from Lemma 12.3.1, since u_1, u_2, u_3 lie in the same component of G'_B. Indeed, $Z := W^1 u_1 \cup W^2 \cup u_2 P p_n \cup Q^2 \cup C'$ is connected and contains u_1, u_2, u_3, and one may follow a similar reasoning as above together with $d_G(p_0, p_n) > 2 \cdot (14K + 1)$ by (12.1) to show that Z is also disjoint from B'. □

We are now in a position to prove Theorem 34'. For this, by Theorem 12.2.5, it suffices to show that K_4^- has property $(*)$.

Lemma 12.3.4. *K_4^- has property $(*)$ (with respect to $R(K) := f_0'(K) := 42K+1$, $f_1'(K) := 28K+1$ and $f_1''(K) = 1$).*

Proof. Let $K \in \mathbb{N}_{\geqslant 1}$, and let G be a graph with no K-fat K_4^- minor. Let B be a ball in G of radius at most $42K+1$, and let C be a component of $G - B$.

If every two vertices in $\partial_G C$ are at most $42K$ apart, then the desired honest $(42K+1)$-component-feasible $(42K+1, 1)$-radial graph-decomposition $(H^C, \mathcal{V}^C) := (H, \mathcal{V})$ of $Y^C := Y := G[N_G(C) \cup \partial_G C]$ is given by defining H as a K_2 on two vertices h, h' and setting $V_h := N_G(C)$, $V_{h'} := N_G(C) \cup \partial_G C$.

Otherwise, we apply Lemma 12.3.3 to the ball B and the component C to obtain an honest decomposition $(P, \mathcal{V})$ of an induced subgraph Y of $G[C, 1]$, modelled on some path $P = p_0 \dots p_n$ such that $(P, \mathcal{V})$ has outer-radial width at most $28K+1$, has radial spread at most $28K+1$, and satisfies (i) and (ii). We obtain the desired honest $(28K+1)$-component-feasible $(28K+1, 28K+1)$-radial graph-decomposition $(H^C, \mathcal{V}^C) := (H, \mathcal{V})$ of $Y^C := Y$ by adding a node h and edges hp_0, hp_n to P and setting $V_h := N_G(C)$. Indeed, (i) ensures that $(H, \mathcal{V})$ is an (honest) graph-decomposition of Y and the radial spread of each vertex in $N_G(C)$ in $(H, \mathcal{V})$ is at most 1. The radial spread of every other vertex in Y did not change, that is, it is still at most $28K+1$. Since $V_h = N_G(C) \subseteq B$, we have $\mathrm{rad}_G(V_h) \leqslant \mathrm{rad}_G(B) \leqslant 42K+1$, and thus $(H, \mathcal{V})$ has outer-radial width at most $42K+1$. □

Proof of Theorem 34'. By Theorem 12.2.5, Lemma 12.3.4 immediately yields Theorem 34'. □

Proof of Theorem 34. Let G be a graph with no K-fat K_4^- minor. By Lemma 11.2.1, Theorem 34' yields that there exists a graph H with no K_4^- minor which is $(84K+2, 84K+2)$-quasi-isometric to G. Hence, it follows from Lemma 9.2.1 that G is $(84K+2, 3 \cdot (84K+2)^2)$-quasi-isometric to H. Thus, Theorem 34 holds with $f(K) := (84K+2, 3 \cdot (84K+2)^2)$. □

12.4 Proof of the fat minor conjecture for $X = K_4$

In this section we prove Theorem 33, which we restate here in the terminology of graph-decompositions.

Theorem 33'. *Every graph G with no K-fat K_4 minor for $K \in \mathbb{N}_{\geqslant 1}$ admits an honest $(25235K + 71, 22)$-radial decomposition $(H, \mathcal{V})$ modelled on a graph H with no K_4 minor.*

Note that by Lemmas 9.2.1 and 11.2.1, Theorem 33' immediately implies Theorem 33.

Two-terminal graphs are graphs with two distinguished (not necessarily distinct) vertices h_1, h_2, its *terminals*, which we refer to as *source* and *sink*. We denote by $\mathcal{H}_{SP}$ the class of all two-terminal graphs H with terminals $h_1, h_2 \in V(H)$ that satisfy $H + h_1h_2 \in \mathrm{Forb}_{\preccurlyeq}(K_4)$.[6] We remark that the finite 2-connected graphs in $\mathcal{H}_{SP}$ are precisely the 2-connected series-parallel graphs.

A *parallel composition* of two-terminal graphs is obtained by identifying the sources and identifying the sinks. A *series composition* of a pair of two-terminal graphs is obtained by identifying the sink of one of them with the source of the other one. Since K_4 is 3-connected, the following facts about $\mathcal{H}_{SP}$ follow easily.

Proposition 12.4.1. *$\mathcal{H}_{SP}$ is closed under the following operations:*

- *parallel and series composition,*
- *subdividing edges,*
- *adding a path of length at least 2 between any two adjacent vertices, and*
- *(infinite) union over graphs $H_0 \subseteq H_1 \subseteq \ldots \in \mathcal{H}_{SP}$.*

Proof. Since K_4 is 3-connected, the first three assertions follow easily. For the fourth assertion, note that since K_4 is finite, $H + h_1h_2$, with $H := \bigcup_{i\in\mathbb{N}} H_i$, contains a model of K_4 if and only if it contains one with finite branch sets. By the definition of H as the union of the $H^i \in \mathcal{H}_{SP}$, this implies that $H + h_1h_2$ has no K_4 minor, and hence $H \in \mathcal{H}_{SP}$. $\square$

Throughout this section we fix

$$\begin{aligned}
R_0(K) &:= 129 \cdot 5K + 5K = 130 \cdot 5K,\\
R_0'(K) &:= 3R_0(K) + 5K + 1,\\
\ell(K) &:= 2R_0(K) + 5K + 2,\\
R_1(K) &:= 4 \cdot (2(\ell(K) + 22K + 1) + 11K + 2), \text{ and}\\
R_2(K) &:= 2R_1(K) + 5K + 3.
\end{aligned}$$

Let us briefly sketch the proof of Theorem 33'. For this, let $K \in \mathbb{N}_{\geqslant 1}$, and let G be a graph with no K-fat K_4 minor. By Theorem 12.2.5 (with $R(K) := R_2(K)$, $f_0'(K) := R_2(K) + 2R_0'(K) + 2$ and $f_1'(K) = f_1''(K) := 7$), and because $R_2(K)+2R_0'(K)+2 = 25235K+71$ and $7+2\cdot 7+1 = 22$, it suffices to show that K_4 satisfies property $(*)$, that is for every ball B in G of radius $\leqslant R_2(K)$ and every component C of $G - B$ there exists a $R_2(K)$-component-feasible, $(R_2(K) + 2R_0'(K) + 2, 7)$-radial partial decomposition $\mathcal{H}^C$ of G modelled on a graph $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$ with support $Y^C \subseteq G[C, 1]$ such that $\partial_G C \subseteq V(Y^C)$ (see Lemma 12.4.4). So let such a ball B and component C of $G - B$ be given. To define $\mathcal{H}^C$, we distinguish two cases. If the boundary $\partial_G D$ of some component D of $C - B_G(B, 22K + 1)$ contains three vertices that are pairwise far apart, then we obtain the desired

[6] Note that we allow $h_1 = h_2$, in which case h_1h_2 is a loop. Thus, we have for two-terminal graphs H with $h_1 = h_2$ that $H \in \mathcal{H}_{SP}$ if and only if $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$.

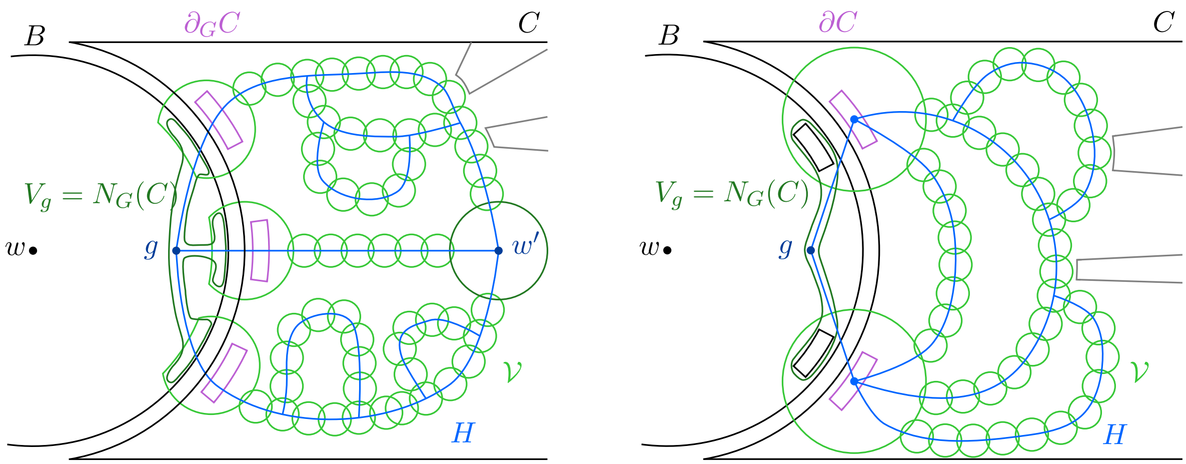

(a) Lemma 12.4.2

(b) Lemma 12.4.3

FIGURE 12.4: Depicted are the graph-decompositions $(H, \mathcal{V})$ with support Y in Lemmas 12.4.2 and 12.4.3 if G has no K-fat K_4 minor. By Lemma 12.4.2 (iv) and Lemma 12.4.3 (iv) the neighbourhood of every component of $C - Y$ is contained in some bag in $\mathcal{V}$.

partial decomposition with support $Y^C \subseteq G[C, 1]$ by applying the following Lemma 12.4.2 to C and $C^* := D$.

Lemma 12.4.2. *Let $K \in \mathbb{N}_{\geqslant 1}$, let B be a ball in a graph G around a vertex w of radius $r \in \mathbb{N}_{\geqslant 1}$, and let C be a component of $G - B$. Suppose there is a component C^* of $C - B_G(B, 22K + 1)$ such that $\partial_G C^*$ contains three vertices that are pairwise at least $R_1(K)$ apart.*

If G has no K-fat K_4 minor, then there is an induced subgraph Y of $G[C, 1]$ with $\partial_G C \subseteq V(Y)$ which admits an honest decomposition $(H, \mathcal{V})$ modelled on some graph $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$ such that

(i) *$(H, \mathcal{V})$ contains $N_G(C)$ as a bag,*

(ii) $\mathrm{oradw}((H, \mathcal{V})) \leqslant \max\{r + 2R_0'(K) + 2, R_2(K)\}$ *and*

(iii) $\mathrm{irads}((H, \mathcal{V})) \leqslant 7$, *and*

(iv) *$(H, \mathcal{V})$ is R-component-feasible for $R := \max\{r, R_2(K)\}$.*

Otherwise, we define $\mathcal{H}^C$ by considering each component of $C - B_G(B, 22K + 1)$ separately. For this, we first take a K_2 on vertices g, g' and set $V_{g'} := (B_G(B, 22K + 1) \cap V(C)) \cup N_G(C)$ and $V_g := N_G(C)$. Next, to ensure that $\mathcal{H}^C$ is $R_2(K)$-component-feasible, we extend this decomposition into the components D of $C - B_G(B, 22K+1)$ as follows. If some $\partial_G D$ has radius $\leqslant R_2(K) - 1$, then we may simply take a K_2 on vertices g^D, h^D, set $V_{g^D} := N_G(D) =: V_{h^D}$ and identify g^D with g'. Otherwise, if $\partial_G D$ can be partitioned into two sets N_0, N_1 of radius at most $R_1(K)$ that are at least

$5K+2$ far apart, then we obtain a partial decomposition $(H^C, \mathcal{V}^C)$ with support $Y^C \subseteq G[C,1]$ modelled on a graph with no K_4 minor by the next Lemma 12.4.3 (applied to $C := D$), and we can then identify the node of H^D given by Lemma 12.4.3 (i) with g'. This then yields the desired partial decomposition $\mathcal{H}^C$.

Lemma 12.4.3. *Let $K \in \mathbb{N}_{\geqslant 1}$, and let B be a ball in a graph G around a vertex w of radius $r \in \mathbb{N}$, and let C be a component of $G - B$. Suppose there are vertices $v_1, v_2 \in \partial_G C$ which are at least $2R_1(K) + 5K + 2$ apart and $\partial_G C \subseteq B_G(v_1, R_1(K)) \cup B_G(v_2, R_1(K))$.*

If G has no K-fat K_4 minor, then there is an induced subgraph Y of $G[C,1]$ with $\partial_G C \subseteq V(Y)$ which admits an honest decomposition $(H, \mathcal{V})$ modelled on some graph $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$ such that

(i) *$(H, \mathcal{V})$ contains $N_G(C)$ as a bag,*

(ii) $\mathrm{oradw}((H, \mathcal{V})) \leqslant R := \max\{r, R_1(K) + 2R_0(K) + 5K + 2\}$,

(iii) $\mathrm{irads}((H, \mathcal{V})) \leqslant 3$*, and*

(iv) *$(H, \mathcal{V})$ is R-component-feasible; moreover, for every component of D of $G - Y$ which meets, or equivalently is contained in, C, its neighbourhood $N_G(D)$ is contained in some bag V_h of $(H, \mathcal{V})$ with* $\mathrm{rad}(V_h) \leqslant R_1(K) + 2R_0(K) + 5K + 1$.

Let us now deduce Theorem 33' from Lemmas 12.4.2 and 12.4.3. For this, as noted earlier, it suffices to prove that K_4 has property $(*)$.

Lemma 12.4.4. *K_4 has property $(*)$ (with respect to $R(K) = R_2(K)$, $f_0'(K) := R_2(K) + 2R_0'(K) + 2$ and $f_1'(K) = f_1''(K) := 7$).*

Proof. Let $K \in \mathbb{N}_{\geqslant 1}$, and let G be a graph with no K-fat K_4 minor. Let B be a ball in G of radius $r \leqslant R_2(K)$ around a vertex $v \in V(G)$, and let C be a component of $G - B$. Further, let $\mathcal{D}$ be the set of components of $C - B_G(B, 22K+1)$, and first assume that there is some $D \in \mathcal{D}$ such that $\partial_G D$ contains three vertices that are pairwise at least $R_1(K)$ apart. Then we obtain an induced subgraph $Y \subseteq G[C,1]$ with $\partial_G C \subseteq V(Y)$ and the desired honest $R_2(K)$-component-feasible $(f_0'(K), 7)$-partial decomposition $(H, \mathcal{V})$ of Y with $N_G(C) = V_g$ for some $g \in V(H)$ by applying Lemma 12.4.2 to the ball B of radius $r \leqslant R_2(K)$ and the component C of $G - B$, which completes the proof in this case.

Otherwise, let $\mathcal{D}' \subseteq \mathcal{D}$ be the set of components D of $C - B_G(B, 22K+1)$ such that there are two vertices in $\partial_G D$ which are at least $2R_1(K) + 5K + 2$ apart. Then by assumption, every $D \in \mathcal{D}'$ satisfies the premise of Lemma 12.4.3; let $(H^D, \mathcal{V}^D)$ be the partial decomposition with support Y^D obtained from applying Lemma 12.4.3 to the ball $B_G(B, 22K+1)$, which is a ball of radius $r + 22K + 1$ around v, and the component D. Let g^D be a node of H^D with $V^D_{g^D} = N_G(D)$, which exists by (i). Further, for every component $D \in \mathcal{D} \setminus \mathcal{D}'$, let H^D be a K_2 on vertices g^D, h^D and set $V^D_{g^D} := V^D_{h^D} := N_G(D)$.

Now let H be the graph obtained from the disjoint union of the H^D for $D \in \mathcal{D}'$ by first identifying all g^D to a single node, which we call g', and adding a node g and the edge gg'. Note that $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$ since it is the 1-sum of the K_4-minor-free graphs H^D and K_2. Set $V_h := V_h^D$ for all $h \in V(H) \smallsetminus \{g, g'\}$ where D is the unique component in $\mathcal{D}'$ such that $h \in V(H^D)$ and $V_{g'} := (B_G(B, 22K+1) \cap V(C)) \cup N_G(C)$ and $V_g := N_G(C)$. Then, by construction, $(H, \mathcal{V})$ is a partial decomposition with support $Y := G[\bigcup V(Y^D) \cup V_{g'}]$ since $V^D_{g^D} \subseteq V_{g'}$ for all $D \in \mathcal{D}'$.

We claim that $(H, \mathcal{V})$ is as desired. Indeed, the outer-radial width of $(H, \mathcal{V})$ is $R_2(K)+22K+1 \leqslant f_0'(K)$, since $\mathrm{rad}(V_{h^D}), \mathrm{rad}_G(V_g) \leqslant \mathrm{rad}_G(V_{g'}) \leqslant \mathrm{rad}_G(B) + 22K + 1 \leqslant R_2(K) + 22K + 1$ for every $D \in \mathcal{D} \smallsetminus \mathcal{D}'$ and $\mathrm{rad}(V_h) \leqslant R_2(K) + 22K + 1$ for all other $h \in V(H)$ by (ii) of Lemma 12.4.3. Note that the radial spread of $(H, \mathcal{V})$ is $\leqslant 2 \cdot \max_{D \in \mathcal{D}} \mathrm{irads}((H^D, \mathcal{V}^D)) \leqslant 6 \leqslant f_1'(K)$, since $V_g \subseteq V_{g'}$, or it is 1 if $\mathcal{D} = \emptyset$.

By construction, it follows immediately from Lemma 12.4.3 (iv) that for every component of $G - Y$ which meets C there exists a bag V_h containing its neighbourhood. Moreover, clearly, every component of $G - Y$ which does not meet C has only neighbours in $N_G(C) = V_g \subseteq B$. All in all, $(H, \mathcal{V})$ is $R(K)$-component-feasible. □

Proof of Theorem 33'. By Theorem 12.2.5, Lemma 12.4.4 immediately yields Theorem 33'. □

Proof of Theorem 33. Let G be a graph with no K-fat K_4 minor. By Lemma 11.2.1, Theorem 33' yields that there exists a graph H with no K_4 minor which is $(50470K+142, 50470K+142)$-quasi-isometric to G. Hence, it follows from Lemma 9.2.1 that G is $(50470K+142, 3 \cdot (50470K+142)^2)$-quasi-isometric to H. Thus, Theorem 34 holds with $f(K) := (50470K+142, 3 \cdot (50470K+142)^2)$. □

The remainder of this section is devoted to the proofs of Lemmas 12.4.2 and 12.4.3. For this, we first show in Section 12.4.1 an auxiliary lemma that we will use to find K_4 as a K-fat minor in a graph G whenever we cannot find our desired decomposition of G modelled on a graph with no K_4 minor of small outer-radial width. Then, in Sections 12.4.2 and 12.4.3, we prove Lemmas 12.4.2 and 12.4.3, respectively.

12.4.1 Finding a K-fat K_4 minor

Lemma 12.4.5. *Let G be a graph, $K, r_1, r_2 \in \mathbb{N}_{\geqslant 1}$, and let v_1, v_2 be two vertices of G that are at least $r_1 + r_2 + 5K$ apart. Set $B_i := B_G(v_i, r_i)$ for $i \in [2]$, and suppose there are three B_1–B_2 paths P_1, P_2 and P_3 which are pairwise at least $5K$ apart.*

If G has no K-fat K_4 minor, then the $\mathring{P}_i$ lie in distinct components of $G - (B_1 \cup B_2)$.

Proof. Let us first note that since the paths P_1, P_2, P_3 are pairwise at least $5K$ apart, also $r_i \geqslant \lceil 5K/2 \rceil > 2K$ for $i \in [2]$. For $i \in [3]$, let $Q^i := q_0^i \dots q_{n_i}^i$ be a shortest $B_G(v_1, r_1 - 2K)$–$B_G(v_2, r_2 - 2K)$ path in $G[B_1] \cup P_i \cup G[B_2]$; in particular, $n_i := ||P_i|| + 4K$.

Suppose that at least two $\mathring{P}_i$ lie in the same component C of $G - (B_1 \cup B_2)$. We show that G has a K-fat K_4 minor. For this, assume without loss of generality that $P_1 \subseteq G[C,1]$. Let W be a P_1–$B_G(Q^2 \cup Q^3, K)$ path in C; by symmetry, we may assume that W ends in $B_G(Q^2, K)$. Further, let $W' = w_0 \dots w_m$ be a shortest P_1–Q^2 path in $G[Q^2, K] \cup W$. Note that W' ends in a vertex in $q^2_K Q^2 q^2_{n_2-K} = Q^2 \cap G[P_2, K]$, as $W \subseteq C \subseteq G - (B_1 \cup B_2)$. In particular, since Q^2 is a shortest $B_G(v_1, r_1 - 2K)$–$B_G(v_2, r_2 - 2K)$ path in $G[B_1 \cup V(P_2) \cup B_2]$ it follows that $w_{m-K} W' w_m \subseteq G[P_2, 2K]$.

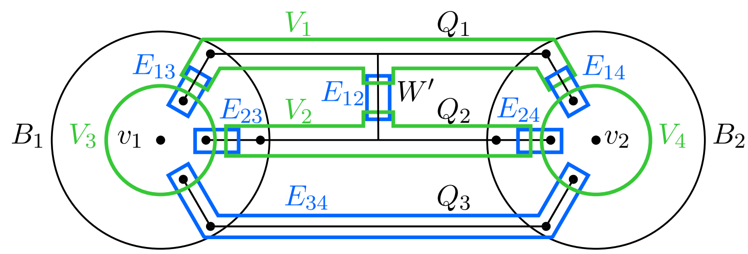

FIGURE 12.5: The K-fat K_4 minor in the proof of Lemma 12.4.5.

By defining the branch sets as

$$V_1 := V(q^1_K Q^1 q^1_{n_1-K}) \cup V(w_0 W' w_K),\ V_2 := V(q^2_K Q^2 q^2_{n_2-K} w_{m-K} W' w_m),$$
$$V_3 := B_G(v_1, r_1 - 2K) \text{ and } V_4 := B_G(v_2, r_2 - 2K).$$

and the branch paths as

$$E_{12} := w_K W' w_{m-K},\ E_{i3} := q^i_0 Q^i q^i_K,\ E_{i4} := q^i_{n_i-K} Q^i q^i_{n_i} \text{ and } E_{34} := Q^3$$

for $i \in [2]$ we obtain a K-fat K_4-model in G (see Figure 12.5). Indeed, the sets V_i and the paths E_{ij} obviously form a model of K_4. It remains to show that this K_4-model is K-fat.

By construction, we have $V_3, V(E_{i3}) \subseteq B_1$ and $V_4, V(E_{i4}) \subseteq B_2$ for $i \in [2]$, and thus by assumption

$$d_G(V_3, V_4), d_G(E_{i3}, E_{j4}), d_G(V_3, E_{i4}), d_G(V_4, E_{i3}) \geqslant d_G(B_1, B_2) \geqslant 5K \geqslant K$$

for $i, j \in [2]$. Moreover, $d_G(V(Q^1) \cup V_1, V(Q^2) \cup V_2) \geqslant d_G(P_1, P_2) - 2 \cdot 2K \geqslant 5K - 2 \cdot 2K = K$ since $V(Q^i) \cup V_i \subseteq B_G(P_i, 2K)$ for $i \in [2]$. This yields that

$$d_G(V_1, V_2), d_G(V_1, E_{2i}), d_G(V_2, E_{1i}), d_G(E_{1i}, E_{2i}) \geqslant K$$

for $i \in \{3,4\}$. As above, we have $d_G(Q^3, V(Q^i) \cup V_i) \geqslant K$ since also $Q^3 \subseteq G[P_3, 2K]$ and $d_G(P_3, P_i) \geqslant 5K$ for $i \in [2]$. Further, $d_G(Q^3, W) \geqslant K$ by the choice of W. Thus,

$$d_G(E_{34}, V_i), d_G(E_{34}, E_{ij}), d_G(E_{34}, E_{12}) \geqslant K$$

for $i \in [2]$ and $j \in \{3,4\}$. As C is a component of $G - (B_1 \cup B_2)$, we have $d_G(C, B_G(v_i, r_i - K)) = K + 1$ for $i \in [2]$. Since $E_{12} \subseteq W \subseteq C$ this implies that

$$d_G(E_{12}, E_{ij}), d_G(E_{12}, V_j) \geqslant d_G(C, B_G(v_i, r_i - K)) = K + 1 \geqslant K$$

for $i \in [2]$ and $j \in \{3,4\}$. Finally, we obtain

$$d_G(V_i, V_j) \geqslant d_G(B_G(C, K), B_G(v_{j-2}, r_{j-2} - 2K)) = 2K - K = K$$

for $i \in [2]$ and $j \in \{3,4\}$. This completes the proof. □

12.4.2 Proof of Lemma 12.4.3

Let us briefly sketch the proof of Lemma 12.4.3. For this, let us first recall its premises: Let G be a graph with no K-fat K_4 minor for $K \in \mathbb{N}_{\geqslant 1}$, let B be a ball in G, and let C be a component of $G - B$. Suppose that there are vertices $v_1, v_2 \in \partial_G C$ such that the sets $B_i := (v_i, R_1(K))$ for $i \in [2]$ induce a partition of $\partial_G C$ and such that $d_G(B_1, B_2) \geqslant 5K + 2$. We then construct for every component D of $C - (B_1 \cup B_2)$ a partial decomposition $(H^D, \mathcal{V}^D)$ of G with support $Y \subseteq G[D, 1]$ modelled on a graph $H^D \in \mathcal{H}_{SP}$ with terminals h_1^D, h_2^D such that $V_{h_i^D}^D = N_G(D) \cap B_i$ (Lemma 12.4.8). To obtain the desired partial graph-decomposition $(H, \mathcal{V})$ to satisfy the conclusion of Lemma 12.4.3, we then glue all the decompositions $(H^D, \mathcal{V}^D)$ together. More precisely, we obtain H from the disjoint union of the H^D by first identifying all the h_1^D to a node h_1 as well as all the h_2^D to a node h_2, and then adding a vertex g and the edges gh_1, gh_2. As H hence arises as the parallel composition of graphs in $\mathcal{H}_{SP}$, it is again in $\mathcal{H}_{SP}$ by Proposition 12.4.1 and hence has no K_4 minor. The desired partial graph-decomposition is then given by $(H, \mathcal{V})$ where $V_{h_i} := (B_i \cap V(C)) \cup (N_G(C) \cap B_G(B_i, 1))$, $V_g := N_G(C)$ and $V_h := V_h^D$ for all other $h \in V(H)$ where H^D is the unique graph containing h.

In the proof of Lemma 12.4.8, we construct the partial graph-decompositions $(H^D, \mathcal{V}^D)$ recursively via Lemma 12.4.7 below. For this, we will need the coarse version of Menger's theorem for two paths (Theorem 27 in Chapter 10), which we recall here for convenience.

Theorem 12.4.6. *Let G be a graph, $X, Y \subseteq V(G)$, and let Q be a shortest X–Y path in G. For all $d \in \mathbb{N}_{\geqslant 1}$, either there exist two disjoint $X - Y$ paths P_1, P_2 that are at least d apart or there*

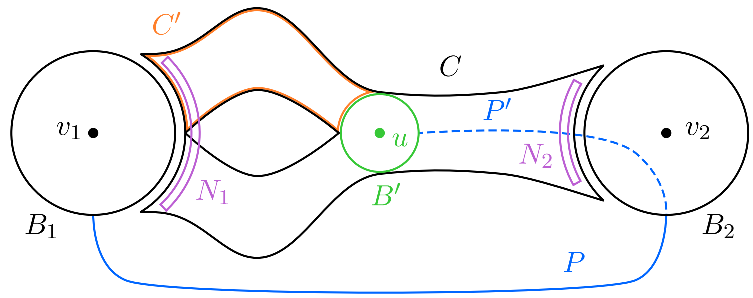

FIGURE 12.6: The situation in Lemma 12.4.7: A component C of $G - (B_1 \cup B_2)$ and the ball $B' \supseteq B$ around u. Depicted in blue is the path P and its dashed extension P', as required for the 'moreover' part applied to the orange component C'.

exists $z \in V(Q)$ such that $B_G(z, 129d)$ intersects every $X - Y$ path.[7]

Lemma 12.4.7. *Let G be a graph, $K, r_1, r_2 \in \mathbb{N}_{\geqslant 1}$, and let v_1, v_2 be two vertices of G that are at least $r_1 + r_2 + 5K + 2$ apart. Set $B_i := B_G(v_i, r_i)$ for $i \in [2]$. Let C be a component of $G - (B_1 \cup B_2)$ which attaches to B_1 and to B_2, and suppose there exists a B_1–B_2 path P in G such that $d_G(P, G[C, 1]) \geqslant 5K$. Set $N_i := N_G(C) \cap B_i$ for $i \in [2]$. If G has no K-fat K_4 minor, then there exists a vertex $u \in B_G(C, 1)$ such that every N_1–N_2 path through C intersects $B_G(u, 129 \cdot 5K)$.*

Moreover, for $B' := B_G(u, 130 \cdot 5K)$, $i \in [2]$ and for every component C' of $G - (B_i \cup B')$ that is contained in C and attaches to B_i and B', there is a B_i–B' path P' such that $d_G(P, G[C', 1]) \geqslant 5K$.

See Figure 12.6 for a depiction of the situation in Lemma 12.4.7.

Proof. Let Q be a shortest N_1–N_2 path through C. As $d_G(v_1, v_2) \geqslant r_1 + r_2 + 5K + 2$, the balls $B_G(B_1, \lceil (5/2)K \rceil)$ and $B_G(B_2, \lceil (5/2)K \rceil)$ are disjoint and not joined by an edge. Thus, Q is a shortest N_1–N_2 path in $G' := G[C, 1] \cup G[B_1 \cup B_2, \lceil (5/2)K \rceil]$. By applying Theorem 12.4.6 in G' to the sets N_1 and N_2 with $d = 5K$, we obtain either a vertex u of $Q \subseteq G[C, 1]$ such that every N_1–N_2 path in G' intersects $B_G(u, 129 \cdot 5K)$ or two N_1–N_2 paths Q_1 and Q_2 in G' such that $d_{G'}(Q_1, Q_2) \geqslant 5K$.

In the former case we are done (except for the 'moreover'-part) since u lies on a shortest N_1–N_2 path through C and $B' := B_G(u, 130 \cdot 5K) \supseteq B_G(u, 129 \cdot 5K)$; so suppose for a contradiction that the latter holds. Since $B_G(B_1, \lceil (5/2)K \rceil)$ and $B_G(B_2, \lceil (5/2)K \rceil)$ are disjoint and not joined by an edge, the internal vertices of any N_1–N_2 path in G' are in C; in particular, $\mathring{Q}_1, \mathring{Q}_2 \subseteq C$. Moreover, $d_G(Q_1, Q_2) \geqslant 5K$ because $d_{G'}(Q_1, Q_2) \geqslant 5K$ and any C-path in G (and hence any $\mathring{Q}_1$–$\mathring{Q}_2$ path) meeting $G - G'$ has length at least $2 \cdot (\lceil (5/2)K \rceil + 1) \geqslant 5K + 2$. Since also

[7] That the vertex z may be chosen on some shortest X–Y path is not stated in Theorem 27, but it follows easily from its proof. For convenience, we also remark that if z is any vertex such that $B_G(z, 129d)$ intersects every X–Y path, then $B_G(z, 129d)$ in particular intersects Q in some vertex z'. Then also $B_G(z', 2 \cdot 129d)$ intersects every X–Y path, so the assertion follows directly from Theorem 27 if we increase the radius of the ball by a factor of 2.

$d_G(Q_1 \cup Q_2, P) \geqslant d_G(G[C,1], P) \geqslant 5K$ by assumption, the paths Q_1, Q_2, P are pairwise at least $5K$ apart. Thus, as $\mathring{Q}_1$ and $\mathring{Q}_2$ lie in the same component C of $G-(B_1 \cup B_2)$, applying Lemma 12.4.5 yields that G has a K-fat K_4 minor, which is a contradiction.

It remains to show the 'moreover'-part. For this, let $i \in [2]$ and let C' be a component of $G-(B_i \cup B')$ which attaches to both B_i and B'. If B' does not meet B_{3-i}, we extend the path P to a B_i–B' path P' by adding first a P–C path P_1 through B_{3-i} and then a P_1–B' path P_2 in C. Otherwise, we extend P by a P–B' path P_1 through B_{3-i}. We claim that $d_G(P', G[C',1]) \geqslant 5K$. For this, let S be a shortest $G[C',1]$–P' path in G. We show that S has length at least $5K$. Then $d_G(G[C',1], P') \geqslant ||S|| \geqslant 5K$, which yields the claim. Since $d_G(G[C',1], P) \geqslant d_G(G[C,1], P) \geqslant 5K$ as $C' \subseteq C$, we are done if S ends in P. Hence, we may assume that S ends in P_1 or P_2.

We distinguish two cases. First, assume that S meets $B_G(w, 129 \cdot 5K) \subseteq B'$. Since $P_1 \cup P_2$ does not meet $B' = B_G(B_G(w, 129 \cdot 5K), 5K)$ except in its last vertex, $P_1 \cup P_2$ has distance at least $5K$ from $B_G(w, 129 \cdot 5K)$. Hence, as S meets both $P_1 \cup P_2$ and $B_G(w, 129 \cdot 5K)$ by assumption, it has length at least $5K$.

Now assume that S avoids $B_G(w, 129 \cdot 5K)$. Note first that also C' and $P_1 \cup P_2$ avoid $B_G(w, 129 \cdot 5K)$. Moreover, $C', P_1 \cup P_2 \subseteq G[C,1] \cup G[B_1 \cup B_2]$. Since $B_G(w, 129 \cdot 5K)$ separates B_1 and B_2 in $G' \supseteq G[C,1] \cup G[B_1 \cup B_2]$, and $C' \subseteq C$ attaches to B_i while $P_1 \cup P_2$ meets B_{3-i}, the two subgraphs C' and $P_1 \cup P_2$ lie in distinct components of $(G[C,1] \cup G[B_1 \cup B_2]) - B_G(w, 129 \cdot 5K)$. Since S avoids $B_G(w, 129 \cdot 5K)$ and C is a component of $G-(B_1 \cup B_2)$, it follows that S leaves $G[C,1] \cup G[B_1 \cup B_2]$ through B_i and enters it again through B_{3-i}. Hence, S contains a B_1–B_2 path. Since $d_G(v_1, v_2) \geqslant 2r_1 + r_2 + 5K + 2$, we have $d_G(B_1, B_2) \geqslant 5K + 2$, and thus S has length at least $5K$. □

Lemma 12.4.8. *Let G be a graph, and $K, r_1, r_2 \in \mathbb{N}_{\geqslant 1}$ such that $r_1 \leqslant r_2$. Let v_1, v_2 be two vertices of G that are at least $r_1 + r_2 + 5K + 2$ apart. Set $B_i := B_G(v_i, r_i)$ for $i \in [2]$, let C be a component of $G-(B_1 \cup B_2)$ which attaches to B_1 and B_2, and suppose there exists a B_1–B_2 path P in G such that $d_G(P, G[C,1]) \geqslant 5K$.*

If G has no K-fat K_4 minor, then there exists an induced subgraph Y of $G[C,1]$ which admits an honest decomposition $(H, \mathcal{V})$ modelled on a graph $H \in \mathcal{H}_{SP}$ with terminals h_1, h_2 such that

(a) $V_{h_i} = N_G(C) \cap B_i$ *for $i \in [2]$,*

(b) $\mathrm{rad}_G(V_h) \leqslant R_0'(K)$ *for all $h \in V(H)$ with $h \notin B_H(\{h_1, h_2\}, 1)$,*

(c) $d_G(v_i, v) \leqslant r_i + 2 \cdot R_0(K) + 5K + 1$ *for $i \in [2]$ and all $v \in V_h$ with $h \in B_H(h_i, 1)$,*

(d) $\mathrm{irads}(\mathcal{H}) \leqslant 3$*; moreover, $d_H(h_i, h) \leqslant 3$ for all $h \in V(H_v)$ and $v \in N_G(C) \cap B_i$ for $i \in [2]$, and*

(e) *for every component C' of $G - Y$ which meets, or equivalently is contained in, C, its neighbourhood $N_G(C')$ is contained in some bag V_h of $(H, \mathcal{V})$.*

We remark that the proof of Lemma 12.4.8 is the only point in this chapter where it actually makes a difference that we not only consider finite but also infinite graphs. In fact, for finite graphs, a much simpler inductive argument yields the desired graph-decomposition $(H, \mathcal{V})$:

Proof sketch of Lemma 12.4.8 for finite graphs. We may apply Lemma 12.4.7 to C, v_1, v_2 and K, r_1, r_2, which yields a vertex $u \in B_G(C, 1)$. We now apply induction to the (now strictly smaller since C is finite) components D of $C - (B_1 \cup B')$ and $C - (B_2 \cup B')$ for $B' := B_G(u, 130 \cdot 5K)$ that attach to both balls, to obtain graph-decompositions $(H^D, \mathcal{V}^D)$ with terminals h_1, u or u, h_2 of H^D, respectively. We then obtain the desired graph-decomposition $(H, \mathcal{V})$ by gluing the $(H^D, \mathcal{V}^D)$ together in u and h_1 or h_2.

Proof of Lemma 12.4.8 for arbitrary graphs. We define the desired graph-decomposition $(H, \mathcal{V})$ recursively (see Figure 12.7 for a sketch). Initialise H^0 as a K_2 on vertices $h_1 := v_1$ and $h_2 := v_2$. In particular, $H^0 \in \mathcal{H}_{SP}$ with terminals h_1 and h_2. Set $V_{h_i} := N_G(C) \cap B_i$ and $r_{h_i} := r_i$ for $i \in [2]$, and note that $N_G(C) \subseteq V_{h_1} \cup V_{h_2}$. Set further $\mathcal{C}^0 := \{C\}$, $w_1^C := h_1$ and $w_2^C := h_2$. We warn the reader that in the following the vertices of the H^i are vertices of G, i.e. $V(H^i) \subseteq V(G)$. Hence, we may divert from our usual convention, and denote the vertices of H by u, w.

Now suppose that for $n \in \mathbb{N}$ we have constructed graphs $H^0 \subseteq \ldots \subseteq H^n \in \mathcal{H}_{SP}$ with terminals h_1 and h_2, and sets $V_u \subseteq B_G(C, 1)$ for all $u \in V(H^n)$. Assume further that, for every $i \leqslant n$, we fixed

(I) the collection $\mathcal{C}^i$ of all components D of $C - \bigcup_{u \in H^i} V_u$ whose neighbourhood is not contained in a V_u for some $u \in V(H^i)$,

(II) for every $D \in \mathcal{C}^i$, an edge $w_1^D w_2^D$ of H^i,

(III) for every $D \in \mathcal{C}^{i-1} \setminus \mathcal{C}^i$, some $u_D \in V(H^i) \setminus V(H^{i-1})$ such that $w_1^D u_D, u_D w_2^D \in E(H^i)$, if $i \geqslant 1$,

(IV) for every $u \in V(H^i) \setminus V(H^{i-1})$, some $C_u \in \mathcal{C}^{i-1}$ with $u \in B_G(C_u, 1)$ and $u_{C_u} = u$, if $i \geqslant 1$, and

(V) for all $u \in V(H^i) \setminus \{h_1, h_2\}$, the bag $V_u := B_G(u, R_0(K)) \cap B_G(C_u, 1)$ and $r_u := R_0(K)$,

such that

(i) for every edge e of $G[\bigcup_{u \in H^i} V_u]$ that is not an edge of $\bigcup_{u \in H^i} G[V_u]$ there is an edge $uu' \in E(H^i)$ such that $e \in G[V_u \cup V_{u'}]$;

(ii) every $D \in \mathcal{C}^i$ is a component of $G - (V_{w_1^D} \cup V_{w_2^D})$;

(iii) for every $u \in V(H^i) \setminus V(H^{i-1})$ every $V_{w_1^{C_u}}$–$V_{w_2^{C_u}}$ path through C_u intersects $B_G(u, 129 \cdot 5K)$;

(iv) if $i \geqslant 1$, then for every $D \in \mathcal{C}^{i-1}$ we have $D \notin \mathcal{C}^i$ if and only if $d_G(w_1^D, w_2^D) \geqslant r_{w_1^D} + r_{w_2^D} + 5K + 2$;

(v) for every $D \in \mathcal{C}^i \setminus \mathcal{C}^{i-1}$ and the unique component $D' \in \mathcal{C}^{i-1} \setminus \mathcal{C}^i$ with $D \subseteq D'$ we have $w_1^D = w_1^{D'}$, $w_2^D = u_{D'}$ or $w_1^D = u_{D'}$, $w_2^D = w_2^{D'}$;

(vi) for every $D \in \mathcal{C}^i$ there exists a $B_G(w_1^D, r_{w_1^D})$–$B_G(w_2^D, r_{w_2^D})$ path P with $d_G(P, G[D,1]) \geqslant 5K$.

In the following, we will refer to components $D \in \mathcal{C}^n$ as in (iv), that is those which satisfy $d_G(w_1^D, w_2^D) \geqslant r_{w_1^D} + r_{w_2^D} + 5K + 2$, as *long*.

We now obtain H^{n+1} as follows. For every long component $D \in \mathcal{C}^n$, let $u_D \in B_G(D,1)$ be a vertex obtained from the application of Lemma 12.4.7 to D, w_1^D, w_2^D and $r_{w_1^D}, r_{w_2^D}$, which is possible due to (vi). Then H^{n+1} is obtained from H^n by adding for every long component $D \in \mathcal{C}^n$ the vertex u_D and the edges $w_1^D u_D$ and $u_D w_2^D$.

Let the collection $\mathcal{C}^{n+1}$, the sets V_{u_D} and the integer r_{u_D} be as required by (I) and (V), respectively. The construction obviously ensures (III) and (iv). For (IV), we set $C_{u_D} := D$ for every $u_D \in V(H^{n+1}) \setminus V(H^n)$, which also yields (iii). For (II), let $D' \in \mathcal{C}^{n+1} \setminus \mathcal{C}^n$. Then D' is contained in a long component $D \in \mathcal{C}^n$. As the neighbourhood of D' is not contained in a V_u for $u \in V(H^{n+1})$, the ball $B_G(u_D, 129 \cdot 5K)$ meets every $V_{w_1^D}$–$V_{w_2^D}$ path through D and $B_G(u_D, 129 \cdot 5K) \cap V(D) \subseteq V_{u_D}$, the component $D' \subseteq D$ attaches to V_{u_D} and precisely one of‘$V_{w_1^D}$ and $V_{w_2^D}$, say $V_{w_1^D}$. Then we set $w_1^{D'} := w_1^D$ and $w_2^{D'} := u_D$, which immediately ensures (ii) and (v). The moreover-part of Lemma 12.4.7 yields (vi). For (i), we first note that every edge of $G[\bigcup_{u \in H^{n+1}} V_u]$ that has no endvertex in $V_{u_D} \cap V(D)$ for any $D \in \mathcal{C}^n$ was already an edge of $G[\bigcup_{u \in H^n} V_u]$, and hence satisfies (i) by (i) of $(H^n, \mathcal{V}^n)$. Now let e be any edge of $G[\bigcup_{u \in H^{n+1}} V_u]$ which is not in $G[\bigcup_{u \in H^n} V_u]$. Then there is $D \in \mathcal{C}^n$ such that $e = xy$ has one endvertex, say x, in $V_{u_D} \cap V(D)$, so e has its other endvertex y in $G[D,1]$. Hence, by (ii) of $(H^n, \mathcal{V}^n)$, the vertex y is in $V_{w_1^D} \cup V_{w_2^D}$. By construction, $w_1^D u_D, u_D w_2^D \in E(H^{n+1})$, and hence e satisfies (i). Also H^{n+1} is still in $\mathcal{H}_{SP}$ with terminals h_1, h_2 by (II) and Proposition 12.4.1. This completes the verification that H^{n+1} is as desired.

Let $H' := \bigcup_{i \in \mathbb{N}} H^i$ be the limit of the H^i. By Proposition 12.4.1, we have $H' \in \mathcal{H}_{SP}$ with terminals h_1, h_2. Now let H be obtained from H' by first deleting all edges uu' of H' with $d_G(u, u') \geqslant r_u + r_{u'} + 5K + 2$ and then subdividing each remaining edge precisely once. By Proposition 12.4.1, $H \in \mathcal{H}_{SP}$ with terminals h_1 and h_2. We claim that $(H, \mathcal{V})$ with the V_h already defined for all $u \in V(H')$ and $V_h := V_u \cup V_w$ for all new subdivision vertices h on an edge uw of H' is a graph-decomposition of $Y := G[\bigcup_{u \in H} V_u] \subseteq G[C,1]$ as desired.

We first show that $(H, \mathcal{V})$ is a graph-decomposition of Y. Indeed, (H1) holds, as the parts $G[V_u]$ cover Y by the definition of Y and the edges relevant for (i) have not been deleted. Also $(H, \mathcal{V})$ satisfies (H2): Let v be in $V(Y)$, and let n be minimal such that $v \in \bigcup_{h \in H^n} V_h$. Assume that $n \geqslant 1$. Then $v \in V_{h_v}$ for some $h_v \in V(H^n) \setminus V(H^{n-1})$. By (V) and the minimality of n, the vertex v lies in $C_{h_v} \in \mathcal{C}^{n-1}$. Thus, the choice of h_v is unique by (IV) and (V). If $n = 0$, then we have $v \in N_G(C) \cap B_i$ for a unique $i \in [2]$ and $h_v = h_i$. We first prove that $H'_v = H'[\{h \in V(H') \mid v \in V_h\}]$

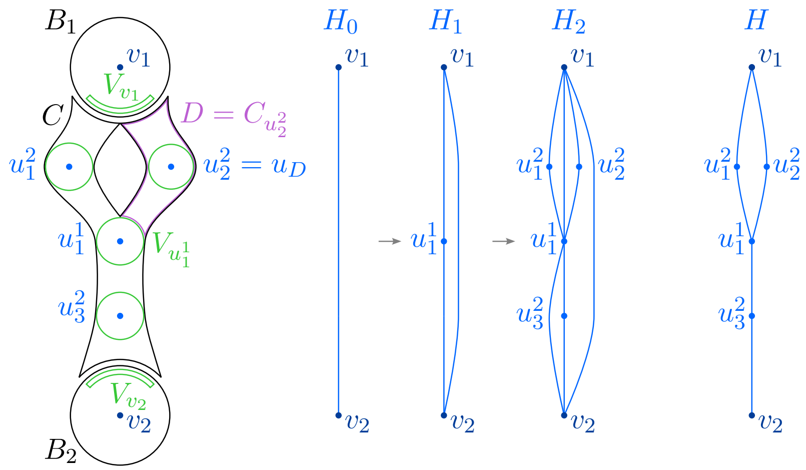

FIGURE 12.7: Depicted are the graphs H^0, H^1, H^2, H in blue and the bags V_u for the nodes u of H^2 for the component C of $G - (B_1 \cup B_2)$ in green. In this example, we have $w_1^D = v_1, w_2^D = u_1^1$ for the component $D = C_{u_2^2}$ in purple, and $H' = H^2$.

is contained in $H'[h_v, 1]$.

For this, let $h' \in V(H') \smallsetminus \{h_v\}$ with $v \in V_{h'}$ be given, and let $m \in \mathbb{N}$ be minimal such that $C_{h'} \in \mathcal{C}^m$. Since n is minimal and h_v is unique, we have that $m > n$ and that $C_{h'}$ attaches to v. For every ℓ with $m \geqslant \ell > n$, let C^ℓ be the unique component in $\mathcal{C}^\ell$ which contains $C_{h'}$; in particular, $C_{h'} = C^m$. Since $C_{h'} \in \mathcal{C}^m \smallsetminus \mathcal{C}^{m-1}$, it follows from (I), (iv) and (V) that no C^ℓ is in $\mathcal{C}^{\ell-1}$. Also, every C^ℓ with $\ell > n$ attaches to v as $v \in V_{h_v}$ and $C^m = C_{h'}$ attaches to v. By iteratively combining this with (ii) and (v), we obtain that $h_v = w_2^{C^{n+1}} = \cdots = w_2^{C^m}$ (after possibly swapping the names of $w_1^{C^\ell}$ and $w_2^{C^\ell}$). Moreover, by (IV), $h' = u_{C_{h'}} = u_{C^m}$, and thus $h_v h' = w_2^{C^m} u_{C^m}$ is an edge of H' by (III), which completes the proof that $H'_v \subseteq H'[h_v, 1]$.

Then (H2) of $(H, \mathcal{V})$ follows: For every node h' of H' with $v \in V_{h'}$, since $d_G(h_v, h') \leqslant r_{h_v} + r_{h'}$ as V_{h_v} and $V_{h'}$ intersect in v, we did not delete the edge $h_v h'$ in the construction of H. Since H is obtained from H' by subdividing each remaining edge precisely once, and since we associated to every subdivision vertex u of H the union V_u of the two bags associated with its neighbours, $H'_v \subseteq H'[h_v, 1]$ implies that $H_v = H[\{h \in V(H) \mid v \in V_h\}]$ is contained in $H[h_v, 3]$.[8] In particular, H_v is connected, which yields (H2). Moreover, $\mathrm{irads}((H, \mathcal{V})) \leqslant 3$ and also (d) follows immediately.

We now check the other required properties. By the definition of V_{h_1}, V_{h_2} at the beginning, $(H, \mathcal{V})$ satisfies (a). We claim that (b) and (c) hold. Indeed, for every $u \in V(H')$ we have $d_G(v, u) \leqslant r_u$ for all $v \in V_u$, where $r_u = R_0(K)$ for $u \neq h_1, h_2$ and $r_u = r_i$ for $u = h_i$. Now let $u \in V(H)$ be a subdivision vertex of an edge ww' of H', say with $r_w \geqslant r_{w'}$. Then $d_G(v, w) \leqslant r_w + 2r_{w'} + 5K + 1$,

[8] We remark that H_v is not necessarily contained in $H[h_v, 2]$: for any given $h' \in H'[h_v, 1]$ and any edge $e \neq h'h_v$, the bag corresponding to the subdivision vertex of e in H also contains v but has distance 3 to h_v.

which yields (b) and (c).

It thus remains to verify (e). For this, suppose towards a contradiction that there is a component C' of $G - Y$ which meets C and whose neighbourhood is not contained in some V_h. In particular, C' is a component of $C - Y$. Now first suppose that there are no two vertices in H' whose union of their bags contains $N_G(C')$. Then there are three vertices $u_1, u_2, u_3 \in N_G(C)$ and nodes g_1, g_2, g_3 such that V_{g_i} contains u_i but no other u_j. Pick some n such that all g_i are contained in $V(H^n)$. Then the component of $C - \bigcup_{h \in H^n} V_h$ that contains C' contradicts that H^n satisfied (ii). Thus, there are nodes $g_1, g_2 \in V(H')$ such that $N_G(C') \subseteq V_{g_1} \cup V_{g_2}$. In particular, C' is a component of $C - \bigcup_{h \in H^n} V_h$ for all $n \in \mathbb{N}$ with $g_1, g_2 \in V(H^n)$. Thus, $C' \in \mathcal{C}^n$ for all such n, as its neighbourhood is not even contained in some V_h with $h \in V(H')$. Let m be the minimal such n. By (II) and (ii), C' is a component of $G - (V_g \cup V_{g'})$ and $gg' = w_1^{C'} w_2^{C'}$ is an edge of H^n. If we did not delete gg' when constructing H, then we have a contradiction, as the bag V_u associated with the subdivision vertex u of H on the edge gg' then contains the entire neighbourhood of C'. Otherwise, we have $d_G(g, g') > r_g + r_{g'} + 5K + 2$, which by (iv) contradicts that $C' \in \mathcal{C}^{m+1}$. □

Proof of Lemma 12.4.3. Set $B_i := B_G(v_i, R_1(K))$ for $i \in [2]$. Further, let P_i be a shortest w–B_i path and set $W := P_1 \cup P_2$. Then W is a B_1–B_2 walk, and thus contains a B_1–B_2 path P. By construction, we have $V(P) \subseteq V(W) \subseteq B_G(w, r - R_1(K) + 1)$ and hence $d_G(P, G[C, 1]) \geqslant R_1(K) - 1 \geqslant 5K$. Since also $d_G(v_1, v_2) \geqslant 2R_1(K) + 5K + 1$ by assumption, we may apply Lemma 12.4.8 to every component C' of $G - (B_1 \cup B_2)$ that has a neighbour in both B_1 and B_2 and that meets, or equivalently is contained in, C. Since we are done if G has a K-fat K_4 minor, we may assume that Lemma 12.4.8 yields for every such C' a partial graph-decomposition $\mathcal{H}^{C'} = (H^{C'}, \mathcal{V}^{C'})$ of G with support $Y^{C'} \subseteq G[C', 1] \subseteq G[C, 1]$ that satisfies (a) to (e).

We now define the desired decomposition $(H, \mathcal{V})$ of $Y := G[(\bigcup V(Y^{C'})) \cup ((B_1 \cup B_2) \cap V(C)) \cup N_G(C)]$. The graph H is obtained from the disjoint union of the $H^{C'}$ and three new vertices h_1, h_2 and g by identifying all $h_1^{C'}$ with h_1 as well as all $h_2^{C'}$ with h_2 and adding the edges $h_1 g$ and $h_2 g$. Note that this is a parallel composition of graphs $H^{C'} \in \mathcal{H}_{SP}$ and a path of length 2; so H is again in $\mathcal{H}_{SP}$ with terminals h_1 and h_2 by Proposition 12.4.1; in particular, $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$. Then $(H, \mathcal{V})$ with $V_g := N_G(C)$, $V_{h_i} := (B_i \cap V(C)) \cup (N_G(C) \cap B_G(B_i, 1))$ for $i \in [2]$ and $V_h := V_h^{C'}$ for all other $h \in V(H)$ where C' is the unique component with $h \in H^{C'}$ is a decomposition of Y. $(H, \mathcal{V})$ obviously satisfies (H1), as the definition of the V_{h_i} ensures that the $G[V_{h_i}]$ cover the edges from $\partial_G C$ to $N_G(C)$. Moreover, (H2) holds for $(H, \mathcal{V})$, since the V_{h_i} are disjoint from all components C' of $G - (B_1 \cup B_2)$ which meet C.

Since $\partial_G C \subseteq B_1 \cup B_2$, we have $\partial_G C \subseteq V_{h_1} \cup V_{h_2} \subseteq V(Y)$. By construction, $(H, \mathcal{V})$ clearly satisfies (i) for $g \in V(H)$. Moreover, by (b) and (c) of the $\mathcal{H}^{C'}$ and the definition of V_{h_i}, we have $\mathrm{rad}(V_h) \leqslant R_1(K) + 2R_0(K) + 5K + 2$ for all $h \neq g \in V(H)$, and thus $(H, \mathcal{V})$ satisfies (ii). By construction, (iii) follows from property (d) of the $\mathcal{H}^{C'}$.

Lastly, we show that $(H, \mathcal{V})$ satisfies (iv). Every component of $G - Y$ which does not meet C is contained in $N_G(C) = V_g$. Now let D be an arbitrary component of $G - Y$ which meets, or equivalently is contained in, C. Let C' be the component of $G - (B_1 \cup B_2)$ which contains D. If C' attaches only to B_1 or only to B_2, then $N_G(D)$ is contained in V_{h_1} or V_{h_2}, respectively. So we may assume that C' attaches to both B_1 and B_2. Then $N_G(D)$ is contained in $V_h^{C'}$ for some $h \in V(H^D)$ by (e). Thus, by construction of $(H, \mathcal{V})$ we have that $N_G(D) \subseteq V_h$ for some $h \in V(H) \smallsetminus \{g\}$, and hence $\mathrm{rad}_G(V_h) \leqslant R_1(K) + 2R_0(K) + 5K + 2$, as we have shown for (ii) above. Hence, $(H, \mathcal{V})$ is as desired. □

12.4.3 Proof of Lemma 12.4.2

Let us briefly sketch the proof of Lemma 12.4.2. For this, let us first recall its premises: Let G be a graph with no K-fat K_4 minor for $K \in \mathbb{N}_{\geqslant 1}$, let B be some ball in G around a vertex w of radius $r \in \mathbb{N}$, and let C be a component of $G - B$. Suppose that C^* is a component of $C - B_G(B, 22K+1)$ such that there are three vertices in $\partial_G C^*$ that are pairwise far apart in G. We first find a vertex $w' \in V(C^*)$ such that there are at least three $B_G(w', 22K)$–$B_G(w, r - \ell(K))$ paths that are pairwise at least $11K$ apart (Corollary 12.4.10). These paths then have to lie in distinct components of $G' := G - (B_1(w', 22K) \cup B_G(w, r - \ell(K)))$ by Lemma 12.4.5. In particular, we can conclude that, for every component of G' that attaches to both balls, one of those three paths is still at least $5K$ away. We then apply Lemma 12.4.8 to $B_1 := B_G(w', 22K)$ and $B_2' := B_G(w, r - \ell(K))$ to obtain, for every component D of G' that attaches to B_1 and B_2', a partial decomposition modelled on a graph $H^D \in \mathcal{H}_{SP}$ with terminals h_1^D, h_2^D. We modify these decompositions by contracting all nodes of H^D whose bags contain vertices from $B_G(B, 1)$ to a single vertex h_2, which ensures that we may enlarge the bag associated with h_2 so that it contains B (Lemma 12.4.11). We then glue these modified decompositions together by identifying the h_1^D and the h_2^D. In the end, we adopt the decomposition so that it is R-component-feasible and obtain the desired partial decomposition of $G[C, 1]$ as its restriction to $G[C, 1]$.

We first show that there exists a vertex w' as described above.

Lemma 12.4.9. *Let $\ell, d, r \in \mathbb{N}_{\geqslant 1}$ with $\ell \geqslant 6d$, and let B be a ball in a graph G around a vertex w of radius r. Assume that C is a component of $G - B$ whose boundary $\partial_G C$ contains three vertices that are pairwise at least $4 \cdot (2\ell + d + 2)$ apart. Then there exist a vertex $w' \in V(C)$ and three $B_G(w, r - \ell)$–$B_G(w', 2d)$ paths that are pairwise at least d apart.*

Proof. Let $u_1, u_2, u_3 \in \partial_G C$ be three vertices that are pairwise at least $4 \cdot (2\ell + d + 2)$ apart. Let U_i, for $i \in [3]$, be the set of vertices $v \in V(C)$ such that every shortest v–B path meets $\partial_G C$ in a vertex

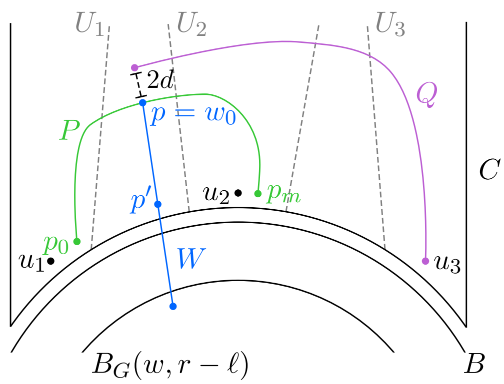

FIGURE 12.8: Setup in the beginning of the proof of Lemma 12.4.9

of distance less than $2\ell + d + 2$ to u_i. Note that the sets U_i are pairwise disjoint, as a shortest v–B path meets $\partial_G C$ in precisely one vertex and this vertex has distance less than $2\ell + d + 2$ to at most one u_i, since the pairwise distance of the u_i is at least $2 \cdot (2\ell + d + 2)$.

Let $P' = p'_0 \dots p'_n$ be a U_1–$(U_2 \cup U_3)$ path in C; by symmetry, we may assume that P' ends in U_2. We extend P' to a $(\partial_G C \cap U_1)$–$(\partial_G C \cap U_2)$ path $P = p_0 \dots p_m$ by adding a shortest $(\partial_G C \cap U_1)$–p'_0 path and a shortest p'_n–$(\partial_G C \cap U_2)$ path. Note that these paths are contained in U_1 and U_2, respectively, and hence are disjoint. Thus P is indeed a path. Further, let Q be a u_3–$B_G(P, 2d)$ path in C and let p be some vertex of P such that Q ends in $B_G(p, 2d)$. Then $p \notin U_3$ by the assumption on P. Let $W = w_0 \dots w_k$ be a shortest p–$B_G(w, r - \ell)$ path in G, and let p' be the unique vertex in $W \cap \partial_G C$ (see Figure 12.8). Since $p \notin U_3$, we may choose W so that $d_G(p', u_3) \geqslant 2\ell + d + 2$. Moreover, as $d_G(u_1, u_2) \geqslant 4 \cdot (2\ell + d + 2)$, we cannot have $d_G(p', u_1), d_G(p', u_2) < 2 \cdot (2\ell + d + 2)$; by symmetry, we may thus assume $d_G(p', u_2) \geqslant 2 \cdot (2\ell + d + 2)$. In particular, this implies that

$$d_G(p_m, p') \geqslant d_G(p', u_2) - d_G(u_2, p_m) \geqslant 2 \cdot (2\ell + d + 2) - (2\ell + d + 2) = 2\ell + d + 2. \tag{12.4}$$

Let x_3 be the first vertex on Q such that $x_3 \in B_G(W, 2d)$, and let x_2 be the last vertex on P such that $x_2 \in B_G(W, 2d)$. Further, let $i_j \in \{0, \dots, k\}$, for $j \in \{2, 3\}$ be an index such that $x_j \in B_G(w_{i_j}, 2d)$. As the following reasoning will be symmetric in P and Q, we may assume without loss of generality that $i_3 \leqslant i_2$. We claim that $w' := w_{i_2}$ is as desired (see Figure 12.9a).

For this, let Q_2 be a shortest p_m–$B_G(w, r - \ell)$ path and Q_3 a shortest u_3–$B_G(w, r - \ell)$ path. Further, let S_3 be a shortest x_3–w_{i_3} path and let s_3 be the first vertex on the path $S'_3 := x_3 S_3 w_{i_3} W w_{i_2}$ that is contained in $B_G(w_{i_2}, 2d)$ (see Figure 12.9a). We now define our final three $B_G(w, r - \ell)$–$B_G(w', 2d)$ paths P_1, P_2, P_3 (see Figure 12.9b):

$$P_1 := w_{i_2 + 2d} W w_k, \; P_2 := x_2 P p_m Q_2 \text{ and } P_3 := s_3 \bar{S}'_3 x_3 \bar{Q} u_3 Q_3.$$

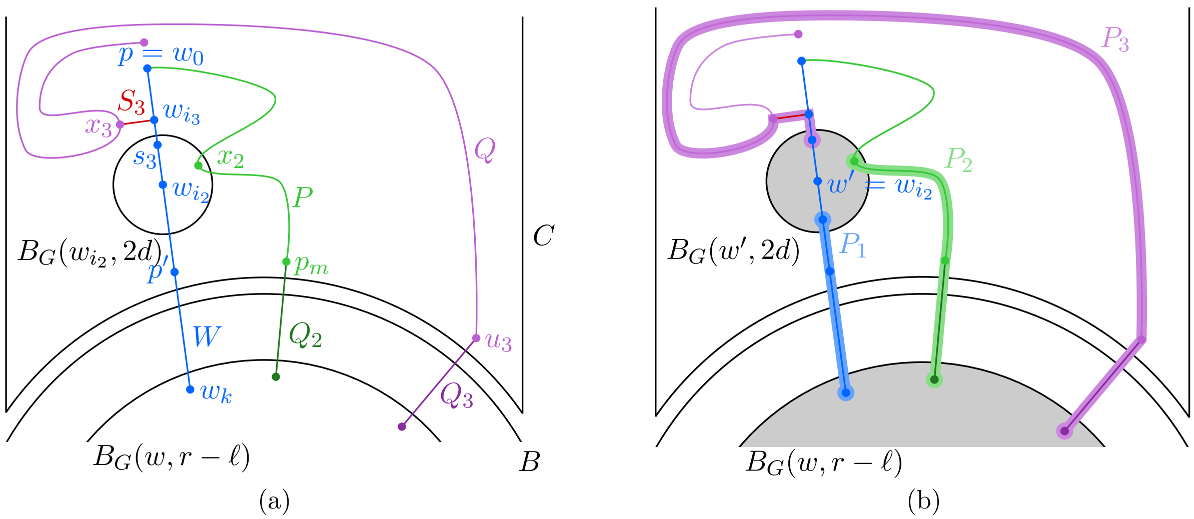

FIGURE 12.9: Choice of w' and construction of the paths P_1, P_2, P_3 in the proof of Lemma 12.4.9.

We claim that the P_i are as desired, i.e. $d_G(P_i, P_j) \geqslant d$ for every two distinct $i, j \in [3]$. First, we show $d_G(P_1, P_2) \geqslant d$: By the choice of x_2, we have $d_G(W, x_2 P p_m) \geqslant 2d \geqslant d$. Further, since Q_2 starts in p_m and both shortest $\partial_G C$–$B_G(w, r-\ell)$ paths Q_2 and $p' W w_k$ have length $\ell + 1$, we have

$$d_G(p' W w_k, Q_2) \geqslant d_G(p_m, p') - ||p' W w_k|| - ||Q_2|| \geqslant (2\ell + d + 2) - 2(\ell + 1) = d. \qquad (12.5)$$

Since $d_G(p_m, p') \geqslant 2\ell + d$ by (12.4), $\ell \geqslant d$, Q_2 is a shortest $\partial_G C$–$B_G(w, r-\ell)$ path and $p' \in \partial_G C$, we have $d_G(Q_2, B_G(p', d) \cap V(W)) \geqslant d$. As W is a shortest p–$B_G(w, r-\ell)$ path, the remainder of W in C has distance $\geqslant d$ to B and hence to Q_2. All in all, $d_G(W, Q_2) \geqslant d$, and hence $d_G(P_1, P_2) \geqslant d$.

Next, we show $d_G(P_1, P_3)$: Analogously to $d_G(W, Q_2) \geqslant d$, we also find $d_G(W, Q_3) \geqslant d$. Moreover, $d_G(W, x_3 \bar{Q} u_3) \geqslant 2d \geqslant d$ by the choice of x_3. Thus, to see that $d_G(P_1, P_3) \geqslant d$, it suffices to verify $d_G(w_{i_2+2d} W w_k, S_3') \geqslant d$: Since W is a shortest path in G and $i_3 \leqslant i_2$, we have $d_G(w_{i_2+2d} W w_k, w_{i_3} W w_{i_2}) = d_G(w_{i_2+2d}, w_{i_2}) = 2d$. Moreover, by the choice of $x_3 \in V(Q)$ and S_3, the x_3–w_{i_3} path S_3 has length $||S_3|| = d_G(W, x_3) = 2d$. Thus, the distance from the first half of S_3 to $w_{i_2+2d} W w_k$ is $\geqslant d_G(x_3, W) - ||S_3||/2 = 2d - d = d$, and the distance from the second half of S_3 to $w_{i_2+2d} W w_k$ is $\geqslant d_G(w_{i_3}, w_{i_2+2d} W w_k) - ||S_3||/2 \geqslant 2d - d = d$.

Finally, we show $d_G(P_2, P_3) \geqslant d$: By the choice of Q, we have $d_G(P, Q) \geqslant 2d$. Similarly to the previous argument, we find $d_G(x_2 P p_m, S_3') \geqslant d$, as $d_G(W, x_2 P p_m), d_G(P, Q) \geqslant 2d$, and thus the distance from the first half of S_3 to $x_2 P p_m$ is $\geqslant d_G(x_3, P) - ||S_3||/2 \geqslant 2d - d = d$, and the distance from the second half of S_3 to $x_2 P p_m$ is $\geqslant d_G(w_{i_3}, x_2 P p_m) - ||S_3||/2 \geqslant 2d - d = d$. Moreover, analogously to (12.5), we find

$$d_G(Q_2, Q_3) \geqslant d_G(p_m, u_3) - ||Q_2|| - ||Q_3|| \geqslant d_G(U_2 \cap \partial_G C, u_3) - 2(\ell+1) \geqslant (2\ell+d+2) - 2(\ell+1) = d.$$

Further, $d_G(x_2 P p_m, Q_3), d_G(u_3 Q x_3, Q_2) \geqslant d$ since $d_G(P, u_3), d_G(Q, p_m) \geqslant d_G(P, Q) \geqslant 2d$ and $P, Q \subseteq C \subseteq G - B$, while Q_2, Q_3 are shortest $\partial_G C$–$B_G(w, r-\ell)$ paths starting in p_m, u_3, respectively. It thus remains to show $d_G(S_3, Q_2) \geqslant d$: For this, it suffices to prove $d_G(x_3, Q_2) \geqslant 3d$, as S_3 has length $2d$ and starts in x_3. If $d_G(x_3, B) > 3d$, then $d_G(x_3, Q_2) \geqslant 3d$ follows immediately, as $Q_2 \subseteq G[B, 1]$. Thus, we may assume $d_G(x_3, B) \leqslant 3d$. Suppose for a contradiction that $d_G(x_3, Q_2) < 3d$. Then $d_G(x_3, p_m) \leqslant 2d_G(x_3, Q_2) < 6d$, since $x_3 \in V(C)$ and Q_2 is a shortest p_m–$B_G(w, r-\ell)$ path. Also $d_G(w_{i_3}, p') \leqslant d_G(w_{i_3}, x_3) + d_G(x_3, B) \leqslant 5d$, since W is a shortest path. Thus, $d_G(p_m, p') \leqslant d_G(p_m, x_3) + d_G(x_3, w_{i_3}) + d_G(w_{i_3}, p') < 13d$, which contradicts $d_G(p_m, p') \geqslant 2\ell + d$ as $\ell \geqslant 6d$. □

Corollary 12.4.10. *Let $K, \ell, r \in \mathbb{N}_{\geqslant 1}$ with $\ell \geqslant 66K$, and let B be a ball in a graph G around a vertex w of radius r. Suppose there is a component C of $G - B$ such that $\partial_G C$ contains three vertices that are pairwise at least $4 \cdot (2\ell + 11K + 2)$ far apart.*

If G has no K-fat K_4 minor, then there exists a vertex $w' \in V(C)$ such that for every component D of $G - (B' \cup A)$ that attaches to $B' := B_G(w, r - \ell)$ and $A := B_G(w', 22K)$ there exists a B'–A path P with $d_G(P, G[D, 1]) \geqslant 5K$.

Proof. We note that in particular $r \geqslant \ell$ as $\partial_G C$ contains two vertices which are $\geqslant 4 \cdot (2\ell + 11K + 2)$ apart. By Lemma 12.4.9 (with $d := 11K$) there exists a vertex $w' \in V(C)$ and three B'–$B_G(w', 22K)$ paths P_1, P_2, P_3 such that $d_G(P_i, P_j) \geqslant 11K$ for all $i \neq j \in [3]$. Set $A := B_G(w', 22K)$, and let D be a component of $G - (B' \cup A)$ that attaches to B' and A.

We show that there is some $i \in [3]$ such that $d_G(P_i, G[D, 1]) \geqslant 5K$. Towards a contradiction, suppose that $d_G(P_i, D) \leqslant 5K$ for all $i \in [3]$. It then follows that all three paths P_i lie in the unique component D' of $G - (B_G(w, r - \ell - 3K) \cup B_G(w', 19K))$ that contains D. For each $i \in [3]$, let P'_i be a shortest $B_G(w, r - \ell - 3K)$–$B_G(w', 19K)$ path in $P_i \cup G[B'] \cup G[A]$. Then $d_G(P'_i, P'_j) \geqslant 11K - 2 \cdot 3K = 5K$, which by Lemma 12.4.5 ($d_G(w, w') > r_1 + r_2 + 5K + 1$ with $r_1 = r - \ell - 3K, r_2 = 19K$, since $d_G(w', w) \geqslant r$ and $\ell \geqslant 23K$) implies that G has a K-fat K_4 minor, which contradicts our assumption. □

Lemma 12.4.11. *Let $K, r_1, r_2 \in \mathbb{N}_{\geqslant 1}$ with $r_1 \leqslant R_0(K)$ and $r_2 > \ell(K)$, and let v_1, v_2 be two vertices of a graph G which are at least $r_1 + r_2 + 2$ apart. Set $B_i := B_G(v_i, r_i)$ for $i \in [2]$ and $B'_2 := B_G(v_2, r_2 - \ell(K) - 1)$. Let C be the component of $G - B_2$ that contains v_1. Let D be a component of $G - (B_1 \cup B'_2)$ which attaches to B_1 and B'_2, and suppose there is a B_1–B'_2 path P in G such that $d_G(P, G[D, 1]) \geqslant 5K$.*

If G has no K-fat K_4 minor, then there exists an induced subgraph Y of $G[B_1 \cup V(D) \cup B_2]$ which admits an honest decomposition $(H, \mathcal{V})$ modelled on some graph $H \in \mathcal{H}_{SP}$ with terminals h_1, h_2 such that

(α) $V_{h_1} = B_1$ *and* $V_{h_2} \supseteq B_2$,

(β) $\operatorname{rad}_G(V_h) \leqslant R'_0(K)$ *for all nodes* $h \neq h_2$ *of* H,

(γ) $\operatorname{rad}_G(V_{h_2}) \leqslant r_2 + 2R'_0(K) + 1$,

(δ) $\operatorname{irads}((H, \mathcal{V})) \leqslant 3$, *and*

(ε) *for every component D' of $G - Y$ that meets C and D, we have* $\operatorname{rad}_G(N_G(D')) \leqslant R_2(K) - 1$ *and there is a node h of H such that $N_G(D') \subseteq V_h$.*

Proof. Since $\ell(K) \geqslant R_0(K) + 5K \geqslant r_1 + 5K$, applying Lemma 12.4.8 to $K, r_1, r'_2 := r_2 - \ell(K) - 1$, v_1, v_2 and D yields an induced subgraph $Y^1 \subseteq G[D, 1]$ and an honest decomposition $(H^1, \mathcal{V}^1)$ of Y^1 modelled on a graph $H^1 \in \mathcal{H}_{SP}$ with terminals h_1, h_2 which satisfies properties (a) – (e) from Lemma 12.4.8.

Set $B_2^{+1} := B_G(v_2, r_2 + 1)$ and

$$\tilde{H} := H^1[\{h \in V(H^1) \mid V_h^1 \cap B_2^{+1} \neq \emptyset\}].$$

We remark that h_2 and $N_{H^1}(h_2)$ are contained in $\tilde{H}$ since $(H^1, \mathcal{V}^1)$ is honest and (a), while $h_1 \notin V(\tilde{H})$ by (a) and because v_1 and v_2 are at least $r_2 + r_1 + 2$ apart.

We claim that $\tilde{H}$ is connected. Indeed, since $h_2 \in V(\tilde{H})$, it is enough to find for every $h \in V(\tilde{H})$ some h–h_2 path in $\tilde{H}$. Let $h \in V(\tilde{H})$ be given, pick some $v \in V_h^1 \cap B_2^{+1} \subseteq B_G(D, 1)$, and let Q be some v–$(N_G(D) \cap B'_2)$ path in $G[B_2^{+1} \cap B_G(D, 1)]$, which exists since B_1 and B_2^{+1} are disjoint and D is a component of $G - (B_1 \cup B'_2)$. Set $Q^1 := Q \cap Y^1$, and note that h_2 is a node of $H_{Q^1} := H[\{h \in V(H^1) \mid V_h^1 \cap Q^1 \neq \emptyset\}]$ since $V_{h_2}^1 = N_G(D) \cap B'_2$ by (a). Moreover, $H_{Q^1} \subseteq \tilde{H}$ by definition of Q. By (H2'), $H_{\tilde{Q}}$ is connected for every component $\tilde{Q}$ of Q^1, and thus the claim follows if Q^1 is connected. Otherwise, the claim follows as the neighbourhood of every component of $D - Y^1$ is contained in some bag of $(H^1, \mathcal{V}^1)$ by (e), and thus $H_{\tilde{Q}}$ and $H_{\hat{Q}}$ intersect for 'consecutive' components $\tilde{Q}$ and $\hat{Q}$ of Q^1, as $Q \subseteq B_G(D, 1)$.

We thus obtain a decomposition $(H, \mathcal{V}^2)$ of Y^1 where $H := H^1/\tilde{H}$ is the graph obtained from contracting $\tilde{H}$ down to a single vertex, which we again call h_2, by merging all bags of $(H^1, \mathcal{V}^1)$ to $V_{h_2}^1$ that contain a vertex of B_2^{+1}, i.e. $V_{h_2}^2 = \bigcup_{h \in \tilde{H}} V_h^1$ and $V_h^2 = V_h^1$ for every node $h \neq h_2$ of H^2. Note that H is still in $\mathcal{H}_{SP}$ with terminals h_1, h_2. We now obtain the desired graph-decomposition $(H, \mathcal{V})$ by letting $V_h := V_h^2 = V_h^1$ for all $h \in V(H) \setminus \{h_1, h_2\}$ and $V_{h_i} := V_{h_i}^2 \cup B_i$ for $i \in [2]$.

We claim that $(H, \mathcal{V})$ is a graph-decomposition of $Y := G[B_1 \cup V(Y^1) \cup B_2] \subseteq G[B_1 \cup V(D) \cup B_2]$. Indeed, $(H, \mathcal{V})$ satisfies (H2) as $(H, \mathcal{V}^2)$ already satisfied (H2), because, by construction, $V_{h_2}^2$ is the only bag of $(H, \mathcal{V}^2)$ that may contain vertices of B_2, and because $B_1 \setminus V^1_{h_1}$ does not meet Y^2 and B_2 as $Y^1 \subseteq G[D, 1]$ and because B_1 and B_2 are disjoint.

To see that (H1) holds, observe that, since $(H, \mathcal{V}^2)$ is a graph-decomposition of Y^1, and $V_{h_2} \supseteq B_2$ and $V_{h_1} = B_1$, it suffices to show that the parts of $(H, \mathcal{V})$ cover all edges between $Y^1 \cap D$ and $B_1 \cup B_2$ as well as all edges between B_1 and B_2. Since $N_G(D) \cap B_1 \subseteq V^2_{h_1}$ and $(H, \mathcal{V}^2)$ satisfies (H1), the parts of $(H, \mathcal{V}^2)$ already covered all edges between B_1 and $Y^1 \cap D$, and so the parts of $(H, \mathcal{V})$ do so as well. Further, the parts of $(H, \mathcal{V})$ cover all edges between B_2 and $Y^1 \cap D$ since $V(Y^1) \cap N_G(B_2) \subseteq V(Y^1) \cap B_2^{+1} \subseteq V_{h_2}$ as ensured by the contraction of $\tilde{H}$ in the construction of $(H^2, \mathcal{V}^2)$. Finally, there are no edges between B_1 and B_2 because $d_G(v_1, v_2) \geqslant r_1 + 2 + r_2$.

We now check that $(H, \mathcal{V})$ satisfies (α) – (ε). By definition, $V_{h_2} = V^2_{h_2} \cup B_2 \supseteq B_2$ and by (a) $V_{h_1} = V^2_{h_1} \cup B_1 = B_1$, and hence $(H, \mathcal{V})$ satisfies (α). To see that $(H, \mathcal{V})$ satisfies (β), let us first note that $\mathrm{rad}(V_{h_1}) = \mathrm{rad}_G(B_1) \leqslant r_1 \leqslant R_0(K)$. So by (b) and (c), we have $\mathrm{rad}_G(V_h) = \mathrm{rad}_G(V^1_h) \leqslant \max\{r_1 + 2R_0(K) + 5K + 1, R'_0(K)\} \leqslant R'_0(K)$ for all nodes $h \neq h_1$ of $V(H^1)$ with $h \notin B_{H^1}(h_2, 1)$. So since $\tilde{H}$ contains $N_{H^1}(h_2)$ and $\tilde{H}$ was contracted to h_2 in the construction of H, this yields (β).

Let us now verify that $(H, \mathcal{V})$ satisfies (γ). For this, let $v \in V_{h_2}$. If $v \in B_2$, then $d_G(v, v_2) \leqslant r_2 \leqslant r_2 + 2R'_0(K) + 1$ as desired. Otherwise, by construction, v is contained in V^1_h for some node h of $\tilde{H} \subseteq H^1$. If $h \in B_{H^1}(h_2, 1)$, then $d_G(v, v_2) \leqslant r'_2 + 2R_0(K) + 5K + 1$ by (c). Otherwise, by (b), we have $d_G(v_2, v) \leqslant r_2 + 1 + 2\,\mathrm{rad}_G(V^1_h) \leqslant r_2 + 1 + 2R'_0(K)$ because V^1_h meets $B_2^{+1} = B_G(B_2, 1)$ as $h \in V(\tilde{H})$. Thus, $(H, \mathcal{V})$ satisfies (γ). Further, (δ) holds for $(H, \mathcal{V})$ by (d).

We are thus left to check (ε). For this, let D' be a component of $G - Y$ that meets C and D; in particular, since $N_G(D) \cap B_1 \subseteq V(Y^1)$, we have that D' is a component of $G - (V(Y^1) \cup B_2)$. Let $\hat{D}$ be the component of $D - Y^1$ that contains D'. Since $D' \cap C \neq \emptyset$, there is a $\hat{D}$–B_1 path Q in C; let q be its first vertex in Y^1; in particular, $q \in N_G(\hat{D})$. Then

$$d_G(q, v_2) \geqslant d_G(C, v_2) > r_2 \geqslant r'_2 + \ell(K) = r'_2 + 2 \cdot R_0(K) + 5K + 2,$$

which by (c) implies that $q \notin V^1_h$ for every $h \in B_{H^1}(h_2, 1)$. By (e), there is a bag V^1_h of $(H^1, \mathcal{V}^1)$ such that $q \in N_G(\hat{D}) \subseteq V^1_h$; by the previous observation, we have $h \notin B_{H^1}(h_2, 1)$, and hence $\mathrm{rad}_G(V^1_h) \leqslant R'_0(K)$ as we have shown in the proof of (β). Now observe that

$$N_G(D') = (N_G(D') \cap V(Y^1)) \cup (N_G(D') \cap B_2) \subseteq N_G(\hat{D}) \cup (N_G(D') \cap B_2) \tag{12.6}$$

where we used that $D' \subseteq \hat{D} - B_2$. Since $D' \subseteq \hat{D}$, $N_G(\hat{D}) \subseteq V^1_h$ and $v_2 \notin V(\hat{D})$, every D'–v_2 path meets V^1_h. Thus, every vertex in $N_G(D') \cap B_2$ (which is empty if $h \notin \tilde{H}$) has distance at most $\ell(K)$ from V^1_h. Since also $N_G(\hat{D}) \subseteq V^1_h$, we have by (12.6) for all $u, w \in N_G(D')$ that

$$d_G(u, w) \leqslant d_G(u, V^1_h) + 2 \cdot \mathrm{rad}_G(V^1_h) + d_G(V^1_h, w) \leqslant \ell(K) + 2 \cdot R'_0(K) + \ell(K) \leqslant R_2(K) - 1.$$

This shows that $(H, \mathcal{V})$ satisfies (ε) and thus concludes the proof. $\square$

Proof of Lemma 12.4.2. Suppose that G has no K-fat K_4 minor for $K \in \mathbb{N}_{\geqslant 1}$; otherwise, we are done. Since $4 \cdot (2(\ell(K) + 22K + 1) + 11K + 2) = R_1(K)$, by Corollary 12.4.10 applied to $B_G(w, r + 22K + 1)$ and C^* (with $\ell := \ell(K) + 22K + 1$) there exists a vertex $w' \in V(C^*)$ such that for every component D of $G - (B_1 \cup B_2')$ that attaches to $B_1 := B_G(w', 22K)$ and $B_2' := B_G(w, r - \ell(K))$ there exists a B_1–B_2' path P such that $d_G(G[D, 1], P) \geqslant 5K$. Let $\mathcal{D}$ be the set of components D of $G - (B_1 \cup B_2')$ which meet C and attach to B_2'. We note that these also attach to B_1, as $w' \in V(C^*) \subseteq V(C)$ and C is connected. Set $B_2 := B$.

Note that $r > \ell(K)$, since there are two vertices in $\partial_G C^*$ which are at least $R_1(K) > 2(\ell(K) + 22K + 1)$ apart. We may thus apply Lemma 12.4.11 to $v_1 := w'$, $v_2 := w$, $r_1 := 22K$, $r_2 := r$ and every component $D \in \mathcal{D}$. This yields for every such D an honest partial graph-decomposition $\mathcal{H}^D = (H^D, \mathcal{V}^D)$ of G with support $Y^D \subseteq G[B_1 \cup V(D) \cup B_2]$ modelled on a graph $H^D \in \mathcal{H}_{SP}$ with terminals h_1^D, h_2^D which satisfies (α) to (ε).

We define the desired decomposition $(H, \mathcal{V})$ in multiple steps and check its desired properties afterwards. First, we define a decomposition $(H^1, \mathcal{V}^1)$ of $Y^1 := G[\bigcup_{D \in \mathcal{D}} V(Y^D)] \subseteq G[B_1 \cup (\bigcup_{D \in \mathcal{D}} V(D)) \cup B_2] \subseteq G[V(C) \cup B_2]$. The graph H^1 is obtained from the disjoint union of the H^D and three new vertices h_1, h_2, g by identifying all h_1^D with h_1 as well as all h_2^D with h_2 and adding the edge $h_2 g$. Then we assign the nodes of H^1 bags $V_g^1 := N_G(C)$, $V_{h_2}^1 := \bigcup_{D \in \mathcal{D}} V_{h_2^D}^D \supseteq B_2$, and $V_{h_1}^1 := \bigcup_{D \in \mathcal{D}} V_{h_1^D}^D = B_1$ and $V_h^1 := V_h^D$ for all other $h \in V(H^1)$ where D is the unique component in $\mathcal{D}$ with $h \in V(H^D)$. Then $(H^1, \mathcal{V}^1)$ is a graph-decomposition of Y^1. Indeed, it follows from $G[B_1 \cup B_2] \subseteq Y^D \subseteq G[B_1 \cup V(D) \cup B_2]$ for $D \in \mathcal{D}$ that the Y^D intersect only in B_1 and B_2. Hence, $(H^1, \mathcal{V}^1)$ is a graph-decomposition of Y^1, as every $(H^D, \mathcal{V}^D)$ is a graph-decomposition of Y^D and $B_1 = V_{h_1^D}^D$ and $B_2 \subseteq V_{h_2^D}^D$.

Next, we adjust $(H^1, \mathcal{V}^1)$ to a graph-decomposition $(H, \mathcal{V}^2)$ of $Y^2 \supseteq Y^1$: Let C' be a component of $G - Y^1$ which meets C. In particular, C' is contained in $G - (B_1 \cup B_2)$ because $B_1, B_2 \subseteq V(Y^1)$. Let $D_{C'}$ be the unique component of $G - (B_1 \cup B_2')$ which contains C'. By $\mathcal{C}$ we denote the set of all components C' of $G - Y^1$ which meet C and whose $D_{C'}$ attach to B_2'. In particular, for every $C' \in \mathcal{C}$ we have that $D_{C'} \in \mathcal{D}$ and C' is a component of $G - Y^{D_{C'}}$ that meets C and $D_{C'}$, since $Y^{D_{C'}} \subseteq Y^1$. Hence, it follows from (ε) of $(H^{D_{C'}}, \mathcal{V}^{D_{C'}})$ that there is a node $h_{C'} \in V(H^{D_{C'}})$ such that $N_G(C') \subseteq V_{h_{C'}}^{D_{C'}}$ and $\operatorname{rad}_G(N_G(C')) \leqslant R_2(K) - 1$.

Now we obtain H from H^1 by adding for each component $C' \in \mathcal{C}$ a new node $h'_{C'}$ and the edge $h_{C'} h'_{C'}$. We assign an $h'_{C'}$ the bag $V_{h'_{C'}}^2 := N_G(C') \cup \partial_G C'$ for $C' \in \mathcal{C}$. Moreover, we set $V_h^2 := V_h^1$ for all old nodes h of $H^1 \subseteq H$. Then it is immediate from the construction that $(H, \mathcal{V}^2)$ is a graph-decomposition of $Y^2 := G[V(Y^1) \cup (\bigcup_{C' \in \mathcal{C}} \partial_G(C'))]$. We now obtain our final graph-decomposition $(H, \mathcal{V})$ of $Y := Y^2 \cap G[C, 1]$ as the restriction of $(H, \mathcal{V}^2)$ to Y.

It remains to check that $(H, \mathcal{V})$ has the desired properties. By definition, $Y \subseteq G[C,1]$. We also have $\partial_G C \subseteq V(Y)$. Indeed, every vertex $v \in \partial_G C \subseteq N_G(B_2)$ which was not already contained in $V(Y^1) \supseteq B_2$ lies in $\partial_G C'$ for some component C' of $G - Y^1$. Then $C' \in \mathcal{C}$, i.e. the component $D_{C'}$ of $G - (B_1 \cup B_2')$ attaches to B_2': because $v \in V(C') \subseteq V(D_{C'})$ and B_2 and B_1 are disjoint, every shortest v–$N_G(B_2')$ path in G is through $B_2 \setminus B_2'$ and hence contained in $D_{C'}$. Now by construction, $v \in V_{h'_{C'}}$. Moreover, the graph H is in $\mathrm{Forb}_{\preccurlyeq}(K_4)$: The graph H^1 is obtained from the $H^D \in \mathcal{H}_{SP}$ by parallel compositions of the graphs $H^D \in \mathcal{H}_{SP}$ followed by a series composition with a path $h_2 g$ with terminals h_2 and g; so H^1 is again in $\mathcal{H}_{SP}$ with terminals h_1 and g by Proposition 12.4.1. Now the graph H is given by 1-sums of the single edges $h_{C'} h'_{C'}$ with the graph $H^1 \in \mathcal{H}_{SP} \subseteq \mathrm{Forb}_{\preccurlyeq}(K_4)$. Hence, $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$.

Since $V_g = V_g^2 = V_g^1 = N_G(C)$ by construction, $(H, \mathcal{V})$ satisfies (i). For (ii), we note that by (β) and (γ) of the $(H^D, \mathcal{V}^D)$ we have $\mathrm{oradw}((H^1, \mathcal{V}^1)) \leqslant r + 2R_0'(K) + 1$. For nodes $h \in H_2$ that are already in H_1, we have $V_h^1 = V_h^2$, and hence $\mathrm{rad}_G(V_h^2) \leqslant \mathrm{rad}_G(V_h^1)$. By construction of $(H^2, \mathcal{V}^2)$, every node in $h \in H_2 - H_1$ is of the form $h'_{C'}$ for some $C' \in \mathcal{C}'$, so their bags $V^2_{h'_{C'}}$ are contained in the respective $V^1_{h_{C'}} \cup N_G(V^1_{h_{C'}})$, and hence $\mathrm{rad}_G(V^2_{h'_{C'}}) \leqslant \mathrm{rad}_G(V^1_{h_{C'}}) + 1$. Thus, as $\mathrm{rad}_G(V_{h'_{C'}}) = \mathrm{rad}_G(V^2_{h'_{C'}}) \leqslant R_2(K)$, the graph-decomposition $(H, \mathcal{V})$ satisfies (ii). Also $(H, \mathcal{V})$ satisfies (iii): The $(H^D, \mathcal{V}^D)$ have radial spread at most 3. Thus, $(H^1, \mathcal{V}^1)$ has radial spread at most 6 by construction. This yields that $(H, \mathcal{V}^2)$ has radial spread at most 7 which yields that its restriction $(H, \mathcal{V})$ also has radial spread at most 7.

We claim that $(H, \mathcal{V})$ satisfies (iv). Indeed, let $\tilde{C}$ be a component of $G - Y$. First, assume that $\tilde{C}$ does not meet C. Since $N_G(C) \subseteq V(Y) \subseteq B_G(C,1)$, its neighbourhood $N_G(\tilde{C})$ is contained in $N_G(C) = V_g$. Also, we have $\mathrm{rad}_G(N_G(C)) \leqslant \mathrm{rad}_G(B) \leqslant r$, since C is a component of $G - B$ and $B = B_G(w, r)$. So we may assume from now on that $\tilde{C}$ meets C. Since $N_G(C), B_1 \subseteq V(Y)$, the component $\tilde{C}$ is contained in $G - (B_1 \cup B_2)$, and thus in C. As $Y = Y^2 \cap G[C,1]$, $Y^2 \supseteq Y^1$ and $\tilde{C} \subseteq C$, there is a unique component C' of $G - Y^1$ which contains $\tilde{C}$. If $D_{C'}$ does not attach to B_2', its neighbourhood $N_G(D_{C'})$ is contained in $B_1 = V^1_{h_1} = V^2_{h_1} = V_{h_1}$, and thus $D_{C'} = C' = \tilde{C}$ by their respective definition. So we may assume that $D_{C'}$ attaches to B_2', i.e. $C' \in \mathcal{C}$. Then the construction of $(H, \mathcal{V}^2)$ ensured that $V_{h'_{C'}} = V^2_{h'_{C'}} = N_G(C') \cup \partial_G C'$. Also $\tilde{C}$ is a component of $C' - \partial_G C'$, since $Y^2 = G[V(Y^1) \cup \bigcup_{C' \in \mathcal{C}} \partial_G C']$. Thus, we have $N_G(\tilde{C}) \subseteq \partial_G C' \subseteq V^2_{h'_{C'}} = V_{h'_{C'}}$. We have already seen above that $\mathrm{rad}_G(V_{h'_{C'}}) \leqslant R_2(K)$. □

13 Asymptotic half-grid and full-grid minors

We prove that every locally finite, quasi-transitive graph with a thick end whose cycle space is generated by cycles of bounded length contains the full-grid as an asymptotic minor and as a diverging minor. This in particular includes all locally finite Cayley graphs of finitely presented groups, and partially solves problems of Georgakopoulos and Papasoglu and of Georgakopoulos and Hamann.

Additionally, we show that every (not necessarily quasi-transitive) graph of finite maximum degree which has a thick end and whose cycle space is generated by cycles of bounded length contains the half-grid as an asymptotic minor and as a diverging minor.

This chapter is based on [9] and joint work with Matthias Hamann.

13.1 Introduction

All graphs in this chapter may be infinite, unless otherwise stated.

Fat minors are a coarse or metric variant of graph minors. They first appeared in works of Chepoi, Dragan, Newman, Rabinovich and Vaxes [37] and of Bonamy, Bousquet, Esperet, Groenland, Liu, Pirot and Scott [23]. They play an important role in many (open) problems at the intersection of structural graph theory and coarse geometry – an area which can be described as 'coarse graph theory'.

A *model* of a graph X in a graph G is a collection of connected *branch sets* and *branch paths* in G such that after contracting each branch set to a vertex, and each branch path to an edge, we obtain a copy of X. A model of X is *K-fat* (in G), for some $K \in \mathbb{N}$, if its branch sets and paths are pairwise at least K apart, except that we do not require this for incident branch set-path pairs (see also Section 9.5 for the definition). We say that X is a *(K-fat) minor* of G if G contains a (K-fat) model of X. The graph X is an *asymptotic minor* of G if X is a K-fat minor of G for every $K \in \mathbb{N}$. An important advantage of asymptotic minors over the usual minors is that they are preserved under quasi-isometries, and in particular, it does not depend on the choice of a finite generating set whether a Cayley graph of a finitely generated group contains a fixed graph as an asymptotic minor [75].

Recently, Georgakopoulos and Papasoglu [75] gave an overview of the area of 'coarse graph theory', where they presented results and open problems regarding the interplay of geometry and graphs, many of which concern fat minors. These problems have already attracted quite some attention; some (partial) solutions can be found in [8, 13, 37, 38, 69, 70, 99]. Our main contribution is a partial resolution of a problem of Georgakopoulos and Papasoglu about asymptotic grid minors in quasi-transitive graphs [75, Problem 7.3]. To state this problem, we first need some definitions.

An *end* of a graph G is an equivalence class of rays where two rays in G are equivalent if there are infinitely many pairwise disjoint paths between them in G. An end is *thick* if it has infinitely many pairwise disjoint rays. The *full-grid* is the graph on $\mathbb{Z} \times \mathbb{Z}$ in which two vertices (m, n) and (m', n') are adjacent if and only if $|m - m'| + |n - n'| = 1$, and the *half-grid*[1] is its induced subgraph on $\mathbb{N} \times \mathbb{Z}$.

One of the cornerstones of infinite graph theory is *Halin's Grid Theorem* [81, Satz 4′], which asserts that every graph with a thick end contains the half-grid as a minor. Following this approach, Heuer [89] characterised the graphs containing the full-grid as a minor. These graphs form a proper subclass of the graphs with a thick end: while it is clearly true that every graph with a full-grid minor has a thick end, the converse is false in general, as the half-grid itself already witnesses. However, as it turned out, if we only consider graphs which are *quasi-transitive*, i.e. graphs whose vertex set has only finitely many orbits under its automorphism group, then these two graph classes coincide. Indeed, Georgakopoulos and Hamann [74] showed that every locally finite, quasi-transitive graph with a thick end contains the full-grid as a minor.

Georgakopoulos and Papasoglu [75] asked whether this result can be generalised to the coarse setting in the following sense.

Problem 13.1.1 ([75, Problem 7.3])**.** *Let G be a locally finite Cayley graph of a one-ended finitely generated group. Must the half-grid be an asymptotic minor of G? Must the full-grid be an asymptotic minor of G?*

Note that every Cayley graph of a group is (quasi-)transitive. Moreover, the unique end of a one-ended, quasi-transitive graph is always thick [35, 122].

Our main theorem partially answers both questions in the affirmative, under the additional assumption that G is a locally finite Cayley graph of a finitely presented group. In fact, we show the following result.

Theorem 35. *Let G be a locally finite, quasi-transitive graph whose cycle space is generated by cycles of bounded length. If G has a thick end, then the full-grid is an asymptotic minor of G.*

(We refer the reader to Section 13.2.2 for the definitions concerning the cycle space.)

Note that Theorem 35 includes all locally finite Cayley graphs of finitely presented groups. Examples such as inaccessible graphs and groups [58, 59] or Diestel-Leader graphs [52, 68] indicate

[1] Note that usually the grid on $\mathbb{N}^2$ is referred to as the half-grid. However, for us it will be more convenient to work with the grid on $\mathbb{N} \times \mathbb{Z}$. It is easy to see that our results about the half-grid also hold for the grid on $\mathbb{N}^2$.

that the geometry of arbitrary locally finite, quasi-transitive (or Cayley) graphs may be far more involved. This is why generalising Theorem 35 to locally finite Cayley graphs of arbitrary finitely generated groups or even to all locally finite, quasi-transitive graphs may be much harder, and will require a different approach to that presented in this chapter (see the sketch of the proof in Section 13.3 for details).

For the proof of Theorem 35 we construct, for every such graph G, a single model of the full-grid (see Theorem 13.3.4), which can be turned into a K-fat model of the full-grid, for every $K \in \mathbb{N}$, by deleting some of its branch sets and paths. Moreover, it can be turned into a model of the full-grid that *diverges*: for any two diverging sequences of vertices and/or edges of the full-grid also their branch sets/paths diverge in G (see Section 13.2.3 for the definition).

Theorem 36. *Let G be a locally finite, quasi-transitive graph whose cycle space is generated by cycles of bounded length. If G has a thick end, then the full-grid is a diverging minor of G.*

This partially solves a question of Georgakopoulos and Hamann [74, Problem 4.1].

Krön and Möller [96, Theorem 5.5] proved that a locally finite, quasi-transitive connected graph has no thick end if and only if it is quasi-isometric to a tree. Thus, instead of assuming that the graph G in Theorems 35 and 36 has a thick end, we may assume that G is not quasi-isometric to a tree (see Section 13.7.2 for details).

As a first step in the proof of Theorem 35, we find the half-grid as an asymptotic minor. For this, we do not need the transitivity assumption on G. Indeed, we prove the following theorem.

Theorem 37. *Let G be a graph of finite maximum degree whose cycle space is generated by cycles of bounded length. If G has a thick end, then the half-grid is an asymptotic minor of G.*

Note that every graph satisfying the premise of Theorems 35 and 36 has finite maximum degree as it is locally finite and quasi-transitive.

Similar as in the proof of Theorem 35, we again construct a single model of the half-grid (see Theorem 13.3.3), which can be turned into a K-fat model of the half-grid, for every $K \in \mathbb{N}$, and into a diverging model of the half-grid.

Theorem 38. *Let G be a graph of finite maximum degree whose cycle space is generated by cycles of bounded length. If G has a thick end, then the half-grid is a diverging minor of G.*

This partially solves a question of Georgakopoulos and Hamann [74, Problem 4.2].

This chapter is structured as follows. In Section 13.2 we recall some important definitions. Section 13.3 consists of three parts. We first introduce some new definitions in Sections 13.3.1 and 13.3.2. We then give a sketch of the proofs of Theorems 35 to 38 in Sections 13.3.3 and 13.3.4, where we also state Theorems 13.3.3 and 13.3.4, our stronger results on half-grid and full-grid minors, which we already briefly mentioned above. In Sections 13.3.5 and 13.3.6 we derive Theorems 35 to 38 from Theorems 13.3.3 and 13.3.4. Section 13.4 contains some preparatory work about diverging and quasi-geodesic rays. We then prove Theorems 13.3.3 and 13.3.4 in Sections 13.5 and 13.6, respectively. We finish in Section 13.7 by discussing some related problems.

13.2 Preliminaries

A graph G is *quasi-transitive* if the automorphism group of G acts on $V(G)$ with only finitely many orbits, that is if $V(G)$ can be partitioned into finitely many sets $U_0, \dots, U_n$ such that for all $i \in \{0, \dots, n\}$ and $u, v \in U_i$ there exists an automorphism φ of G such that $\varphi(u) = v$. The *stabilizer* of a subgraph X of G consists of precisely those automorphisms of G that map X to itself.

A graph G is *accessible* if there exists some $n \in \mathbb{N}$ such that every two distinct ends $\varepsilon, \varepsilon'$ of G can be *distinguished* by a set U of at most n vertices of G, i.e. no component of $G - U$ contains rays from both ε and ε'.

Two rays R, S in G *diverge* if for every $n \in \mathbb{N}$ they have tails $R' \subseteq R$, $S' \subseteq S$ satisfying $d_G(R', S') > n$. A double ray R in G *diverges* if every two disjoint tails of R diverge.

13.2.1 (Hexagonal) grids

The *full-grid*, denoted by FG, is the graph on $\mathbb{Z}^2$ in which two vertices (m, n) and (m', n') are adjacent if and only if $|m - m'| + |n - n'| = 1$. The *hexagonal full-grid* is obtained from FG by deleting every other rung, as shown in Figure 13.1. The *(hexagonal) half-grid*, denoted by HG, is the induced subgraph of the (hexagonal) full-grid on vertex set $\mathbb{N} \times \mathbb{Z}$.

We call the double rays R^i of the (hexagonal) full- and half-grid its *vertical double rays* and the edges e_{ij} its *horizontal edges* (see Figure 13.1).

13.2.2 Cycle space

Let G be a graph. The *edge space* of G is the vector space over the 2-element field $\mathbb{F}_2$ of all functions $E(G) \to \mathbb{F}_2$: its elements correspond to the subsets of $E(G)$ and vector addition corresponds to symmetric difference. The *cycle space* of G is the subspace of the edge space of G spanned by all the cycles in G – more precisely, by their edge sets; for simplicity, we will not distinguish between the edge sets in the cycle space and the subgraphs they induce in G.

We say that the cycle space of G is *generated by cycles of bounded length* if there is some $n \in \mathbb{N}$ such that the cycles in G of length at most n generate the cycle space of G.

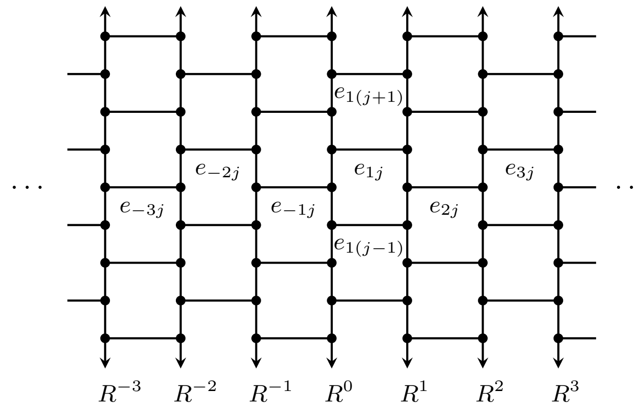

FIGURE 13.1: The hexagonal full-grid with vertical double rays R^i and horizontal edges e_{ij}.

Theorem 13.2.1 ([86, Corollary 3.2]). *Every locally finite, quasi-transitive graph whose cycle space is generated by cycles of bounded length is accessible.*

13.2.3 Asymptotic and diverging minors and subdivisions

Let G, X be graphs. Then X is an *asymptotic minor* of G, denoted by $X \prec_\infty G$, if X is a K-fat minor of G for all $K \in \mathbb{N}$. Let ε be an end of G. If X is a one-ended graph, then we write $X \prec_K^\varepsilon G$ if G contains a K-fat model $(\mathcal{V}, \mathcal{E})$ of X such that every ray in $\bigcup_{x \in V(X)} G[V_x] \cup \bigcup_{e \in E(X)} E_e$ is an ε-ray. Similarly, we write $X \prec_\infty^\varepsilon G$ if $X \prec_K^\varepsilon G$ for all $K \in \mathbb{N}$.

A model $(\mathcal{V}, \mathcal{E})$ of X in G *diverges* if for every two sequences $(x_n)_{n\in\mathbb{N}}$ and $(y_n)_{n\in\mathbb{N}}$ of vertices and/or edges of X such that $d_X(x_n, y_n) \to \infty$, we have $d_G(U_n, W_n) \to \infty$ where $U_n := V_{x_n}$ if $x_n \in V(X)$ and $U_n := V(E_{x_n})$ if $x_n \in E(X)$ and analogously $W_n := V_{y_n}$ or $W_n := V(E_{y_n})$.

A *subdivision* of a graph X is a graph which arises from X by replacing every edge in X by a new path between its endvertices such that no new path has an inner vertex in $V(X)$ or on any other new path. The original vertices of X are the *branch vertices* of the subdivision and the new paths are its *branch paths*. Let G be a graph and let $H \subseteq G$ be a subdivision of X with branch vertices v_x for $x \in V(X)$ and branch paths E_e for $e \in E(X)$. Then H is *K-fat* (in G) if there are sets $V_x \subseteq V(H)$ with $v_x \in V_x$ for $x \in V(X)$ and paths $E'_e \subseteq E_e$ for $e \in E(X)$ such that $((V_x)_{x\in V(X)}, (E'_e)_{e\in E(X)})$ is a K-fat model of X. The subdivision H of X *diverges* (in G) if the model $((\{v_x\})_{x\in V(X)}, (E_e)_{e\in E(X)})$ of X in G diverges.

13.3 Further definitions and a sketch of the proofs in this chapter

In this section we first introduce ultra fat minors and escaping subdivisions of certain graphs (see Sections 13.3.1 and 13.3.2). We then give in Sections 13.3.3 and 13.3.4 a sketch of the proofs of Theorems 35 to 38. There, we also state two stronger theorems, Theorems 13.3.3 and 13.3.4, from which we then derive Theorems 35 to 38 in Sections 13.3.5 and 13.3.6.

13.3.1 Ultra fat minors

We say that a model $((V_i)_{i\in\mathbb{N}}, (E_{ij})_{i\neq j\in\mathbb{N}})$ of $K_{\aleph_0}$ in a graph G is *ultra fat* if

- $d_G(V_i, V_j) \geqslant \min\{i, j\}$ for all $i \neq j \in \mathbb{N}$,
- $d_G(E_{ij}, E_{k\ell}) \geqslant \min\{i, j, k, \ell\}$ for all $i, j, k, \ell \in \mathbb{N}$ with $\{i, j\} \neq \{k, \ell\}$, and
- $d_G(V_i, E_{k\ell}) \geqslant \min\{i, k, \ell\}$ for all $i, k, \ell \in \mathbb{N}$ with $i \notin \{k, \ell\}$.

Further, we say that $K_{\aleph_0}$ is an *ultra fat minor* of G, and write $K_{\aleph_0} \prec_{UF} G$, if G contains an ultra fat model of $K_{\aleph_0}$. The idea is that an ultra fat model of $K_{\aleph_0}$ in a graph G witnesses that G contains $K_{\aleph_0}$ as an asymptotic minor. Indeed, if $((V_i)_{i\in\mathbb{N}}, (E_{ij})_{i\neq j\in\mathbb{N}})$ is an ultra fat model of $K_{\aleph_0}$ in G, then $((V_i)_{i\in\mathbb{N}_{\geqslant K}}, (E_{ij})_{i\neq j\in\mathbb{N}_{\geqslant K}})$ is a K-fat model of $K_{\aleph_0}$ in G. In particular, we have the following observation.

Observation 13.3.1. *If a graph G contains $K_{\aleph_0}$ as an ultra fat minor, then it contains every countable graph as an asymptotic minor. Moreover, if $K_{\aleph_0} \prec^{\varepsilon}_{UF} G$ for some end ε of G, then also $X \prec^{\varepsilon}_{\infty} G$ for every one-ended, countable graph X.* □

Moreover, the following observation is immediate from the definitions.

Observation 13.3.2. *If a graph G contains $K_{\aleph_0}$ as an ultra fat minor, then it contains every countable graph as a diverging minor, and in particular, it contains every countable graph of maximum degree at most 3 as a diverging subdivision. Moreover, if $K_{\aleph_0} \prec^{\varepsilon}_{UF} G$ for some end ε of G, then we may choose the diverging minor / subdivision so that all its rays lie in ε.* □

13.3.2 Escaping subdivisions

We call the double rays in a subdivision of the hexagonal half- or full-grid corresponding to the vertical double rays R^i of the hexagonal half- or full-grid its *vertical (double) rays*, and the branch paths corresponding to the horizontal edges e_{ij} its *horizontal paths*, and we usually denote the former by S^i and the latter by P_{ij}. Whenever we introduce a subdivision of the hexagonal half- or full-grid with vertical double rays S^i without specifying the vertex sets of the S^i, we tacitly assume that $S^i = \dots s^i_{-1}s^i_0s^i_1\dots$ and that their tails $S^i_{\geqslant 0}$ are the image of the 'upper' half of the vertical double ray R^i of the hexagonal half- or full-grid.

Let G be a graph and let $H \subseteq G$ be a subdivision of the hexagonal half-grid with vertical double rays S^i and horizontal paths P_{ij}. We say that H is *escaping* if there are $0 := M_0 < M_1 < \ldots \in \mathbb{N}$ such that $M_i > M_{i-1} + 2i$ for all $i \geqslant 1$ and

(i) $S^i \subseteq G[S^0, M_i] - B_G(S^0, M_{i-1} + 2i)$ for all $i \in \mathbb{N}_{\geqslant 1}$, and

(ii) $P_{1j} \subseteq G[S^0, M_1]$ and $P_{ij} \subseteq G[S^0, M_i] - B_G(S^0, M_{i-2} + i)$ for all $i \in \mathbb{N}_{\geqslant 2}$ and $j \in \mathbb{Z}$.

A subdivision $H \subseteq G$ of the hexagonal full-grid with vertical double rays S^i and horizontal paths P_{ij} is *escaping* if the S^i and P_{ij} with $i \geqslant 0$ form an escaping subdivision of the hexagonal half-grid as well as the S^i and P_{ij} with $i \leqslant 0$, and if there is some $M \in \mathbb{N}$ such that the S^i with $i > 0$ are contained in a different component of $G - B_G(S^0, M)$ than the S^i with $i < 0$.

13.3.3 Sketch of the proofs of Theorems 37 and 38

We will prove Theorems 37 and 38 simultaneously by showing the following stronger result.

Theorem 13.3.3. *Let ε be a thick end of a graph G with finite maximum degree whose cycle space is generated by cycles of bounded length. Then either $K_{\aleph_0} \prec^{\varepsilon}_{UF} G$ or G contains an escaping subdivision H of the hexagonal half-grid whose rays all lie in ε.*

By Observations 13.3.1 and 13.3.2, an ultra fat model of $K_{\aleph_0}$ contains a diverging and a K-fat model of the half-grid for every $K \in \mathbb{N}$. So to derive Theorems 37 and 38 from Theorem 13.3.3 it suffices to show that an escaping subdivision of the hexagonal half-grid also contains a diverging and a K-fat subdivision of the hexagonal half-grid (see Sections 13.3.5 and 13.3.6).

For the proof of Theorem 13.3.3, we first show that G contains for every thick end ε a diverging double ε-ray R (see Theorem 13.4.1), and we then set $S^0 := R$. Second, we show that G contains double rays $S^1, S^2, \ldots$ such that the S^i are contained in increasingly distant 'thickened cylinders' around R of the form $G[R, M_i] - B_G(R, M_{i-1} + 2i)$ for some $M_0 < M_1 < \ldots \in \mathbb{N}$, as required by (i) for the vertical double rays of an escaping subdivision of the hexagonal half-grid. Finally, we connect the S^i by infinitely many paths so that infinitely many of them either form the vertical double rays of an escaping subdivision of the hexagonal half-grid or they form the branch sets of an ultra fat model of $K_{\aleph_0}$ (see Lemma 13.5.4).

Let us describe the second step in more detail. We will choose the S^i recursively, starting from the diverging double ε-ray $R = \ldots r_{-1} r_0 r_1 \ldots (=: S^0)$. For this, we first show that $C[\partial_G C, \lfloor \frac{\kappa-2}{2} \rfloor] = C \cap G[R, L + \lfloor \frac{\kappa}{2} \rfloor]$ is connected for every $L \in \mathbb{N}$ and every component C of $G - B_G(R, L)$, where $\kappa \in \mathbb{N}$ is such that the cycle space of G is generated by cycles of length $\leqslant \kappa$ (see Lemma 13.5.1). Note that this is the only part in in the proofs of Theorems 35 to 38 where we use the assumption on the cycle space; nevertheless, the assumption is crucial here, and the rest of the proof relies on this lemma.[2]

We then show that for every $L \in \mathbb{N}$ some component C of $G - B_G(R, L)$ is 'long', i.e. it has a neighbour in $B_G(R_{\geqslant j}, L)$ and in $B_G(R_{\leqslant -j}, L)$ for all $j \in \mathbb{N}$ (see Lemma 13.5.3). Combining that C is long and $C[\partial_G C, \lfloor \frac{\kappa-2}{2} \rfloor]$ is connected then allows us to find a double ray in $C[\partial_G C, \lfloor \frac{\kappa-2}{2} \rfloor]$, which thus also lies in $G[R, L + \lfloor \frac{\kappa}{2} \rfloor] - B_G(R, L)$. Hence, we may proceed recursively by increasing the radius L of the ball around R by a summand of $\lfloor \frac{\kappa}{2} \rfloor + 2i$ in each step.

13.3.4 Sketch of the proof of Theorems 35 and 36

Similar as before, we prove Theorems 35 and 36 simultaneously, by showing the following stronger result.

Theorem 13.3.4. *Let ε be a thick end of a locally finite, quasi-transitive graph G whose cycle space is generated by cycles of bounded length. Then either $K_{\aleph_0} \prec_{UF}^{\varepsilon} G$ or G contains an escaping subdivision of the hexagonal full-grid whose rays all lie in ε.*

Similar as above, Observations 13.3.1 and 13.3.2 together with results from Section 13.3.5 below will show that it suffices to prove Theorem 13.3.4 in order to obtain Theorems 35 and 36 (see Section 13.3.6).

The proof of Theorem 13.3.4 builds on Theorem 13.3.3. From the proof of Theorem 13.3.3 it follows that we have more control over where the escaping subdivision of the hexagonal half-grid lies (see Theorem 13.3.3', the detailed version of Theorem 13.3.3, in Section 13.5). For this, let ε be a thick end of G, and let R be a diverging double ε-ray. Given a 'thick' component C of $G - B_G(R, K)$ for some $K \in \mathbb{N}$, that is one which includes a long component of $G - B_G(R, L)$ for every $L \geqslant K$, we in fact obtain an escaping subdivision H of the hexagonal half-grid whose first vertical double ray is R and which is 'mostly' contained in C (unless Theorem 13.3.3' yields an ultra fat model of $K_{\aleph_0}$, in which case we are immediately done). Now suppose that for some large enough $L \in \mathbb{N}$ there is another thick component D of $G - B_G(R, L)$. Then Theorem 13.3.3 yields another escaping subdivision H' of the hexagonal half-grid whose first vertical double ray is R and which is 'mostly' contained in D (or an ultra fat model of $K_{\aleph_0}$). Gluing H and H' together along their common first vertical double ray R then yields the desired subdivision of the hexagonal full-grid (see Lemma 13.6.2).

It thus suffices to prove that G contains a diverging double ε-ray R such that, for some large enough $K \in \mathbb{N}$, there are two distinct thick components of $G - B_G(R, K)$. This step is mainly divided into two lemmas (see Lemmas 13.6.8 and 13.6.9). We first show that if R' is a double ε-ray

[2] This is not entirely true. We in fact prove stronger versions of Theorems 35 to 38 (see Section 13.3.6 below), which find the desired minors in a prescribed end. For this, in the case of Theorems 35 and 36, we need the assumption on the cycle space once more, to ensure that the graph is accessible.

which is not only diverging but even quasi-geodesic, then it is enough that for some large enough $K \in \mathbb{N}$ there are distinct components $C \neq D$ of $G - B_G(R', K)$ such that C is thick but D is only 'half-thick' (see Lemma 13.6.8) because then we can use the quasi-transitivity of G to find another quasi-geodesic double ε-ray R such that $G - B_G(R, K)$ has two distinct thick components (see Lemma 13.6.8). Here, a component of $G - B_G(R, K)$ is half-thick if it includes for every $L \geqslant K$ a component of $G - B_G(R, L)$ which is 'half-long', i.e. which has neighbours in $B_G(R_{\geqslant n}, L)$ *or* in $B_G(R_{\leqslant -n}, L)$ for all $n \in \mathbb{N}$.

Next, we show that such a double ray R' exists. For this, we first prove that G contains three ε-rays R_1, R_2, R_3 that intersect pairwise in a single common vertex such that $R_1 \cup R_2 \cup R_3$ is quasi-geodesic (see Theorem 13.4.5). Applying (the detailed version Theorem 13.3.3' of) Theorem 13.3.3 to the quasi-geodesic, and hence diverging, double ray $R_1 \cup R_2$ then yields an escaping subdivision H of the hexagonal half-grid whose first vertical double ray is $R_1 \cup R_2$. Now for every $K \in \mathbb{N}$, by the definition of escaping, H will lie 'mostly' in one component C_K of $G - B_G(R, K)$, which then needs to be thick. We then analyse where R_3 lies in relation to H. If, for some large enough $L \in \mathbb{N}$, the ray R_3 has a tail in a component $D_L \neq C_L$ of $G - B_G(R_1 \cup R_2, L)$, then we are done since D_L needs to be half-thick as R_3 diverges from $R_1 \cup R_2$ but lies in the same end as R_1 and R_2.

Otherwise, again since R_3 diverges from $R_1 \cup R_2$, it has a tail in C_K for all $K \in \mathbb{N}$. We then distinguish two cases. First assume that R_3 is far away from H. Then, since R_3 has a tail in each C_K, we can connect R_3 and H by infinitely many paths. These paths together with R_3 then yield infinitely many H-paths that 'jump over' H. We then use these paths together with H to find an ultra fat model of $K_{\aleph_0}$. Otherwise, R_3 lies close to H. Then either R_3 separates H into an 'upper half' containing (a tail of) R_1 and a 'lower half' containing (a tail of) R_2, and then $R_1 \cup R_3$ (or symmetrically $R_2 \cup R_3$) is the desired double ray R', or there are infinitely many H-paths that 'jump over' R_3, which then again yield an ultra fat model of $K_{\aleph_0}$ (see Lemma 13.6.9).

13.3.5 Obtaining fat and diverging minors from escaping subdivisions

In this section we describe how one can turn an escaping subdivision H of the hexagonal half-grid (full-grid) into a K-fat or diverging minor of the hexagonal half-grid (full-grid).

Lemma 13.3.5. *Let G be a locally finite graph, and let $H \subseteq G$ be an escaping subdivision of the hexagonal half-grid (full-grid) whose vertical double ray S^0 diverges. Then the following assertions hold for all $K \in \mathbb{N}$:*

(i) *H contains a subdivision of the hexagonal half-grid (full-grid) which is K-fat in G, and*

(ii) *H contains a subdivision of the hexagonal half-grid (full-grid) which diverges in G.*

In fact, the subdivisions which we obtain from Lemma 13.3.5 (i) and (ii) will have the property that their sets of vertical double rays are a subset of the vertical double rays of H.

Since we will have to delete some of the branch paths from H in the proof of Lemma 13.3.5, we need the following auxiliary result.

Proposition 13.3.6. *Let H be a subdivision of the hexagonal half-grid (full-grid) with vertical double rays S^i and horizontal paths P_{ij}. Let H' be obtained from H by deleting some of the P_{ij}. If H' still contains, for every $i \in \mathbb{N} \setminus \{0\}$ ($i \in \mathbb{Z} \setminus \{0\}$), infinitely many P_{ij} with $j \in \mathbb{N}$ and infinitely many P_{ij} with $j \in \mathbb{Z}_{\leqslant 0}$, then H' contains a subdivision H'' of the hexagonal half-grid (full-grid) whose vertical double rays are the S^i and whose set of horizontal paths is a subset of the P_{ij}.*

Proof. To obtain the desired graph H'', one may recursively select paths $P_{ij} \subseteq H'$ with sufficiently large $|j|$ to represent the edges f_k in the order indicated in Figure 13.2 (and similarly for the hexagonal full-grid in the order $f_1, f_2, f_2', f_3, f_3', f_4, f_4', f_5, f_6, f_7, f_7', \ldots$). □

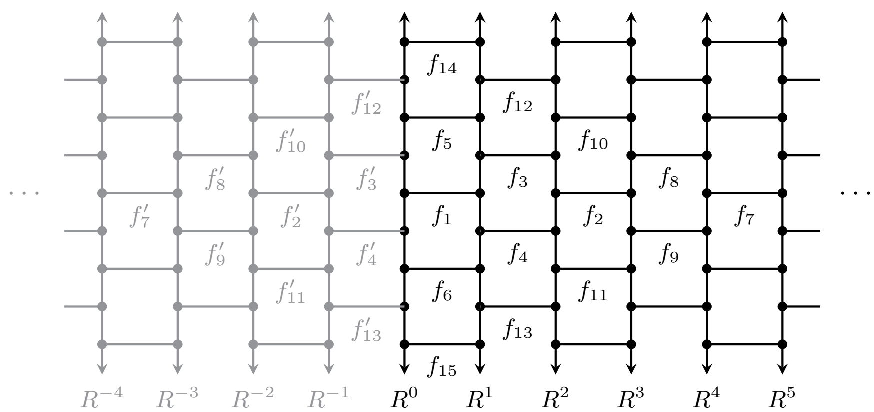

FIGURE 13.2: The hexagonal half-grid (full-grid) with an enumeration of its horizontal edges as needed for the proof of Proposition 13.3.6.

For the proof of Lemma 13.3.5 we need the following two auxiliary results, which assert that escaping subdivisions have some additional properties.

Lemma 13.3.7. *Let $R = \ldots r_{-1} r_0 r_1 \ldots$ be a diverging double ray in a locally finite graph G, let $L \in \mathbb{N}$, and let S be a ray in $G[R, L]$. Suppose there are infinitely many pairwise disjoint $R_{\geqslant 0}$–S paths P_i in $G[R, L]$. Then S has a tail T such that $T \subseteq G[R_{\geqslant 0}, L]$, and all but finitely many P_i are contained in $G[R_{\geqslant 0}, L]$.*

Proof. Since R diverges, there is some $n \in \mathbb{N}$ such that $R_{\geqslant n}$ and $R_{\leqslant -n}$ have distance at least $2L + 2$ from each other; in particular, $G[R_{\geqslant n}, L]$, and $G[R_{\leqslant -n}, L]$ are disjoint and not joined by an edge. Hence, they are separated in $G[R, L]$ by $B_G(r_{-n+1} R r_{n-1}, L)$. Since $B_G(r_{-n+1} R r_{n-1}, L)$ is finite as G is locally finite, it follows that S is eventually contained in either $G[R_{\geqslant n}, L]$ or

$G[R_{\leqslant -n}, L]$ and that at most finitely many P_i meet $B_G(r_{-n+1}Rr_{n-1}, L)$. As the P_i are disjoint and start in $R_{\geqslant 0}$, and hence infinitely many P_i start in $R_{\geqslant n}$, it follows that S has a tail T such that $T \subseteq G[R_{\geqslant n}, L] \subseteq G[R_{\geqslant 0}, L]$ and that all but finitely many P_i lie in $G[R_{\geqslant 0}, L]$. □

Corollary 13.3.8. *Let G be a locally finite graph and let H be an escaping subdivision of the hexagonal half-grid with vertical double rays S^i and horizonal paths P_{ij} such that S^0 diverges. Then*

(i) *for all $i, k \in \mathbb{N}$ there is $\ell \in \mathbb{N}$ such that $S^i_{\geqslant \ell} \subseteq G[S^0_{\geqslant k}, M_i]$ and $S^i_{\leqslant -\ell} \subseteq G[S^0_{\leqslant -k}, M_i]$, and*

(ii) *for all $i, k \in \mathbb{N}$ there is $\ell \in \mathbb{N}$ such that $P_{ij} \subseteq G[S^0_{\geqslant k}, M_i]$ and $P_{i(-j)} \subseteq G[S^0_{\leqslant -k}, M_i]$ for all $j \geqslant \ell$.*

Proof. Let $i, k \in \mathbb{N}$ be given. Set $R := S^0$, where we enumerate $R = \dots r_{-1}r_0r_1 \dots$ so that $r_0 = s^0_k$. Applying Lemma 13.3.7 to R, $L := M_i$ and $S := S^i_{\geqslant 0}$ and the paths $P_j := P_{ij}$ for $j \geqslant 0$ yields some $m \in \mathbb{N}$ such that $S^i_{\geqslant m}$ and all P_{ij} with $j \geqslant m$ are contained in $G[R_{\geqslant 0}, M_i] = G[S^0_{\geqslant k}, M_i]$. Similarly, we find some $n \in \mathbb{N}$ such that $S^i_{\leqslant -n}$ and all P_{ij} with $j \leqslant -n$ are contained in $G[S^0_{\leqslant -k}, M_i]$. Then $\ell := \max\{m, n\}$ is as desired. □

The next lemma finds a diverging subdivision in an escaping subdivision of the hexagonal half- or full-grid.

Lemma 13.3.9. *Let H be an escaping subdivision of the hexagonal half-grid (full-grid) in a locally finite graph G with vertical double rays S^i and horizontal paths P_{ij} such that S^0 diverges. Then there exists an escaping subdivision $H' \subseteq H$ of the hexagonal half-grid (full-grid) whose vertical double rays are the S^i and whose horizontal paths P'_{ij} are a subcollection of the P_{ij} such that H' diverges and such that for every two non-incident edges of the hexagonal half-grid (full-grid) their images in H' are at least K apart in G if they are contained in $H'_K := \bigcup_{i \geqslant K} S^i \cup \bigcup_{i > K, j \in \mathbb{Z}} P'_{ij}$.*

Proof. We only give the proof for the hexagonal full-grid; the construction for the hexagonal half-grid is analogous. For the sake of this proof, we denote the horizontal edges e_{ij} of the hexagonal full-grid by $f_{i(2j)}$ if $i \in 2\mathbb{Z}+1$, and by $f_{i(2j-1)}$ if $i \in 2\mathbb{Z}$, and we enumerate the P_{ij} accordingly. Let x^{i-1}_j, x^i_j denote the endvertices of P_{ij} on S^{i-1} and S^i, respectively. We will recursively select the branch paths Q_{ij} of the edges f_{ij} amongst the $P_{k\ell}$ such that

(1) $d_G(Q_{ij}, Q_{k\ell}) \geqslant \max\{|i|, |j|, |k|, |\ell|\}$ for all $i, j, k, \ell \in \mathbb{Z}$ such that $\{i, j\} \neq \{k, \ell\}$,

(2) $d_G(Q_{ij}, S^k y^k_{j-1} \cup y^k_{j+1} S^k) \geqslant \max\{|i|, |j|, |k|\}$ for all $i, j, k \in \mathbb{Z}$, and

(3) $d_G(S^i y^i_{j-1}, y^k_j S^k) \geqslant \max\{|i|, |k|, |j|\}$ for all $i, j, k \in \mathbb{Z}$,

where y^{i-1}_j and y^i_j denote the endvertices of Q_{ij} on S^{i-1} and S^i, respectively. Let H' denote the graph obtained from the union of the S^i and the Q_{ij}. Clearly, H' is still an escaping subdivision of the hexagonal full-grid (whose horizontal paths P'_{ij} are essentially the Q_{ij}, except that they are again enumerated as usual). It follows from (1)–(3) and (i) of escaping subdivisions that H'_K

has the desired property. Moreover, H' diverges. Indeed, the distances in G between the images U_n, W_n in H' of vertices and/or edges a_n, b_n of the hexagonal full-grid which form diverging sequences $(a_n)_{n\in\mathbb{N}}, (b_n)_{n\in\mathbb{N}}$ in H grow by (1)–(3) unless the U_n, W_n are of the form $Q_{ij}, y^k_{j-1}S^k y^k_j$ or $Q_{ij}, y^k_j S^k y^k_{j+1}$ or $y^i_{j-1}S^i y^i_j, y^k_{j-1}S^k y^k_j$. But their distances grow because of (i) and (ii) of escaping subdivisions; we omit the details.

We now describe how we choose the paths Q_{ij}. First, we set $Q_{i0} := P_{i0}$ for all $i \in 2\mathbb{N}+1$ and $Q_{i0} := P_{(i-1)0}$ for $i \in 2\mathbb{Z}_{<0}+1$. Now let $n \in \mathbb{N}$ be given, and assume that we have already chosen paths Q_{ij} for all $|j| < n$. Assume further, without loss of generality, that n is even. We now describe how we choose the paths Q_{in}, where $i \in 2\mathbb{Z}+1$ since n is even. The choice of the paths $Q_{i(-n)}$ can be done analogously after the choice of the Q_{in}.

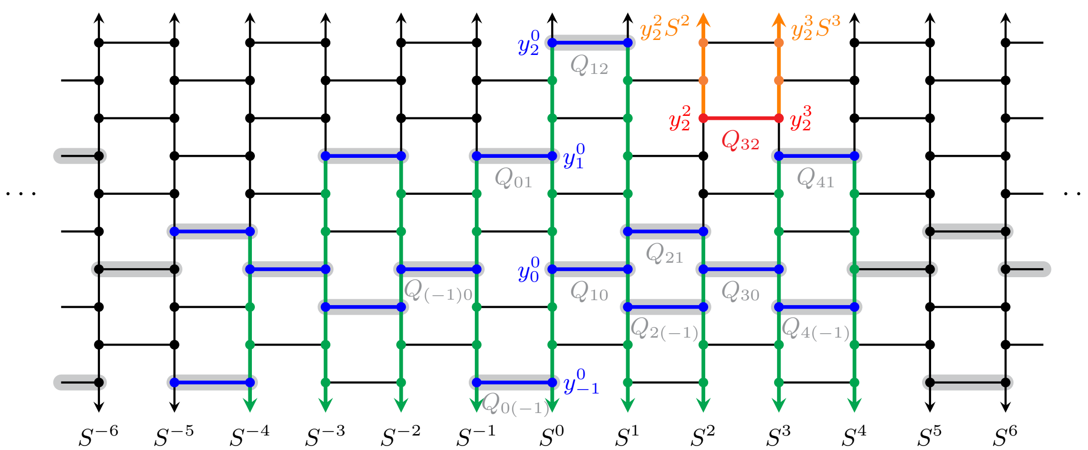

FIGURE 13.3: Depicted in blue and green are the subgraphs X_{32} and Y_{32} that are used to choose Q_{32}. The paths Q_{ij} that are chosen before Q_{32} are shown in grey.

So let $i \in 2\mathbb{N}+1$ be given, and assume that we have already chosen paths Q_{kn} for all $k \in 2\mathbb{N}+1$ with $k < i$. We will now select a path P_{ij} to be Q_{in}; again, the choice of the $Q_{(-i)n}$ can be done analogously after the choice of the Q_{in}.

By (i) and (ii), we have $d_G(P_{ij}, P_{k\ell}), d_G(S^i, P_{k\ell}) \geqslant \max\{i, |k|\}$ for all $k \notin \{i-1, i, i+1\}$ and $j, \ell \in \mathbb{Z}$. Hence, no matter which P_{ij} we choose to be Q_{in}, we will have that $d_G(Q_{in}, Q_{k\ell}) \geqslant \max\{i, n, |k|, |\ell|\}$ and $d_G(y^{i-1}_n S^{i-1}, Q_{k\ell}), d_G(y^i_n S^i, Q_{k\ell}) \geqslant \max\{i, n, |k|\}$ for all k with $|k| \geqslant k_{in} := \max\{i+2, n\}$ and $|\ell| \leqslant n$. So when choosing Q_{in}, we only need to consider those finitely many $Q_{k\ell}$ with $|k| < k_{in}$ and $|\ell| \leqslant n$. Since $X_{in} := \bigcup_{|k|<k_{in}, |\ell|<n} Q_{k\ell} \cup \bigcup_{0<k<i} Q_{kn}$ is finite and G is locally finite, all but finitely many P_{ij} have the property that $x^{i-1}_j S^{i-1} \cup P_{ij} \cup x^i_j S^i$ has distance at least $\max\{i, n\}$ from X_{in} (see Figure 13.3).

Moreover, by (i) and (ii), we have $d_G(P_{ij}, S^k), d_G(S^i, S^k) \geqslant \max\{i, |k|\}$ for all $k \neq i$. Hence, no matter which P_{ij} we choose to be Q_{in}, we will have that $d_G(Q_{in}, S^k), d_G(y_n^i S^i, S^k), d_G(y_n^{i-1} S^{i-1}, S^k)$ $\geqslant \max\{i, n, |k|\}$ for all k with $|k| \geqslant \max\{i+1, n\}$ and $|\ell| \leqslant n$. So when choosing Q_{in}, we only need to consider those finitely many S^k with $|k| < \max\{i+1, n\}$. By Corollary 13.3.8 (i) and (ii) and because S^0 diverges, all but finitely many of the P_{ij} have the property that $x_n^{i-1} S^{i-1} \cup P_{ij} \cup x_j^i S^i$ has distance at least $\max\{i, n\}$ from $Y_{in} := \bigcup_{|k|<k_{in}} S^k y^k_{(n-1)}$. Hence, we can pick a path P_{ij} whose endvertices on S^i and S^{i-1} appear on S^{i-1} and S^i after the endvertices of $Q_{(i-1)(n-1)}$ and $Q_{(i+1)(n-1)}$, respectively, such that $x_j^{i-1} S^{i-1} \cup P_{ij} \cup x_j^i S^i$ has distance at least $\max\{i, n\}$ from $X_{in} \cup Y_{in}$, and we set $Q_{in} := P_{ij}$. Clearly, Q_{in} is as desired. □

We are now ready to prove Lemma 13.3.5.

Proof of Lemma 13.3.5. (ii): By Lemma 13.3.9 every escaping subdivision H of the hexagonal half- or full-grid whose vertical double ray S^0 diverges contains a diverging subdivision H' of the hexagonal half- or full-grid, respectively, as a subgraph.

(i): Assume first that H is an escaping subdivision of the hexagonal half-grid, and let $H' \subseteq H$ be obtained from H by applying Lemma 13.3.9. Further, let $\tilde{H}$ be obtained from H' by deleting all horizontal paths that are not of the form P'_{ij} for $i \in 2\mathbb{Z}$ and $j \in 3\mathbb{Z}$ or $i \in 2\mathbb{Z}+1$ and $j \in 3\mathbb{Z}+1$. Then also the subgraph $\tilde{H}_K$ of $\tilde{H}$ consisting of all S^i with $i \geqslant K$ and all $P'_{ij} \subseteq \tilde{H}$ with $i > K$ is a subdivision of the hexagonal half-grid, and we claim that it is K-fat. Indeed, to turn $\tilde{H}_K$ into a K-fat model of the hexagonal half-grid we may choose the following sets V_x as branch sets, for every $x \in V(\tilde{H}_K)$ of degree 3. Let $x \in V(S^i)$ and let E_e and E_f be the branch paths of $H'_K := \bigcup_{i \geqslant K} S^i \cup \bigcup_{i>K, j\in\mathbb{Z}} P'_{ij} \subseteq H'$ starting at x that are contained in S^i. We then choose as V_x the union over E_e and E_f. By construction and since the images in H'_K of any two non-incident edges of the hexagonal half-grid have distance at least K in G, it follows that the model is K-fat.

Second, assume that H is an escaping subdivision of the hexagonal full-grid, let $H' \subseteq H$ be obtained from H by applying Lemma 13.3.9, and let $\tilde{H} \subseteq H'$ be defined as above. Then the graph $\tilde{H}_{K+1}$ defined as above is K-fat for every $K \in \mathbb{N}$ by the argument above. Similarly, it follows by (the symmetry of) the construction of H' in the proof of Lemma 13.3.9 that also the subgraph $\tilde{H}_{-K-1}$ consisting of all S^i with $i \leqslant -K-1$ and all $P'_{ij} \subseteq \tilde{H}$ with $i \leqslant -K-1$ is K-fat for all $K \in \mathbb{N}$. Since G is locally finite, we then find infinitely many S^{-K-1}–S^{K+1} paths W_j in $\tilde{H}$ which are pairwise at least K apart. By the assumptions on $\tilde{H}$ and since $\tilde{H}$ is escaping, gluing the W_i with $\tilde{H}_{K+1}$ and $\tilde{H}_{-K-1}$ together yields (after possibly applying Proposition 13.3.6) a subdivision H'' of the hexagonal full-grid. By construction, H'' is K-fat. □

13.3.6 Proof of the main results given Theorems 13.3.3 and 13.3.4

In this section we derive Theorems 35 to 38 from Theorems 13.3.3 and 13.3.4; in fact, we show the following more detailed versions.

Theorem 35'. *Let G be a locally finite, quasi-transitive graph whose cycle space is generated by cycles of bounded length. Then $FG \prec_\infty^\varepsilon G$ for every thick end ε of G.*

Proof of Theorem 35 and Theorem 35' given Theorem 13.3.4. Combining Observation 13.3.1 and Lemma 13.3.5 (i), we obtain that Theorem 13.3.4 yields Theorem 35' and hence also Theorem 35, where we note that the subdivision obtained from Lemma 13.3.5 (i) has all its rays in the same end as the full-grid obtained from Theorem 13.3.4. □

Theorem 36'. *Let ε be a thick end of a locally finite, quasi-transitive graph G whose cycle space is generated by cycles of bounded length. Then G contains a diverging subdivision of the hexagonal full-grid whose rays all lie in ε.*

Proof of Theorem 36 and Theorem 36' given Theorem 13.3.4. Combining Observation 13.3.2 and Lemma 13.3.5 (ii), we obtain that Theorem 13.3.4 yields Theorem 36', where we note that the subdivision obtained from Lemma 13.3.5 (ii) has all its rays in the same end as the full-grid obtained from Theorem 13.3.4. For Theorem 36, note that every diverging subdivision of the hexagonal full-grid can be contracted into a diverging minor of the full-grid. □

Theorem 37'. *Let G be graph of finite maximum degree whose cycle space is generated by cycles of bounded length. Then $HG \prec_\infty^\varepsilon G$ for every thick end ε of G.*

Proof of Theorem 37 and Theorem 37' given Theorem 13.3.3. Combining Observation 13.3.1 and Lemma 13.3.5 (i), we obtain that Theorem 13.3.3 yields Theorem 37', and hence also Theorem 37, where we note that the subdivision obtained from Lemma 13.3.5 (i) has all its rays in the same end as the half-grid obtained from Theorem 13.3.3. □

Theorem 38'. *Let ε be a thick end of a graph G of finite maximum degree whose cycle space is generated by cycles of bounded length. Then G contains a diverging subdivision of the hexagonal half-grid whose rays all lie in ε.*

Proof of Theorem 38 and Theorem 38' given Theorem 13.3.3. Combining Observation 13.3.2 and Lemma 13.3.5 (ii), we obtain that Theorem 13.3.3 yields Theorem 38', where we note that the subdivision obtained from Lemma 13.3.5 (ii) has all its rays in the same end as the half-grid obtained from Theorem 13.3.3. For Theorem 38', note that a diverging subdivision of the hexagonal half-grid can easily be contracted into a diverging minor of the half-grid. □

13.4 Diverging double rays and quasi-geodesic 3-stars of rays in thick ends

In this section we prove two theorems about double rays and 3-stars of rays in thick ends, which we need for the proofs of Theorems 35 to 38.

13.4.1 Diverging double rays

Georgakopoulos and Papasoglu [75, Theorem 8.16] showed that every connected graph of finite maximum degree which has an infinite set of pairwise disjoint rays has a diverging double ray (whose tails may lie in two distinct ends). For the proofs of Theorems 37 and 38 we will need the following variant of that theorem, which lets us find the diverging double ray in any thick end we like.

Theorem 13.4.1. *Let G be a graph of finite maximum degree, and let ε be a thick end of G. Then G has a diverging double ε-ray.*

The proof of Theorem 13.4.1 uses the same idea as the one of [75, Theorem 8.16] by Georgakopoulos and Papasoglu, in that Lemma 13.4.3 and Lemma 13.4.4 below are variants of [75, Corollary 8.15 and Lemma 8.17]. However, our proof is more involved, as we need to take care that the tails of the double ray lie in the prescribed end.

Essentially, we will deduce Theorem 13.4.1 from the coarse Menger's theorem for two paths (see Theorem 27 in Chapter 10); the version we state here is from Georgakopoulos and Papasoglu [75, Theorem 8.1] and a little more general than Theorem 27. A *metric graph* is a pair (G, ℓ) of a graph G and an assignment of edge-lengths $\ell\colon E(G) \to \mathbb{R}_{>0}$.

Theorem 13.4.2. *Let G be a metric graph, and let $X, Y \subseteq V(G)$. For every $K > 0$, there is either*

(i) *a set $B \subseteq V(G)$ of diameter $\leqslant K$ such that $G - B$ contains no path joining X to Y, or*

(ii) *two X–Y paths at distance at least $d := K/272$ from each other.*

It follows by a compactness argument that we may replace in Theorem 13.4.2 the set Y by an end ε and the two paths in (ii) by ε-rays.

Lemma 13.4.3. *Let (G, ℓ) be a metric graph such that $B_{(G,\ell)}(v, n)$ is finite for all $v \in V(G)$ and $n \in \mathbb{N}$. Let A be a finite set of vertices in G, and let ε be an end of G. For every $K > 0$, there is either*

(i) *a set $B \subseteq V(G)$ of diameter $\leqslant K$ such that $G - B$ contains no ε-ray starting at A, or*

(ii) *two ε-rays starting at A at distance at least $d := K/272$ from each other.*

Proof. In the following, we abbreviate $d_{(G,\ell)}$ and $B_{(G,\ell)}$ with d_G and B_G, respectively.

By the assumption on ℓ and since A is finite, the balls $B_G(A, n)$, for $n \in \mathbb{N}$, are finite. Hence, there exists, for every $n \in \mathbb{N}$ a unique component C_n of $G - B_G(A, n)$ such that every ε-ray has a tail in C_n. Their neighbourhoods $N_G(C_n) \subseteq B_G(A, n)$ are finite; so we may set $k_n := |N_G(C_n)|$ and enumerate $N_G(C_n) =: \{v_1^n, \ldots, v_{k_n}^n\}$.

We apply Theorem 13.4.2 to the sets $X := A$ and $Y := N_G(C_n)$ in G. If, for some $n \in \mathbb{N}$, Theorem 13.4.2 yields a set $B \subseteq V(G)$ of diameter $\leqslant K$ that separates A and $N_G(C_n)$, then by the definition of C_n this B is as in (i). Hence, we may assume that for every $n \in \mathbb{N}$, we find two A–$N_G(C_n)$ paths P_n, Q_n that are at least $d := K/272$ apart in G.

For all $m \leqslant n \in \mathbb{N}$, we define a k_m-tuple

$$t_n^m = ((t_1^{nm}, \tilde{t}_1^{nm}, s_1^{nm}), \ldots, (t_{k_m}^{nm}, \tilde{t}_{k_m}^{nm}, s_{k_m}^{nm})) \in T_m := (\{0, \ldots, d\}^2 \times \{-4, \ldots, 4\})^{k_m}$$

of triples as follows. We let t_i^{nm} be the distance $d_G(P_n, v_i^m)$ between P_n and v_i^m if it is less than d; otherwise we set $t_i^{nm} := d$. Analogously, we let $\tilde{t}_i^{nm}$ be the distance $d_G(Q_n, v_i^m)$ between Q_n and v_i^m if it is less than d; otherwise we set $\tilde{t}_i^{nm} := d$. Further, if $t_i^{nm} \neq 0$ and $\tilde{t}_i^{nm} \neq 0$, then we set $s_i^{nm} := 0$. Otherwise, it follows that precisely one of P_n and Q_n meets v_i^m, and we then let s_i^{nm} encode whether v_i^m meets P_n ($s_i^{nm} \in \{1, 2, 3, 4\}$) or Q_n ($s_i^{nm} \in \{-4, -3, -2, -1\}$) and whether its predecessor and successor on P_n or Q_n both lie in C_m ($|s_i^{nm}| = 1$), both lie in $G - C_m$ ($|s_i^{nm}| = 2$), or its predecessor lies in C_m and its successor lies in $G - C_m$ ($|s_i^{nm}| = 3$) or vice versa ($|s_i^{nm}| = 4$).

Since all T_m are finite, there exist infinite index sets $\mathbb{N} \supseteq I_0 \supseteq I_1 \supseteq \ldots$ such that, for all $m \in \mathbb{N}$, all t_n^m with $n \in I_m$ are equal. We pick, for every $m \in \mathbb{N}$, some $i_m \in I_m$. Now set

$$\tilde{P}_{i_m} := P_{i_m} \cap (G[C_{m-1}, 1] - C_m) \;\text{ and }\; \tilde{Q}_{i_m} := Q_{i_m} \cap (G[C_{m-1}, 1] - C_m),$$

and let $P := \bigcup_{m \in \mathbb{N}} \tilde{P}_{i_m}$ and $Q := \bigcup_{m \in \mathbb{N}} \tilde{Q}_{i_m}$. We claim that P and Q are at least d apart in G and that they both contain an ε-ray that starts in A. It then follows that these rays are as in (ii).

First, we show that P and Q are at least d apart in G. For this, recall that $d_G(P_n, Q_n) \geqslant d$ for all $n \in \mathbb{N}$ by the choice of P_n, Q_n. Now let $m \leqslant n \in \mathbb{N}$ be given. We show that $d_G(\tilde{P}_{i_m}, \tilde{Q}_{i_n}) \geqslant d$; the other case is symmetric. Clearly, if $m = n$, then $d_G(\tilde{P}_{i_m}, \tilde{Q}_{i_n}) \geqslant d$ holds by the choice of P_{i_m}, Q_{i_n}, so we may assume that $m < n$. Set $\ell := d_G(\tilde{P}_{i_m}, \tilde{Q}_{i_n})$, and let $W = w_0 \ldots w_\ell$ be a shortest $\tilde{P}_{i_m}$–$\tilde{Q}_{i_n}$ path. Then W meets $N_G(C_m)$ in a vertex v_j^m because $\tilde{Q}_{i_n} \subseteq G[C_{n-1}, 1] \subseteq G[C_m, 1]$ and

$\tilde{P}_{i_m} \subseteq G - C_m$ as $m < n$. It follows that

$$\begin{aligned}\ell &= d_G(w_0, v_j^m) + d_G(v_j^m, w_\ell) \geqslant d_G(\tilde{P}_{i_m}, v_j^m) + d_G(v_j^m, \tilde{Q}_{i_n}) \geqslant d_G(P_{i_m}, v_j^m) + d_G(v_j^m, Q_{i_n}) \\ &\geqslant t_j^{i_m m} + \tilde{t}_j^{i_n m} = t_j^{i_n m} + \tilde{t}_j^{i_n m} = \min\{d_G(P_{i_n}, v_j^m), d\} + \min\{d_G(v_j^m, Q_{i_n}), d\} \\ &\geqslant \min\{d_G(P_{i_n}, Q_{i_n}), d\} = d.\end{aligned}$$

where we used $t_j^{i_m m} = t_j^{i_n m}$ since $i_m, i_n \in I_m$. Hence, $d_G(P, Q) \geqslant d$ as desired.

So to conclude the proof, it remains to show that P and Q both contain an ε-ray that starts at A. We show the claim for P; the other case is symmetric. By definition, it is clear that P meets A in a unique vertex a, which is the endvertex of P_{i_1} in A; in particular, a has degree 1 in P. Hence, it suffices to show that all other vertices in P have degree 2 in P, as then P contains a ray that starts in a, and which then has to lie in ε since $P - C_m = \bigcup_{n \leqslant m} \tilde{P}_{i_n}$ is finite for all $m \in \mathbb{N}$. By definition of P, every vertex of P that is not contained in some $N_G(C_m)$ is contained in precisely one $\tilde{P}_{i_m}$, and has thus degree 2 in P. So let some $v_j^m \in V(P) \cap N_G(C_m)$ be given. Then $s_j^{i_m m} \in \{1, 2, 3, 4\}$, and it follows that v_j^m has degree 2 in G because $s_j^{i_m m} = s_j^{i_{m+1} m}$ by the choice of i_m, i_{m+1}. □

The remainder of the proof of Theorem 13.4.1 is now analogous to the one of [75, Theorem 8.16]. More precisely, we have the following auxiliary lemma.

Lemma 13.4.4. *Let G be a graph of finite maximum degree, and let ε be a thick end of G. Then there is a finite set A of vertices in G and an assignment of edge-lengths $\ell \colon E(G) \to \mathbb{R}_{>0}$ with the following properties:*

(i) *no ball of radius 1 in the corresponding metric d_ℓ separates A from ε,*

(ii) $\lim_{e \in E(G)} \ell(e) = 0$*, and*

(iii) *every ball of finite radius in d_ℓ is finite.*

Proof. The proof is analogous to [75, Lemma 8.17] with just one exception: we choose the sequence $(R_n)_{n \in \mathbb{N}}$ of pairwise disjoint rays so that every R_n is an ε-ray, which is possible because ε is thick.

Note that (iii) follows easily from the proof, since the S^n are 'thickened rings' $B_G(o, r_n) \setminus B_G(o, r_{n-1})$ around a vertex $o \in V(G)$ and because $\sum_{i \in \mathbb{N}} 1/n$ is infinite. □

Proof of Theorem 13.4.1. The proof is analogous to [75, Theorem 8.16] with just one exception: instead of [75, Corollary 8.15 & Lemma 8.17] we apply Lemmas 13.4.3 and 13.4.4. □

13.4.2 Quasi-geodesic 3-stars of rays

By Theorem 13.4.1 every graph G of finite maximum degree contains for every thick end ε a diverging double ε-ray. For the proofs of Theorems 35 and 36, we need the following result, which strengthens Theorem 13.4.1 in the special case where G is quasi-transitive and accessible.

Theorem 13.4.5. *Let ε be a thick end of a locally finite, accessible, quasi-transitive graph G. Then there exists $c \in \mathbb{N}_{\geqslant 1}$ and ε-rays R_1, R_2, R_3 in G such that $R_1 \cap R_2 = R_1 \cap R_3 = R_2 \cap R_3 = \{v\}$ for some $v \in V(G)$ and such that $R_1 \cup R_2 \cup R_3$ is c-quasi-geodesic in G.*

For the proof of Theorem 13.4.5, we first need the following auxiliary lemma.

Lemma 13.4.6. *Let G be a locally finite, accessible, quasi-transitive graph that contains a thick end ε. Then there exists a connected, one-ended, quasi-geodesic subgraph H of G such that every ray in H is an ε-ray in G and such that the stabilizer of H acts quasi-transitively on H.*

Proof. By a result of Diestel, Jacobs, Knappe and Kurkofka [48, Lemma 7.12] and in particular its proof [49, Appendix A], there exists a connected, induced, one-ended subgraph H of G whose rays all lie in ε such that every component of $G - H$ has finite neighbourhood in H, such that there are only finitely many orbits of such components under the stabilizer Γ of H in the automorphism group of G and such that Γ acts quasi-transitively on H. It remains to prove that H is quasi-geodesic. Since there are only finitely many orbits of components of $G - H$ under Γ, each such component has finite neighbourhood in H and because H is connected, there exists $c \in \mathbb{N}_{\geqslant 1}$ such that for every component C of $G - H$ every two vertices in $N_G(C)$ have distance at most c in H.

We claim that H is c-quasi-geodesic. Indeed, let $x, y \in V(H)$ be given, and let P be a shortest x–y path in G. Further, let $Q_0, \dots, Q_m$ be the maximal non-trivial subpaths of P that are internally disjoint from H. Then every Q_i is internally contained in some component C of $G - H$ and starts and ends in $N_G(C)$. By the choice of c, there exists a path Q_i' in H of length at most c which has the same endvertices as Q_i. It follows that the union W over $P \cap H$ and the Q_m' is connected, contained in H, contains x, y, and hence contains an x–y path. Since all Q_i are non-trivial and have thus length at least 1, it follows that $d_H(x,y) \leqslant |E(W)| \leqslant |E(P \cap H)| + c(m+1) \leqslant c \cdot |E(P)| = c \cdot d_G(x,y)$ as desired. □

We also need Lemma 11.4.3 from Chapter 11, which we recall here for convenience.

Lemma 13.4.7. *Let X be a c-quasi-geodesic subgraph of some graph G for some $c \in \mathbb{N}_{\geqslant 1}$. If P is a shortest v–X path in G for some vertex $v \in V(G)$, then $X \cup P$ is $(2c+1)$-quasi-geodesic in G.*

We can now prove Theorem 13.4.5.

Proof of Theorem 13.4.5. Let us first assume that G is one-ended. In this case, we apply two compactness arguments. First, a standard compactness argument (see e.g. [123, Proposition 5.2]) implies the existence of a geodesic double ray R in the locally finite and quasi-transitive graph G.

For the second compactness argument, we first show that $G[R, K] \neq G$ for all $K \in \mathbb{N}$. Since G is locally finite, the set $B_G(r_1 R r_{2K+1}, K)$ is finite. As R is geodesic, the sets $B_G(R_{\leqslant 0}, K)$ and

$B_G(R_{\geqslant 2K+2}, K)$ are disjoint and not joined by an edge. Hence, every $B_G(R_{\leqslant 0}, K)$–$B_G(R_{\geqslant 2K+2}, K)$ path meets either $B_G(r_1Rr_{2K+1}, K)$ or $G - G[R, K]$. But since both $R_{\leqslant 0}$ and $R_{\geqslant 2K+2}$ lie in the unique end of G, there are infinitely many disjoint such paths, of which at most finitely many can meet the finite set $B_G(r_1Rr_{2K+1}, K)$. Hence, $G - G[R, K]$ is non-empty.

Thus, there exists vertices in G of arbitrary distance from R. Let x_i be a vertex at distance i from R, let r_{j_i} be a vertex of R with $d_G(x_i, r_{j_i}) = d_G(x_i, R)$, and let $P_i = p_0^i \dots p_i^i$ be a shortest x_i–r_{j_i} path. Then $R \cup P_i$ is 3-quasi-geodesic by Lemma 13.4.7. Since G is quasi-transitive, there is an infinite index set $I \subseteq \mathbb{N}$ such that all r_{j_i} lie in the same orbit. Let $s \in V(G)$ be another vertex in that orbit. For all $i \in I$, let φ_i be an automorphism of G that maps r_{j_i} to s. Then, since G is locally finite, there exists an infinite index set $I_1 \subseteq I$ such that $\varphi_i(r_{j_i-1}Rr_{j_i+1} \cup p_0^iPp_1^i)$ coincides for all $i \in I_1$, amongst which we again find an infinite index set $I_2 \subseteq I_1$ such that $\varphi_i(r_{j_i-2}Rr_{j_i+2} \cup p_0^iPp_2^i)$ coincides for all $i \in I_2$ and so on. This results in three internally disjoint, geodesic rays starting in s whose union is 3-quasi-geodesic. Obviously, all three rays must lie in the unique end ε of G, so they are as desired.

Let us now assume that G has more than one end. Since G is accessible, there exists by Lemma 13.4.6 a connected, one-ended, c-quasi-geodesic, quasi-transitive subgraph H of G for some $c \in \mathbb{N}_{\geqslant 1}$ such that every ray in H is an ε-ray in G. By the first case, we find the desired three rays R_1, R_2, R_3 in H whose union is 3-quasi-geodesic. Since H is a c-quasi-geodesic subgraph of G, the rays R_1, R_2, R_3 form a $3c$-quasi-geodesic subgraph of G. □

13.5 Half-grid minors

In this section we prove Theorem 13.3.3; in fact, we show a more detailed version, which we need in the next section for the proof of Theorem 13.3.4.

Let $R = \dots r_{-1}r_0r_1 \dots$ be a double ray in a graph G, and let $K \in \mathbb{N}$. A component C of $G - B_G(R, K)$ is *long* if C has a neighbour in $B_G(R_{\geqslant i}, K)$ and in $B_G(R_{\leqslant -i}, K)$ for all $i \in \mathbb{N}$. Further, C is *thick* if, for every $L \geqslant K$, some long component of $G - B_G(R, L)$ is contained in C.

Theorem 13.3.3'. *Let R be a diverging double ray in a thick end ε of a locally finite graph G whose cycle space is generated by cycles of bounded length. Then either $K_{\aleph_0} \prec_{UF}^{\varepsilon} G$ or G contains an escaping subdivision H of the hexagonal half-grid whose first vertical ray is R.*

In particular, if $K_{\aleph_0} \not\prec_{UF}^{\varepsilon} G$ and C is a thick component of $G - B_G(R, L)$ for some $L \in \mathbb{N}$, then we may choose the vertical double rays S^i of H so that $S^i \subseteq C$ for all $i \geqslant 1$.

Proof of Theorem 13.3.3 given Theorem 13.3.3'. By Theorem 13.4.1, there exists a diverging double ε-ray R in G. Apply Theorem 13.3.3 to R. □

In the remainder of this section we prove Theorem 13.3.3'; see Section 13.3.3 for a sketch of the proof.

Lemma 13.5.1. *Let G be a graph whose cycle space is generated by cycles of length at most $\kappa \in \mathbb{N}$, and let Y be a connected subgraph of G. Then for every component C of $G - Y$ that attaches to Y, the graph $C[\partial_G C, \lfloor \frac{\kappa-2}{2} \rfloor]$ is connected.*

Note that if G is connected, then every component of $G - Y$ attaches to Y.

Proof. Clearly, it suffices to show for every two vertices $v_0, v_1 \in \partial_G C$ that there exists a v_0–v_1 path in $C[\partial_G C, \lfloor \frac{\kappa-2}{2} \rfloor]$. So let $v_0, v_1 \in \partial_G C$ be given, and let u_0 and u_1 be vertices of Y which are adjacent to v_0 and v_1, respectively. Since Y is connected, there exists a u_1–u_0 path Q in Y. Let P be a v_0–v_1 path in C. Then $D := v_0 P v_1 u_1 Q u_0 v_0$ is a cycle in G. By assumption on the cycle space of G, we can write D as a finite sum of cycles $D_1, \dots, D_n$ in G of length at most κ, i.e.

$$D = \sum_{D_i \in \mathcal{D}} D_i$$

where $\mathcal{D} := \{D_1, \dots, D_n\}$. Let $\mathcal{D}' \subseteq \{D_1, \dots, D_n\}$ consist of those D_i that do not lie completely in C, i.e. that contain a vertex of $G - C$. Note that $D_i \cap C \subseteq C[\partial_G C, \lfloor \frac{\kappa-2}{2} \rfloor]$ for all $D_i \in \mathcal{D}'$ since D_i has length at most κ and meets $G - C$. Let

$$H := \bigcup_{D_i \in \mathcal{D}'} D_i \cap C \subseteq C\left[\partial_G C, \left\lfloor \frac{\kappa - 2}{2} \right\rfloor\right] \subseteq C$$

be the subgraph of C consisting of all vertices and edges in C that lie on cycles from $\mathcal{D}'$. Note that $v_0, v_1 \in V(H)$ since $v_0u_0, v_1u_1 \in E(D)$. We claim that v_0 and v_1 lie in the same component of H, which clearly yields the claim. So suppose for a contradiction that v_0 and v_1 lie in distinct components H_0, H_1 of H. Then the set F of edges in G between H_0 and $G - H_0$ is a cut in G that separates H_0 and H_1; in particular, F is finite since G is locally finite and because $H_0 \subseteq H$ is finite as $\bigcup_{D_i \in \mathcal{D}'} D_i \supseteq H_0$ is a finite union of finite cycles. Obviously, F must contain an edge f from $P \subseteq C$. Then f cannot lie in $\sum_{D_i \in \mathcal{D}'} D_i \subseteq \bigcup_{D_i \in \mathcal{D}'} D_i$, since $f \in E(C)$ but $f \notin E(H)$. Hence, as $f \in E(P) \subseteq E(D)$, it lies in

$$H' := D + \sum_{D_i \in \mathcal{D}'} D_i = \sum_{D_i \in \mathcal{D}} D_i + \sum_{D_i \in \mathcal{D}'} D_i = \sum_{D_i \in \mathcal{D} \smallsetminus \mathcal{D}'} D_i \subseteq C,$$

where for the last inclusion we used that $D_i \subseteq C$ for all $D_i \in \mathcal{D} \smallsetminus \mathcal{D}'$ by the choice of $\mathcal{D}'$. In particular, the same argument also yields that $E(P) \cap F \subseteq E(H')$.

As H' is a finite sum of cycles in G, it is an element of the cycle space of G. Thus, H' meets the finite cut F in an even number of edges. As P is a finite path from $v_0 \in V(H_0)$ to

$v_1 \in V(H_1) \subseteq V(G - H_0)$, it meets the finite cut F in an odd number of edges. Combining these two facts with $E(P) \cap F \subseteq E(H')$ yields that H' contains an edge $f' \neq f$ from F which does not lie on P. Since $H' \subseteq C$, the edge f' must lie in C. But since f' is not an edge of $P = D \cap C$, it is not an edge of D either. Hence, f' is an edge of $\sum_{D_i \in \mathcal{D}'} D_i$, and thus an edge of $\bigcup_{D_i \in \mathcal{D}'} D_i$. Since f' is also an edge of C, it lies in H, which is a contradiction to the choice of F. □

Lemma 13.5.2. *Let $R = \dots r_{-1} r_0 r_1 \dots$ be a diverging double ray in an end ε of a locally finite graph G. Then for every $K, n \in \mathbb{N}$ some component of $G - B_G(R, K)$ attaches to $B_G(R_{\leqslant -n}, K)$ and $B_G(R_{\geqslant n}, K)$.*

Proof. Since R diverges, there exists some $m \in \mathbb{N}$ such that $R_{\leqslant -m}$ and $R_{\geqslant m}$ are at least $2K+2$ apart in G. Set $N := \max\{n, m\}$. As G is locally finite, the set $B_G(r_{-N} R r_N, K)$ is finite. Hence, as $R_{\leqslant -N}$ and $R_{\geqslant N}$ are both ε-rays and thus equivalent, there exists an $R_{\leqslant -N}$–$R_{\geqslant N}$ path $P = p_0 \dots p_\ell$ in G that avoids $B_G(r_{-N} R r_N, K)$. Since $p_0 \in V(R_{\leqslant -N})$ and $p_\ell \in V(R_{\geqslant N})$, there is a first vertex p_i of P that is contained in $B_G(R_{\geqslant N}, K)$, and a last vertex p_j with $j \leqslant i$ that is still contained in $B_G(R_{\leqslant -N}, K)$.

We claim that $i \geqslant j+2$, which then implies that $P' := p_{j+1} P p_{i-1}$ is non-empty. As P avoids $B_G(r_{-N} R r_N, K)$ and by the choice of p_i and p_j, it then follows that P' is contained in a component of $G - B_G(R, K)$, which then attaches to $B_G(R_{\leqslant -N}, K)$ and $B_G(R_{\geqslant N}, K)$ via $p_j p_{j+1}$ and $p_{i-1} p_i$, respectively, and which is thus as desired.

So suppose for a contradiction that $i - j \leqslant 1$. Then $d_G(p_j, p_i) \leqslant i - j \leqslant 1$, and thus

$$d_G(R_{\leqslant -N}, R_{\geqslant N}) \leqslant d_G(R_{\leqslant -N}, p_j) + d_G(p_j, p_i) + d_G(p_i, R_{\geqslant N}) \leqslant K + 1 + K = 2K+1,$$

which is a contradiction since $d_G(R_{\leqslant -N}, R_{\geqslant N}) \geqslant 2K+2$ by the choice of N. □

Lemma 13.5.3. *Let R be a diverging double ray in a thick end of a locally finite graph G whose cycle space is generated by cycles of bounded length. Then for every $K \in \mathbb{N}$ some component of $G - B_G(R, K)$ is long.*

Proof. Let $\kappa \in \mathbb{N}$ such that the cycle space of G is generated by cycles of length at most κ. Suppose for a contradiction that no component of $G - B_G(R, K)$ is long. Since R diverges, there exists some $N \in \mathbb{N}$ such that $d_G(R_{\leqslant -N}, R_{\geqslant N}) \geqslant 2K + \kappa + 2$. As G is locally finite, $B_G(r_{-N} R r_N, K)$ is finite. Hence, the set $\mathcal{C}$ of components of $G - B_G(R, K)$ which attach to $B_G(r_{-N} R r_N, K)$ is finite. Since no $C \in \mathcal{C}$ is long by assumption and because $\mathcal{C}$ is finite, there exists some $m \in \mathbb{N}$ such that $N_G(C) \subseteq B_G(R_{\geqslant -m}, K)$ or $N_G(C) \subseteq B_G(R_{\leqslant m}, K)$ for all $C \in \mathcal{C}$.

By Lemma 13.5.2, some component C of $G - B_G(R, K)$ attaches to $B_G(R_{\leqslant -m-1}, K)$ and $B_G(R_{\geqslant m+1}, K)$; in particular, $C \notin \mathcal{C}$ by the choice of m. Let $U^-, U^+ \subseteq \partial_G C$ be the set of vertices in C that send an edge to $B_G(R_{\leqslant -N}, K)$ or to $B_G(R_{\geqslant N}, K)$, respectively. Then $U^- \cup U^+ = \partial_G C$ because $C \notin \mathcal{C}$. Since $C[\partial_G C, \lfloor\frac{\kappa-2}{2}\rfloor]$ is connected by Lemma 13.5.1, this implies that $B_C(U^-, \lfloor\frac{\kappa-2}{2}\rfloor)$ and $B_C(U^+, \lfloor\frac{\kappa-2}{2}\rfloor)$ either intersect non-emptily or there is an edge between them. Hence, there are vertices $u^- \in U^-$ and $u^+ \in U^+$ of distance at most $\lfloor\frac{\kappa-2}{2}\rfloor + 1 + \lfloor\frac{\kappa-2}{2}\rfloor$ from each other. Thus,

$$\begin{aligned} d_G(R_{-N}, R_N) &\leqslant d_G(R_{-N}, u^-) + d_G(u^-, u^+) + d_G(u^+, R_N) \\ &\leqslant (K+1) + \left(\left\lfloor\frac{\kappa-2}{2}\right\rfloor + 1 + \left\lfloor\frac{\kappa-2}{2}\right\rfloor\right) + (K+1) \\ &\leqslant 2K + \kappa + 1 \end{aligned}$$

which is a contradiction since $d_G(R_{-N}, R_N) \geqslant 2K + \kappa + 2$ by the choice of N. □

Lemma 13.5.4. *Let ε be an end of a locally finite graph G. Suppose there are $M_0 < M_1 < \ldots \in \mathbb{N}$ and double ε-rays $S^0, S^1, \ldots$ such that S^0 diverges, such that $S^i \subseteq G[S^0, M_i] - B_G(S^0, M_{i-1})$ for all $i \in \mathbb{N}_{\geqslant 1}$ and such that there are infinitely many disjoint $S^0_{\geqslant 0}$–$S^i_{\geqslant 0}$ paths and infinitely many disjoint $S^0_{\leqslant 0}$–$S^i_{\leqslant 0}$ paths in $G[S^0, M_i]$. Then either $K_{\aleph_0} \prec^{\varepsilon}_{UF} G$, or there are $0 = i_0 < i_1 < \ldots \in \mathbb{N}$ and an escaping subdivision H of the hexagonal half-grid whose vertical double rays are the S^{i_j}.*

Proof. By passing to a subsequence of the S^i if necessary, we may assume that $M_i > M_{i-1} + 2i$ and that

(a) $S^i \subseteq G[S^0, M_i] - B_G(S^0, M_{i-1} + 2i)$ for all $i \in \mathbb{N}$.

Set $T'_0 := S^0_{\geqslant 0}$ and $T''_0 := S^0_{\leqslant 0}$. By assumption and Lemma 13.3.7, every S^i has disjoint tails T'_i and T''_i that are contained in $G[T'_0, M_i]$ and in $G[T''_0, M_i]$, respectively. For each vertex t in T'_i we choose a shortest t–T'_0 path in G, which then has length $\leqslant M_i$ and lies in $G[T'_0, M_i]$. Then infinitely many of these paths are T'_i–T'_0 paths (i.e. they only have their first vertex on T'_i), of which infinitely many are pairwise disjoint since they have length $\leqslant M_i$ and because G is locally finite; let us denote these paths by Q_{ij}.

For every Q_{ij} let $k_{ij} \neq i$ be maximal such that $d_G(Q_{ij}, T'_{k_{ij}}) < k_{ij}$; if no such k_{ij} exists, we set $k_{ij} := 0$. Note that $k_{ij} < i$ since $d_G(Q_{ij}, T'_k) \geqslant d_G(Q_{ij}, S^k) \geqslant k$ for all $k > i$ by (a) and because $Q_{ij} \subseteq G[T'_0, M_i]$ for all $i, j \in \mathbb{N}$. We now obtain T'_i–$T'_{k_{ij}}$ paths Q'_{ij} by concatenating a suitable (initial) subpath of Q_{ij} with a shortest Q_{ij}–$T'_{k_{ij}}$ path. In particular, since G is locally finite, we may assume that the Q'_{ij} for every (arbitrary but fixed) $i \in \mathbb{N}_{\geqslant 1}$ are pairwise disjoint. By the choice of the k_{ij} it follows that

(b) $d_G(Q'_{ij}, T'_k) \geqslant k$ for all $k, j \in \mathbb{N}$ and $i \in \mathbb{N}_{\geqslant 1}$ with $k \notin \{i, k_{ij}\}$.

Moreover, since Q_{ij} is a shortest path between its first vertex and T'_0, it follows that once Q_{ij} meets $G[T'_0, M_{k_{ij}-1} + i]$ it will stay in there. By the definition of Q'_{ij} and k_{ij} and by (a), this implies that

(c) $Q'_{ij} \subseteq G[T'_0, M_i] - B_G(T'_0, M_{k_{ij}-1} + i)$ for all $i \in \mathbb{N}_{\geqslant 1}$ and $j \in \mathbb{N}$.

Let X be the auxiliary graph on the vertex set $\{T'_i \mid i \in \mathbb{N}\}$ where T'_i and $T'_{i'}$ are connected by an edge in X for $i' < i$ if and only if infinitely many of the Q'_{ij} have one endvertex on $T'_{i'}$. Clearly, every T'_i is adjacent to at least one $T'_{i'}$ with $i' < i$, and hence X is connected. Thus, since X is infinite, it either has a vertex of infinite degree or it contains a ray by Lemma 2.5.1.

Let us first assume that there is some $\ell \in \mathbb{N}$ and an infinite subset $I = \{i_0, i_1, \dots\} \subseteq \mathbb{N}$ with $i_0 < i_1 < \dots$ such that T'_ℓ is adjacent in X to all T'_i with $i \in I$. Then the T'_i for $i \in I$ form the branch sets $V_n := V(T'_{i_n})$ of an ultra fat model of $K_{\aleph_0}$ in G. Indeed, we have $d_G(V_n, V_m) \geqslant \min\{n, m\}$ by (a), so it remains to find suitable branch paths. Given any enumeration of $\mathbb{N}^2$, we may choose the branch paths P_{nm} between V_n and V_m recursively. Since G is locally finite, and because at step (n, m) we have only chosen finitely many branch paths $P_{n'm'}$, there exist paths Q'_{i_nj}, $Q'_{i_mj'}$ that both end in T'_ℓ such that the path P_{nm} consisting of $Q'_{i_nj}, Q'_{i_mj'}$ and a suitable subpath of T'_ℓ is at least $\min\{n, m\}$ apart from all earlier chosen branch paths $P_{n'm'}$. Then by construction and (a) and (b) it follows that also $d_G(P_{nm}, V_k) \geqslant k$ for all $k \notin \{n, m\}$, and hence the model of $K_{\aleph_0}$ is ultra fat. Moreover, since its branch sets are the vertex sets of the ε-rays T'_{i_n}, we find $K_{\aleph_0} \prec^{\varepsilon}_{UF} G$.

Hence, we may assume that there are $0 = i_0 < i_1 < \ldots \in \mathbb{N}$ such that every T'_{i_n} is adjacent in X to $T'_{i_{n-1}}$. Then there are, for every $n \in \mathbb{N}$, infinitely many Q'_{i_nj} that end in $T'_{i_{n-1}}$. We reindex these Q'_{i_nj} by $\mathbb{N}_{\geqslant 1} \times \mathbb{N}$, and the S^{i_n} by $\mathbb{N}$.

We now apply the same argument to the tails T''_i of the (reindexed) S^i. This either yields $K_{\aleph_0} \prec^{\varepsilon}_{UF} G$, or we find indices $0 = i_0 < i_1 < \ldots \in \mathbb{N}$ such that, for every $n \in \mathbb{N}$, there are pairwise disjoint $T''_{i_{n-1}}$–T''_{i_n} paths Q''_{i_nj} that satisfy (b) and (c) with T''_k instead of T'_k. Now the S^{i_n} form the vertical double rays of an escaping subdivision of the hexagonal half-grid. Indeed, we can choose for every $n \in \mathbb{N}$ infinitely many $T'_{i_{n-1}}$–T'_{i_n} paths P'_{i_nj} in $\bigcup_{i_{n-1} < k \leqslant i_n} (T'_k \cup \bigcup_{j \in \mathbb{N}} Q'_{kj})$. We also set $P'_{i_n(-j)} := Q''_{i_nj}$ for all $j \in \mathbb{N}$. Then combining (c) of the Q'_{i_nj} and the Q''_{i_nj} with Lemma 13.3.7 yields that

(c') $P'_{i_nj} \subseteq G[S^0, M_i] - B_G(S^0, M_{i_n-1} + i)$

for every $i \in \mathbb{N}_{\geqslant 1}$ and all but finitely $j \in \mathbb{N}$. Since G is locally finite, we now obtain a subdivision of the hexagonal half-grid with vertical double rays the S^{i_n} by recursively selecting paths P_{nj} amongst the P'_{i_nk} to represent the horizontal edges e_{nj} (compare Proposition 13.3.6). □

We are now ready to prove Theorem 13.3.3'.

Proof of Theorem 13.3.3'. Let $\kappa \in \mathbb{N}$ such that the cycle space of G is generated by cycles of length at most κ. Let N_n, for $n \in \mathbb{N}$, be such that $d_G(R_{\geqslant N_n}, R_{\leqslant -N_n}) > n$, which exists since R diverges.

By Lemma 13.5.4, it suffices to show that there are $M_0 < M_1 < \ldots \in \mathbb{N}$ and diverging double ε-rays $R := S^0, S^1, \dots$ such that $S^i \subseteq G[S^0, M_i] - B_G(S^0, M_{i-1})$ for all $i \in \mathbb{N}_{\geqslant 1}$ and such that there are infinitely many disjoint $S^0_{\geqslant 0}$–$S^i_{\geqslant 0}$ paths and infinitely many disjoint $S^0_{\leqslant 0}$–$S^i_{\leqslant 0}$ paths in $G[S^0, M_i]$. We will prove the assertion with $M_0 := 0$ and $M_i := M_0 + \lfloor \kappa/2 \rfloor$ for all $i > 0$.

Let $i > 0$ be given. By Lemma 13.5.3, there exists a long component C_i of $G - B_G(R, M_{i-1})$. For the 'in particular' part we note that if we are given some $L \in \mathbb{N}$ and a thick component C of $G - B_G(R, L)$, then we may set $M_0 := L$ instead of $M_0 := 0$ and choose as C_i always a long component of $G - B_G(R, M_{i-1})$ which is contained in C.

Set $U^+ := \partial_G C_i \cap B_G(R_{\geqslant 0}, M_{i-1} + 1)$ and $U^- := \partial_G C_i \cap B_G(R_{\leqslant 0}, M_{i-1} + 1)$. Since G is locally finite and C_i is long, U^+ and U^- are infinite. As $C_i[\partial_G C_i, \lfloor \frac{\kappa-2}{2} \rfloor]$ is connected by Lemma 13.5.1, applying the Star-Comb Lemma (cf. Lemma 2.5.1) in $C_i[\partial_G C_i, \lfloor \frac{\kappa-2}{2} \rfloor]$ to U^+ and U^-, respectively, yields two combs D^+ and D^-. By Lemma 13.3.7, their spines S^+ and S^- are eventually contained in $G[R_{\geqslant N_{M_i}}, M_i]$ and $G[R_{\leqslant -N_{M_i}}, M_i]$, respectively, i.e. they have tails T^+, T^- such that $T^+ \subseteq G[R_{\geqslant N_{M_i}}, M_i]$ and $T^- \subseteq G[R_{\leqslant -N_{M_i}}, M_i]$. In particular, T^+, T^- are disjoint by the choice of N_{M_i}, so we can link them by a path in the connected $C_i[\partial_G C_i, \lfloor \frac{\kappa-2}{2} \rfloor]$ to obtain a double ray $S^i \subseteq C_i[\partial_G C_i, \lfloor \frac{\kappa-2}{2} \rfloor]$. Clearly, S^i is as desired. Indeed, the infinitely many $S^0_{\geqslant 0}$–$S^i_{\geqslant 0}$ paths can be obtained by extending the paths in D^+ from T^+ to its teeth by shortest paths to S^0, and analogously for T^-. □

13.6 Full-grid minors

In this section we prove Theorem 13.3.4, which we restate here for convenience.

Theorem 13.3.4. *Let ε be a thick end of a locally finite, quasi-transitive graph G whose cycle space is generated by cycles of bounded length. Then either $K_{\aleph_0} \prec^{\varepsilon}_{UF} G$ or G contains an escaping subdivision of the hexagonal full-grid whose rays all lie in ε.*

In fact, we will prove the following variant of Theorem 13.3.4, which implies Theorem 13.3.4:

Theorem 13.6.1. *Let G be a locally finite, quasi-transitive graph whose cycle space is generated by cycles of bounded length. If G has a thick end, then either $K_{\aleph_0} \prec_{UF} G$ or G contains an escaping subdivision of the hexagonal full-grid.*

We remark that instead of adding Theorem 13.6.1 as an intermediate step in the proof of Theorem 13.3.4 we could have also formulated the three lemmas below which we use to construct either an ultra fat model of $K_{\aleph_0}$ or an escaping subdivision of the hexagonal full-grid (Lemmas 13.6.2, 13.6.8 and 13.6.9) so that we may choose a thick end ε, and the lemma then returns the desired structure 'in' ε. However, while this would have been possible in Lemmas 13.6.2 and 13.6.9 without changing their proofs, this is not true for Lemma 13.6.8. There, we would then have to use the fact that G is accessible, to reduce the problem to one-ended graphs.

Proof of Theorem 13.3.4 given Theorem 13.6.1. By Theorem 13.2.1, G is accessible, so we may apply Lemma 13.4.6 to G and ε. Hence, there exists a connected, one-ended, c-quasi-geodesic, quasi-transitive subgraph X of G for some $c \in \mathbb{N}$ such that every ray in X is an ε-ray in G. Applying Theorem 13.6.1 to X yields either an ultra fat model $((V_i)_{i\in\mathbb{N}}, (E_{ij})_{i\neq j\in\mathbb{N}})$ of $K_{\aleph_0}$ in X or an escaping subdivision H of the hexagonal full-grid in X. Since X is a c-quasi-geodesic subgraph of G, we find in the former case that $((V_i)_{i\in c\mathbb{N}}, (E_{ij})_{i\neq j\in c\mathbb{N}})$ is an ultra fat model of $K_{\aleph_0}$ in G. In fact, since every ray in X is an ε-ray in G, we have $K_{\aleph_0} \prec^{\varepsilon}_{UF} G$. Similarly, in the latter case, H contains a subdivision of the hexagonal full-grid which is escaping in G. Indeed, we may choose vertical double rays $\dots, S^{i_{-1}}, S^{i_0}, S^{i_1}, \dots$ of H such that $i_0 = 0$ and $M_{i_j-1} + 2i_j \geqslant c(M_{i_{j-1}} + 2j)$ for $j > 0$, and similarly $M_{i_j+1} + 2|i_j| \geqslant c(M_{i_{j+1}} + 2|j|)$ for $j < 0$. Then adding suitable S^{i_j}–$S^{i_{j+1}}$ paths in H yields a subdivision $H' \subseteq H$ of the hexagonal full-grid which is escaping in G since X is a c-quasi-geodesic subgraph of G. Moreover, all rays in H' are ε-rays. □

In the remainder of this section we prove Theorem 13.6.1. The formal proof of Theorem 13.6.1, which collects the tools from this whole section, can be found at the end of the last subsection, Section 13.6.2.

We first give a brief overview of this section; a more detailed sketch of the proof of Theorem 13.6.1 can be found in Section 13.3.4. Let G be a locally finite, quasi-transitive graph whose cycle space is generated by cycles of bounded length and which has a thick end. Further, let $R = \dots r_{-1}r_0r_1\dots$ be a double ray, and let $K \in \mathbb{N}$. Recall that a component C of $G - B_G(R, K)$ is *long* if C has a neighbour in $B_G(R_{\geqslant i}, K)$ and in $B_G(R_{\leqslant -i}, K)$ for all $i \in \mathbb{N}$. Further, C is *thick* if, for every $L \geqslant K$, some long component of $G - B_G(R, L)$ is contained in C.

In Lemma 13.6.2 below, we show that if G contains a diverging double ray R such that, for some $L \in \mathbb{N}$, $G - B_G(R, L)$ has at least two thick components, then either $K_{\aleph_0} \prec_{UF} G$ or G contains an escaping subdivision of the hexagonal full-grid. Our remaining task then is to prove that G indeed contains such a double ray R. Showing this will be the main effort of this proof, and it will be done in Section 13.6.2 (see Lemmas 13.6.8 and 13.6.9). For this, in Section 13.6.1, we provide with Lemma 13.6.5 a sufficient condition for G to contain $K_{\aleph_0}$ as an ultra fat minor, which enables us to find an ultra fat $K_{\aleph_0}$ minor in G if we cannot find such a double ray R.

Lemma 13.6.2. *Let R be a diverging double ray in a locally finite graph G whose cycle space is generated by cycles of bounded length. Suppose that for some $L \in \mathbb{N}$, there are at least two thick components of $G - B_G(R, L)$. Then either $K_{\aleph_0} \prec_{UF} G$, or G contains an escaping subdivision of the hexagonal full-grid.*

Proof. Let $C \neq D$ be two distinct thick components of $G - B_G(R, L)$. Since we are done if $K_{\aleph_0} \prec_{UF} G$, we may assume that applying Theorem 13.3.3' to R, K and C or D, respectively, yields escaping subdivisions H_C and H_D, respectively, of the hexagonal half-grid. Let $M^C_0 < M^C_1 < \dots$

and $M_0^D < M_1^D < \dots$ witness that H^C and H^D are escaping. Further, let S_i^C, S_i^D and P_{ij}^C, P_{ij}^D be the vertical double rays and horizontal paths of H_C and H_D, respectively. By the 'in particular' part of Theorem 13.3.3', it follows that the S_i^C and the S_i^D are contained in C and D, respectively; in particular, they are disjoint. Moreover, by property (ii) of escaping subdivisions, we have that for $M := \max\{M_1^C + 2, M_1^D + 2\}$ the paths P_{ij}^C and P_{ij}^D with $i \geqslant M$ are contained in C and D, respectively, and are hence disjoint from each other. Let $H'_C \subseteq H_C$ be a subdivision of the hexagonal half-grid with vertical double rays S_0^C and S_i^C for $i \geqslant M$, which we may obtain by choosing as the new branch paths for the horizontal edges e_{1j} infinitely many disjoint S_0^C–S_M^C paths Q_j^C in H'_C. Let H'_D be chosen analogously. Clearly, H'_C and H'_D are still escaping.

Now since $S_0^C = R = S_0^D$, gluing H'_C and H'_D together along R yields a graph H' which is nearly as desired except that the paths Q_j^C and Q_ℓ^D may intersect. But since G is locally finite, we can delete some of the Q_j^C and Q_ℓ^D and apply Proposition 13.3.6 to obtain a subdivision $H \subseteq H'$ of the hexagonal full-grid. By construction, H is escaping. $\square$

13.6.1 Half-grids with crosses

In this section we establish a sufficient condition which ensures that a graph G contains $K_{\aleph_0}$ as an ultra fat minor. This condition essentially requires an escaping subdivision $H \subseteq G$ of the hexagonal half-grid, and infinitely many H-paths in G that 'jump over' the vertical double rays in H. For this, we first need the following two auxiliary statements about $K_{\aleph_0}$ minors in half-grids with certain additional edges.

Lemma 13.6.3. *Let G be obtained from the half-grid by adding all edges of the form $(i,0)(i+1,1)$ and $(i,1)(i+1,0)$ for $i \in \mathbb{N}$. Then G contains a model $(\mathcal{V}, \mathcal{E})$ of $K_{\aleph_0}$ such that $V_i, V(E_{ij}) \subseteq \mathbb{N}_{\geqslant i} \times \mathbb{Z}$ for all $i < j \in \mathbb{N}$ and such that the E_{ij} are pairwise disjoint.*

Proof. One can easily construct the $K_{\aleph_0}$ minor recursively starting from K_1 with branch set $V_1 := \{(0,0)\}$. For this, assume that we have already defined a model of K_n with branch sets $V_1, \dots, V_n$, branch paths E_{ij} and integers $j_n, \ell_n \in \mathbb{N}$ such that $V_i, E_{ij} \subseteq \{i, \dots, \ell_n\} \times \{-j_n, \dots, 0, \dots, j_n\}$ and such that either all V_i meet $\{i, \ell_n\} \times \{-j_n\}$ or all V_i meet $\{0, \ell_n\} \times \{j_n\}$ and such that the E_{ij} are pairwise disjoint. We can then extend the branch sets $V_1, \dots, V_n$ and add a new branch set V_{n+1} as well as new branch paths $E_{(n+1)j}$ as depicted in Figure 13.4 such that $V_1, \dots, V_{n+1}$ are the branch sets and the E_{ij} are the branch paths of a K_{n+1} minor. $\square$

In particular, we have the following corollary:

Corollary 13.6.4. *Let X be obtained from the half-grid by adding disjoint edges f_n between (i_n, k_n) and (j_n, ℓ_n), for every $n \in \mathbb{N}$, such that $i_n, j_n < i_{n'}, j_{n'}$ and $|i_n - j_n| \geqslant 2$ for all $n < n'$. Then X*

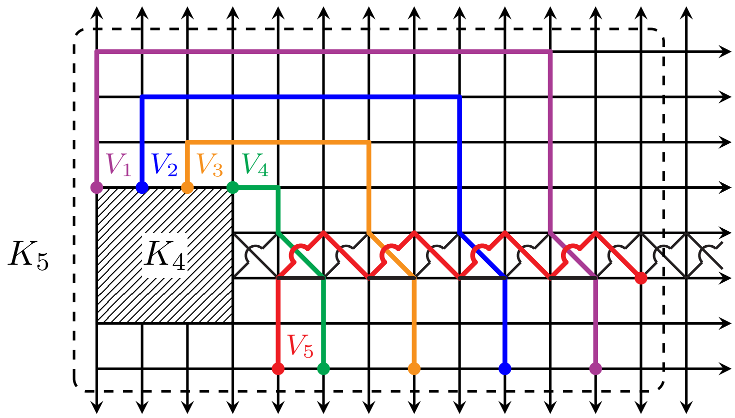

FIGURE 13.4: Sketch of a K_5 minor in a half-grid with all $(i,0)(i+1,1)$ and $(i,1)(i+1,0)$ edges. The branch paths E_{ij} between V_5 and the V_i are the thickened $(i,0)(i+1,0)$ edges.

contains a model $(\mathcal{V},\mathcal{E})$ of $K_{\aleph_0}$ such that $V_i, V(E_{ij}) \subseteq \mathbb{N}_{\geqslant i} \times \mathbb{Z}$ for all $i < j \in \mathbb{N}$ and such that the E_{ij} are pairwise disjoint.

Proof. It is straight forward to check that the lines $\{i_n\} \times \mathbb{Z}$ in X form the vertical double rays of a subdivision X' of the graph G in the premise of Lemma 13.6.3 such that every branch path in X' corresponding to an edge of G between $\{n\}\times\mathbb{Z}$ and $\{n+1\}\times\mathbb{Z}$ is contained in $X[\{i_n,\dots,i_{n+1}\}\times\mathbb{Z}]$. Hence, Lemma 13.6.3 immediately yields the assertion. □

Before we can state the main lemma of this subsection, we first need the following definition. Let H be an escaping subdivision of the hexagonal half-grid in a graph G with vertical double rays S^i and horizontal paths P_{ij}. An H-path Q in G with endvertices on S^i and S^j for some $i,j \in \mathbb{N}$ is *K-fat* for some $K \in \mathbb{N}$ if

- $d_G(Q, S^k) \geqslant \min\{k, K\}$ for all $k \neq i,j \in \mathbb{N}$, and
- $d_G(Q, P_{k\ell}) \geqslant \min\{k, K\}$ for all $k \in \mathbb{N}$ and $\ell \in \mathbb{Z}$.

Lemma 13.6.5. *Let H be an escaping subdivision of the hexagonal half-grid in a locally finite graph G with vertical double rays S^i. Suppose there are infinitely many pairwise disjoint H-paths Q_m with endvertices on S^{i_m} and S^{j_m} for some $i_m, j_m \in \mathbb{N}$ such that Q_m is m-fat and such that $i_m, j_m < i_{m'}, j_{m'}$ and $|i_m - j_m| \geqslant 2$ for all $m < m'$. Then H contains $K_{\aleph_0}$ as an ultra fat minor.*

Proof. Since G is locally finite and the Q_m are finite, we may assume that $d_G(Q_m, Q_{m'}) \geqslant m$ for all $m < m' \in \mathbb{N}$, by possibly deleting some of the Q_m. Further, we may assume, for every $Q_m =: q_0^m \dots q_{n_m}^m$, that $Q_m \cap G[S^{i_m}, i'_m] = q_0^m \dots q_{i'_m}^m$ and $Q_m \cap G[S^{j_m}, j'_m] = q_{n_m - j'_m}^m \dots q_{n_m}^m$ where $i'_m := \min\{\lfloor i_m/3 \rfloor, m\}$ and $j'_m := \min\{\lfloor j_m/3 \rfloor, m\}$. In particular, if Q_m has distance less than

$\lfloor i'_m/3 \rfloor$ to vertices $u, v \in V(S^{i_m})$, then $d_G(u,v) < i'_m$, and similarly for j_m. Indeed, let Q'_m be a subpath of Q_m which is a $B_G(S^{i_m}, i'_m)$–$B_G(S^{j_m}, j'_m)$ path. Then we can replace Q_m by a path that consists of Q'_m and a shortest S^{i_m}–Q'_m path and a shortest Q'_m–S^{j_m} path. Note that to regain that $d_G(Q_m, P_{k\ell}) \geqslant \min\{k, m\}$, we might have to delete some of the $P_{k\ell}$. But since the new Q_m can only be too close to paths $P_{k\ell}$ with $k \in \{i_m, i_m+1, j_m, j_m+1\}$ and because G is locally finite, we only need to delete at most finitely many $P_{k\ell}$ for every k, so by Proposition 13.3.6 this yields a subdivision $H'' \subseteq H'$ of the hexagonal half-grid with the same vertical double rays as H'.

Then the graph $\tilde{H}$ obtained from the union of H'' and the Q_m is a subdivision of a graph X as in the premise of Corollary 13.6.4. Hence, G contains a model $(\mathcal{V}, \mathcal{E})$ of $K_{\aleph_0}$ such that $V_i, V(E_{ij}) \subseteq \tilde{H}_i := H''_{\geqslant i} \cup \bigcup_{m \geqslant i} Q_m$ and such that the E_{ij} are pairwise disjoint. By Lemma 13.3.9 and the assumptions on the Q_m, we may assume for every $K \in \mathbb{N}$ that the images of every two non-incident edges of X have distance at least $\lfloor K/3 \rfloor$ in G if their images in $\tilde{H}$ are contained in $\tilde{H}_K$. Since every vertex of $K_{\aleph_0}$ has degree at least 2, we may assume that if some $V_i \in \mathcal{V}$ meets some branch paths of $\tilde{H}$ in an inner vertex, then it in fact contains it. Similarly, every $E_{ij} \in \mathcal{E}$ contains any branch path of $\tilde{H}$ as soon as it meets an inner vertex of it. Hence, since the V_i, E_{ij} are pairwise disjoint (except for incident branch set - path pairs) and contained in $\tilde{H}_i$, they have distance at least $\lfloor i/3 \rfloor$ to all $V_k, E_{k\ell}$ with $k \geqslant i$. It follows that $((V_i)_{i \in 3\mathbb{N}}, (E_{ij})_{i \neq j \in 3\mathbb{N}})$ is ultra fat. □

13.6.2 Finding two thick components

Let G be a locally finite, quasi-transitive graph whose cycle space is generated by short cycles. In this subsection we show that if G has a thick end, then either $K_{\aleph_0} \prec_{UF} G$, or G contains a quasi-geodesic double ray such that, for some $L \in \mathbb{N}$, there are at least two thick components of $G - B_G(R, L)$. Together with Lemma 13.6.2, this then concludes the proof of Theorem 13.6.1. The proof of this assertion is mainly divided into two lemmas, Lemmas 13.6.8 and 13.6.9 below; see Section 13.3.4 for a sketch of the proof.

To carry out the proofs of Lemmas 13.6.8 and 13.6.9, we need the following two auxiliary statements.

Lemma 13.6.6. *Let $R = \dots r_{-1} r_0 r_1 \dots$ be a quasi-geodesic double ray in a graph G of finite maximum degree whose cycle space is generated by cycles of bounded length. For every $K, N \in \mathbb{N}$ there exist $d = d(K, N), \ell = \ell(K) \in \mathbb{N}$ such that the following holds: If a component C of $G - B_G(R, K)$ has neighbours in $B_G(R_{\leqslant i}, K)$ and $B_G(R_{\geqslant i+N}, K)$ for some $i \in \mathbb{Z}$, then there exist vertices $x \in \partial_G C \cap B_G(r_{i-\ell} R r_i, K+1)$ and $y \in \partial_G C \cap B_G(r_{i+N} R r_{i+N+\ell}, K+1)$ such that $d_C(x, y) \leqslant d$.*

Proof. Let $c, \kappa \in \mathbb{N}$ such that R is c-quasi-geodesic and such that the cycle space of G is generated by cycles of length at most κ. By assumption, the maximum degree $\Delta(G)$ of G is finite. We prove the assertion with $d := N \cdot \Delta(G)^{K+\lfloor\frac{\kappa}{2}\rfloor+1} + \kappa$ and $\ell := c(2K + \kappa + 1)$.

Let $U^+, U^*, U^- \subseteq B_C(\partial_G C, \lfloor\frac{\kappa-2}{2}\rfloor)$ be the set of vertices in C that have distance at most $K + \lfloor\frac{\kappa}{2}\rfloor$ in G from $R_{\geqslant i+N}$, $r_i R r_{i+N}$ and $R_{\leqslant i}$, respectively; and note that by the assumptions on C, the sets U^+ and U^- are non-empty. By Lemma 13.5.1, $C[\partial_G C, \lfloor\frac{\kappa-2}{2}\rfloor]$ is connected. Hence, since $U^+, U^- \subseteq B_C(\partial_G C, \lfloor\frac{\kappa-2}{2}\rfloor)$, there is a U^+–U^- path $P = p_0 \dots p_n$ in $C[\partial_G C, \lfloor\frac{\kappa-2}{2}\rfloor]$. Clearly, since $B_C(\partial_G C, \lfloor\frac{\kappa-2}{2}\rfloor) = U^+ \cup U^* \cup U^-$, the choice of P guarantees that $\mathring{P} = p_1 \dots p_{n-1}$ is contained in U^*, and hence P has length at most $|U^*| + 1 \leqslant N \cdot \Delta(G)^{K+\lfloor\frac{\kappa}{2}\rfloor+1}$. Let $j, j' \in \mathbb{Z}$ such that P starts in $V(C) \cap B_G(r_j, K + \lfloor\frac{\kappa}{2}\rfloor)$ and ends in $V(C) \cap B_G(r_{j'}, K + \lfloor\frac{\kappa}{2}\rfloor)$, and let Q be a shortest $B_G(r_j, K+1)$–P path and Q' a shortest $B_G(r_{j'}, K+1)$–P path. Then the concatenation of Q, P and Q' yields a $B_G(r_j, K+1)$–$B_G(r_{j'}, K+1)$ path P' in G which starts in a vertex $x \in \partial_G C \cap B_G(r_j, K+1)$ and ends in some $y \in \partial_G C \cap B_G(r_{j'}, K+1)$. In particular, Q and Q' have length at most $\lfloor\frac{\kappa-2}{2}\rfloor$, and hence P' has length at most d, which implies $d_C(x, y) \leqslant d$. Moreover, $j \geqslant i + N$ and $j' \leqslant i$. It remains to show that $j \leqslant i + N + \ell$ and $j' \geqslant i - \ell$. Since $\mathring{P}$ lies in U^*, there exists a p_1–$r_i R r_{i+N}$ path Q'' of length at most $K + \lfloor\frac{\kappa}{2}\rfloor$. Then $Q p_0 p_1 Q''$ is an r_j–r_k path for some $k \leqslant i + N$. As $Q p_0 p_1 Q''$ has length at most $(K + \lfloor\frac{\kappa}{2}\rfloor) + 1 + (K + \lfloor\frac{\kappa}{2}\rfloor) \leqslant 2K + \kappa + 1$ and R is c-quasi-geodesic, it follows that $|j - k| = d_R(r_j, r_k) \leqslant c(2K + \kappa + 1) = \ell$, and hence $j \leqslant k + \ell \leqslant i + N + \ell$. The case $j' \geqslant i - \ell$ is analogous. □

Corollary 13.6.7. *Let $L \in \mathbb{N}$, and let R be a quasi-geodesic double ray in some locally finite graph G whose cycle space is generated by cycles of bounded length. Then $G - B_G(R, L)$ has at most finitely many long components.*

In particular, if, for every $K \geqslant L$, $G - B_G(R, K)$ has a long component C_K, then $G - B_G(R, L)$ has a thick component which contains infinitely many C_K.

Proof. Applying Lemma 13.6.6 to R and $K := L$ and $N := 0$ yields some $\ell \in \mathbb{N}$ such that every long component of $G - B_G(R, L)$ has a neighbour in $B_G(r_0 R r_\ell, L)$. Hence, since $B_G(r_0 R r_\ell, L)$ is finite as $r_0 R r_\ell$ is finite and G is locally finite, there are at most finitely many long components.

For the 'in particular'-part note that since G is locally finite, every component of $G - B_G(R, L)$ which contains a long component of $G - B_G(R, K)$ for some $K \geqslant L$ is long. Hence, the assertion follows from the first part by the pigeonhole principle. □

We can now prove the two main lemmas of this section. Given a double ray $R = \dots r_{-1} r_0 r_1 \dots$ in a graph G, a component C of $G - B_G(R, L)$ is *half-long* if C has neighbours in $B_G(R_{\geqslant i}, L)$ *or* in $B_G(R_{\leqslant -i}, K)$ for all $i \in \mathbb{N}$, and C is *half-thick* if, for every $M \geqslant L$, some half-long component of $G - B_G(R, M)$ is contained in C.

Lemma 13.6.8. *Let $c, L \in \mathbb{N}$ and let R be a c-quasi-geodesic double ray in a locally finite, quasi-transitive graph G whose cycle space is generated by cycles of bounded length. Assume that $G - B_G(R, L)$ has distinct components $C \neq D$ such that C is thick and D is half-thick. Then G contains a c-quasi-geodesic double ray S such that $G - B_G(S, L)$ has two thick components.*

Proof. Let $\kappa \in \mathbb{N}$ such that the cycle space of G is generated by cycles of length at most κ. As we are done if D is thick, we may assume that we can enumerate $R =: \dots r_{-1}r_0r_1 \dots$ so that $N_G(D) \subseteq B_G(R_{\geqslant 0}, L)$. Since G is quasi-transitive, there is an infinite index set $I_0 \subseteq \mathbb{N}$ such that all r_i with $i \in I_0$ lie in the same $\mathrm{Aut}(G)$-orbit. Let v be another vertex in that orbit. Then there exists a sequence $(\varphi_i)_{i \in I_0}$ of automorphisms of G such that $\varphi_i(r_i) = v$. Since G is locally finite, there is an infinite index set $I_1 \subseteq I_0$ such that $\varphi_i(r_{i-1}r_ir_{i+1})$ coincides for all $i \in I_1$ amongst which we again find some infinite set $I_2 \subseteq I_1$ such that $\varphi_i(r_{i-2} \dots r_{i+2})$ coincides for all $i \in I_2$ and so on. Now pick for every $n \in \mathbb{N}$ some $i_n \in I_n$, and let I consist of these i_n. This leads to a c-quasi-geodesic double ray S that contains v and such that every subpath of S of length $2\ell \in \mathbb{N}$ that contains v as central vertex is the image of $r_{i_n - \ell}Rr_{i_n + \ell}$ under φ_{i_n} for all $i_n \in I$ with $n \geqslant \ell$. We enumerate $S =: \dots s_{-1}s_0s_1 \dots$ where $s_0 := v$ and $s_1 = \varphi_{i_1}(r_{i_1+1})$.

We claim that S is as desired. For this, we show that, for every $K \geqslant L$, $G - B_G(S, K)$ has long components C'_K, D'_K such that every C'_K–D'_K path in G meets $B_G(S, L)$. Then the assertion follows. Indeed, by Corollary 13.6.7, $G - B_G(S, L)$ has a thick component E which contains infinitely many of the C'_K, D'_K. If infinitely many of the C'_K, D'_K do not lie in E, then applying Corollary 13.6.7 again yields a second thick component $E' \neq E$. Otherwise, at most finitely many of the C'_K, D'_K are not contained in E, which implies that E contains both C'_K and D'_K for some $K \geqslant L$. But since E is connected and avoids $B_G(S, L)$, this contradicts that every C'_K–D'_K path meets $B_G(S, L)$.

So let $K \geqslant L$ be given. It remains to show that $G - B_G(S, K)$ has long components C'_K, D'_K such that every C'_K–D'_K path in G meets $B_G(S, L)$. Since C is thick and D is half-thick there exist components $C_K \subseteq C$ and $D_K \subseteq D$ of $G - B_G(R, K)$ such that C_K is long and D_K is half-long.

Claim 1. *There are long components C'_K, D'_K of $G - B_G(S, K)$ and vertices $x \in \partial_G C'_K$ and $y \in \partial_G D'_K$ such that $\varphi_i^{-1}(x) \in \partial_G C_K$ and $\varphi_i^{-1}(y) \in \partial_G D_K$ for infinitely many $i \in I$.*

Proof. Let us first find a component D'_K. Since $N_G(D) \subseteq B_G(R_{\geqslant 0}, L)$ and $K \geqslant L$, we also have $N_G(D_K) \subseteq B_G(R_{\geqslant 0}, K)$. Let $m \in \mathbb{N}$ such that $N_G(D_K) \cap B_G(r_m, K) \neq \emptyset$. For every $n \in \mathbb{N}$, let ℓ, d_n be given by applying Lemma 13.6.6 to R, K and $N := 2n$ (note that ℓ does not depend on N).

Then there is, for every $n \in \mathbb{N}$ and $i \geqslant m+n$, a path P_{in} in D_K from $B_G(r_{i-n-\ell}Rr_{i-n}, K+1)$ to $B_G(r_{i+n}Rr_{i+n+\ell}, K+1)$ of length at most d_n with endvertices $p_{in} \in \partial_G D_K \cap B_G(r_{i-n-\ell}Rr_{i-n}, K+1)$ and $q_{in} \in \partial_G D_K \cap B_G(r_{i+n}Rr_{i+n+\ell}, K+1)$.

By the choice of I, we have $\varphi_i(r_{i-\ell}Rr_{i+\ell}) = s_{-\ell}Ss_\ell$ for all large enough $i \in I$, and hence these automorphisms φ_i map $B_G(r_{i-\ell}Rr_{i+\ell}, K+1)$ to $B_G(s_{-\ell}Ss_\ell, K+1)$. As G is locally finite, the set $B_G(s_{-\ell}Ss_\ell, K+1)$ is finite. Combining these facts with $p_{i0}, q_{i0} \in B_G(r_{i-\ell}Rr_{i+\ell}, K+1)$ yields that there are $p_0, q_0 \in V(G)$ and an infinite index set $J_0 \subseteq I$ such that $\varphi_i(p_{i0}) = p_0$ and $\varphi_i(q_{i0}) = q_0$ for all $i \in J_0$. Using again that G is locally finite and P_{i0} has length at most d_0 for all $i \in J_0$, we find a p_0–q_0 path $P_0 \subseteq G$ and an infinite index set $J_0' \subseteq J_0$ such that $\varphi_i(P_{i0}) = P_0$ for all $i \in J_0'$. By the same argument, we find a subsets $J_0' \supseteq J_1' \supseteq \dots$ such that, for all $n \in \mathbb{N}$ and $i \in J_n'$, $\varphi_i(P_{in}) = P_n$ for some path $P_n \subseteq G$ with endvertices $p_n, q_n \in V(G)$. Pick for every $n \in \mathbb{N}$ some $i \in J_n'$ and let J consist of these i.

We claim that there is a component D_K' of $G - B_G(S, K)$ which contains infinitely many of the P_n, and which is hence long. For this, we first show that no P_n meets $B_G(S, K)$. So let $n \in \mathbb{N}$ be given, and set $d_n' := n + \ell + c(d_n + 2K + 2)$. By the choice of P_n, we have $P_n = \varphi_i(P_{in})$ for all large enough $i \in J$. Since also $\varphi_i(B_G(r_{i-d_n'}Rr_{i+d_n'}, K)) = B_G(s_{-d_n'}Ss_{d_n'}, K)$ for all large enough $i \in J$ and since P_{in} avoids $B_G(R, K)$, it follows that P_n avoids $B_G(s_{-d_n'}Ss_{d_n'})$. Moreover, P_n is also disjoint from $B_G(S_{<-d_n'}, K)$ and $B_G(S_{>d_n'}, K)$ since $d_n' = n + \ell + c(d_n + 2K + 2)$ and because S is c-quasi-geodesic and P_n has length at most d_n and starts in $B_G(s_{-n-\ell}Ss_{-n}, K+1)$ and ends in $B_G(s_nSs_{n+\ell}, K+1)$.

So every P_n avoids $B_G(S, K)$ and is hence contained in a component of $G - B_G(S, K)$. By the choice of ℓ via Lemma 13.6.6 and because P_n starts in $B_G(S_{\geqslant n}, K+1)$ and ends in $B_G(S_{\leqslant -n}, K+1)$, every component that contains some P_n with $n \geqslant \ell$ attaches to $B_G(s_0Ss_\ell, K)$. Since this set is finite as G is locally finite, infinitely many P_n lie in the same component D_K' of $G - B_G(S, K)$, which then needs to be long. Now since $\varphi_i^{-1}(P_n) = P_{in} \subseteq D_K$ for all $i \in I'$, where I' is an infinite subset of I, we may choose $y \in \partial_G D_K'$ as one endvertex of some P_n which is contained in D_K'.

The proof for C_K' is now completely analogous except that we have to use I' instead of I in order to ensure that $\varphi_i^{-1}(x) \in \partial_G C_K$ and $\varphi_i^{-1}(y) \in \partial_G D_K$ for (the same) infinitely many $i \in I$. ■

Let C_K', D_K' and $x \in \partial_G C_K'$, $y \in \partial_G D_K'$ be given by Claim 1. To finish the proof, we are left to show that every C_K'–D_K' path in G meets $B_G(S, L)$. For this, suppose for a contradiction that there is a C_K'–D_K' path that avoids $B_G(S, L)$. Then, since $C_K' \ni x$ and $D_K' \ni y$ are connected, there also exists an x–y path Q that avoids $B_G(S, L)$. Denote by m the length of Q, let $\ell \in \mathbb{N}$ such that $x, y \in B_G(s_{-\ell}Ss_\ell, K+1)$, and set $m' := c(K + m + L + 2) + \ell$. By the choice of S there is some i among the infinitely many $i \in I$ that satisfy $\varphi_i^{-1}(x) \in \partial_G C_K$ and $\varphi_i^{-1}(y) \in \partial_G D_K$ such that $\varphi_i(r_{i-m'}Rr_{i+m'}) = s_{-m'}Ss_{m'}$. Since Q avoids $B_G(S, L)$, it follows that $\varphi_i^{-1}(Q)$ avoids $B_G(r_{i-m'}Rr_{i+m'}, L)$. But $\varphi_i^{-1}(Q)$ also avoids $B_G(R_{<i-m'}, L) \cup B_G(R_{>i+m'}, L)$: Otherwise, since $\varphi_i^{-1}(Q)$ has length m and starts in $B_G(r_{i-\ell}Rr_{i+\ell}, K+1)$, there would be a path of length at most $L + m + (K+1)$ that joins a vertex from $R_{<i-m'} \cup R_{>i+m'}$ to a vertex from $r_{i-\ell}Rr_{i+\ell}$. Since $m' - \ell = c(L + m + K + 2)$, this would contradict that R is c-

quasi-geodesic. Hence, $\varphi_i^{-1}(Q)$ is an $\varphi_i^{-1}(x)$–$\varphi_i^{-1}(y)$ path which avoids $B_G(R, L)$. But since $\varphi_i^{-1}(x) \in \partial_G C_K \subseteq C_L$ and $\varphi_i^{-1}(y) \in \partial_G D_K \subseteq D_L$, this contradicts that C_L and D_L are distinct components of $G - B_G(R, L)$. □

Lemma 13.6.9. *Let G be a locally finite, quasi-transitive graph with a thick end whose cycle space is generated by cycles of bounded length. If G does not contain $K_{\aleph_0}$ as an ultra fat minor, then there exists $L \in \mathbb{N}$ and a quasi-geodesic double ray R in G such that $G - B_G(R, L)$ has distinct components $C \neq D$ such that C is thick and D is half-thick.*

Proof. Let R_1, R_2, R_3 be given by applying Theorem 13.4.5 to some thick end of the locally finite, quasi-transitive graph G, which is accessible by Theorem 13.2.1. Then applying Theorem 13.3.3 to the quasi-geodesic, and hence diverging, double ray $R_1 \cup R_2$ yields an escaping subdivision H of the hexagonal half-grid with vertical double rays S^i and horizontal paths P_{ij} such that $S^0 = R_1 \cup R_2$. Since H is escaping, there exist $M_0 < M_1 < \ldots \in \mathbb{N}$ such that $M_i > M_{i-1} + 2i$ for all $i \geqslant 1$ and

(i) $S^i \subseteq G[S^0, M_i] - B_G(S^0, M_{i-1} + 2i)$ for all $i \in \mathbb{N}_{\geqslant 1}$, and

(ii) $P_{1j} \subseteq G[S^0, M_1]$ and $P_{ij} \subseteq G[S^0, M_i] - B_G(S^0, M_{i-2} + i)$ for all $i \in \mathbb{N}_{\geqslant 2}$ and $j \in \mathbb{Z}$.

Let us also note that, since $R_1 \cup R_2 \cup R_3$ is quasi-geodesic, we also have that

(iii) the set $V(R_3) \cap B_G(S^0, L)$ is finite for all $L \in \mathbb{N}$.

Hence, by (ii) and (iii) and because G is locally finite, we may assume, by deleting at most finitely many P_{ij} for every $i \in \mathbb{N}$ and applying Proposition 13.3.6, that

(iv) $d_G(R_3, P_{ij}) > 2i$ for all $i \in \mathbb{N}$ and $j \in \mathbb{Z}$.

Let $H_{\geqslant n} \subseteq H$, for $n \in \mathbb{N}$, be the subgraph consisting of all S^i, P_{kj} with $i \geqslant n$ and $k > n$. By (i) and (ii), we have, for every $L \in \mathbb{N}_{\geqslant 1}$, that $d_G(H_{\geqslant L}, S^0) > L$; let C_L be the component of $G - B_G(S^0, L)$ containing $H_{\geqslant L}$. Clearly, the C_L are long and hence, since $C_1 \supseteq C_2 \supseteq \ldots$, the C_L are thick.

If there is some $L \in \mathbb{N}_{\geqslant 1}$ such that $R_3 \cap C_L = \emptyset$, then we are done. Indeed, by (iii), there is, for every $L' \geqslant L$, a component $D_{L'}$ of $G - B_G(S^0, L')$ that contains a tail of R_3. Since R_3 and R_1 belong to the same end, the components $D_{L'}$ are half-long. As clearly $D_{L'} \subseteq D_L$ for all $L' \geqslant L$, we find that D_L is half-thick. Since also $D_L \neq C_L$ by assumption, $R := R_1 \cup R_2$, L, $C := C_L$ and $D := D_L$ are as desired.

Thus, we may assume that $R_3 \cap C_L \neq \emptyset$ for all $L \in \mathbb{N}_{\geqslant 1}$. We distinguish two cases.

Case 1: *For all $K \in \mathbb{N}$ there is some $N_K \in \mathbb{N}$ such that $d_G(R_3, S^i) > K$ for all $i \geqslant N_K$.* (See Figure 13.5.)

We will use R_3 to find fat H-paths as in Lemma 13.6.5, which then implies that G contains $K_{\aleph_0}$ as an ultra fat minor, concluding the first case of the proof.

Without loss of generality let $N_K \geqslant K$ for all $K \in \mathbb{N}$. By (i) and (iii), for every $K \in \mathbb{N}$, the ray R_3 has a tail T_K which avoids $B_G(S^0, M_{N_K-1}+K)$, and which thus satisfies $d_G(T_K, S^i) > K$ for all $i \in \mathbb{N}$ by (i). By (ii) and (iv), we also find $d_G(P_{ij}, T_K) > K$ for all $i \in \mathbb{N}$ and $j \in \mathbb{Z}$, and hence $d_G(H, T_K) > K$.

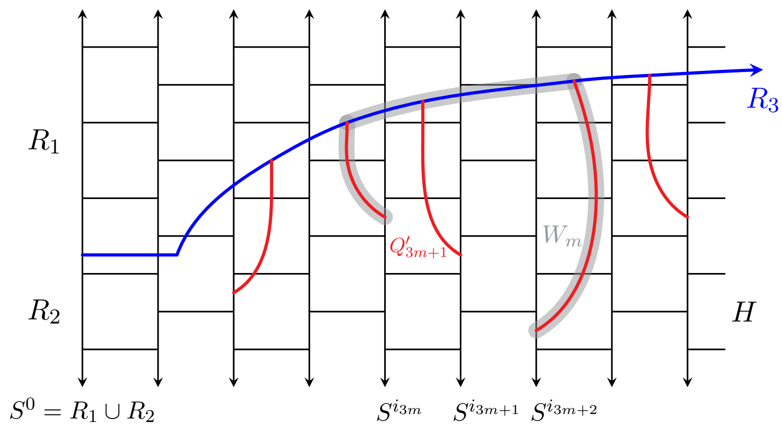

Figure 13.5: Sketch of Case 1 in the proof of Lemma 13.6.9: the hexagonal half-grid H and the ray R_3, which has large distance from H but is connected to H by paths Q'_m. The paths W_m consisting of Q'_{3m}, Q'_{3m+2} and a suitable subpath of R_3 are m-fat H-paths.

By (i), and because $R_3 \cap C_L \neq \emptyset$ as well as $H_{\geqslant L} \subseteq C_L$ for all $L \in \mathbb{N}_{\geqslant 1}$, there exists, for every $m \in \mathbb{N}$, an T_m–$\bigcup_{i \geqslant N_m} B_G(S^i, m)$ path Q_m that ends in some $B_G(S^{i_m}, m)$ for $i_m \geqslant m$ and avoids $B_G(S^0, M_{N_m}+m)$. We extend each Q_m to an T_m–S^{i_m} path Q'_m by adding a shortest S^{i_m}–Q_m path. Since $d_G(Q_m, S^i) > m$ for all $i \neq i_m$ by the choice of Q_m and again by (i), we find $d_G(Q'_m, S^i) > m$ for all $i \neq i_m$. Moreover, by (i), the paths Q'_m still avoid $B_G(S^0, M_{N_m}+m)$. Hence, by (ii), we have that, for every $i \in \mathbb{N}$, the paths Q'_m with $m \geqslant i$ have distance at least m from the paths P_{ij}. Since G is locally finite, this implies that we may assume, by deleting for every $i \in \mathbb{N}$ at most finitely many P_{ij} and applying Proposition 13.3.6, that $d_G(P_{ij}, Q'_m) > m$ for all $m, i \in \mathbb{N}$ and $j \in \mathbb{Z}$. All in all, we find $d_G(H - S^{i_m}, Q'_m) > m$ for all $m \in \mathbb{N}$.

Since Q'_m avoids $B_G(S^0, M_{N_m}+m)$ but is itself eventually contained in some $B_G(S^0, m')$ for $m' > m$ as Q'_m is finite, we may assume, by passing to a subsequence of the Q'_m if necessary, that the Q'_m are pairwise disjoint. Moreover, by (i) and again since Q'_m avoids $B_G(S^0, M_{N_m})$, we may assume that $i_1 < i_2 < \dots$, by once again passing to a subsequence of the Q'_m if necessary.

Now since $d_G(H - S^{i_m}, Q'_m) > m$ and $d_G(H, T_m) > m$, the paths W_m that consist of Q'_{3m}, Q'_{3m+2} and a suitable subpath of T_{3m} are m-fat H-paths with endvertices on $S^{i_{3m}}$ and $S^{i_{3m+2}}$ (see

Figure 13.5). In particular, since the Q'_m are pairwise disjoint, we may assume, by passing to a subsequence of the W_m, that also the W_m are pairwise disjoint. Hence, the W_m are m-fat H-paths as in Lemma 13.6.5, which implies that G contains $K_{\aleph_0}$ as an ultra fat minor. This concludes the first case of the proof.

Case 2: *There exists some K such that $d_G(R_3, S^i) < K$ for infinitely many $i \in \mathbb{N}$.* (See Figure 13.6.)

We first show that we may assume, by passing to a subgraph of H if necessary, that

(a) $d_G(R_3, S^i) < i$ for all $i \geqslant K$, and

(b) for all $K \leqslant j < i$, if $d_G(r, S^i) < i$ for some $r \in V(R_3)$, then $d_G(rR_3, S^j) \geqslant j$.

Indeed, set $i_j := j$ for all $j < K$, and let $i_K \geqslant K$ be minimal such that $d_G(R_3, S^{i_K}) < K$. By (i) and (iii), R_3 has a tail T that is disjoint from $B_G(S^0, M_{i_K} + i_K)$. Since $R_3 - T$ is finite and because of (i), we have $d_G(R_3 - T, S^i) > i$ for all large enough $i \in \mathbb{N}$, and hence, as $d_G(R_3, S^i) < K$ for infinitely many $i \in \mathbb{N}$, there is some $i_{K+1} > i_K$ such that $d_G(T, S^{i_{K+1}}) < i_{K+1}$ and $d_G(R_3 - T, S^i) \geqslant i$ for all $i \geqslant i_{K+1}$. By continuing in this way, we obtain a sequence $0 := i_0 < i_1 < \ldots \in \mathbb{N}$ such that the rays S^{i_j} satisfy (a) and (b). Now pick a subdivision $H' \subseteq H$ of the hexagonal half-grid whose vertical double rays are precisely the S^{i_j}. It is easy to see that H' still satisfies (i) to (iii). By deleting from H' at most finitely many P'_{ij} for every $i \in \mathbb{N}$ and applying Proposition 13.3.6, we also regain property (iv) for H'. Hence, H' is the desired subgraph of H, and we denote H' again by H.

Let $S^i =: \ldots s^i_{-1} s^i_0 s^i_1 \ldots$ such that $d_G(R_3, s^i_0) < i$ and such that $S^i_{\geqslant 0}$ is the 'upper half' and $S^i_{\leqslant 0}$ is the 'lower half' of S^i (see Figure 13.6). Further, let H' be the 'upper half' of H with respect to this enumeration, that is, let H' be the subgraph of H that consists of the $S^i_{\geqslant 0}$ and all paths P_{ij} whose endvertices lie on the $S^i_{\geqslant 0}$. Let H'' be the 'lower half' of H defined analogously.

Now choose for all $i, L \in \mathbb{N}$ maximal tails T'_{iL}, T''_{iL} of $S^i_{\geqslant 0}$ and $S^i_{\leqslant 0}$, respectively, such that

(c) $d_G(R_3, T'_{iL}) \geqslant 2L$ and $d_G(R_3, T''_{iL}) \geqslant 2L$.

Note that by (i) and (iii), the T'_{iL}, T''_{iL} are non-empty. Note further that, by (ii) and (iii), for all $i, L \in \mathbb{N}$ all but finitely many of the P_{ij} avoid $B_G(R_3, L)$. Hence, there are, for every $L \in \mathbb{N}$, escaping subdivisions $H'_L, H''_L \subseteq H_{\geqslant L}$ of the hexagonal 'quarter-grid' whose vertical rays are the T'_{iL} or T''_{iL}, respectively, for $i \geqslant L$ such that $H'_1 \supseteq H'_2 \supseteq \ldots$ as well as $H''_1 \supseteq H''_2 \supseteq \ldots$, and such that H'_L and H''_L avoid $B_G(R_1 \cup R_3, L)$.

Since H'_L and H''_L avoid $B_G(R_1 \cup R_3, L)$, they are contained in components C_L and D_L of $G - B_G(R_1 \cup R_3, L)$, respectively. Clearly, C_L is long and D_L is half-long. As $C_1 \supseteq C_2 \supseteq \ldots$ and $D_1 \supseteq D_2 \supseteq \ldots$, this implies that C_L is thick and D_L is half-thick. If there is some $L \in \mathbb{N}$ such that $C_L \neq D_L$, then we are done as then L, $R := R_1 \cup R_3$, $C := C_L$ and $D := D_L$ are as desired.

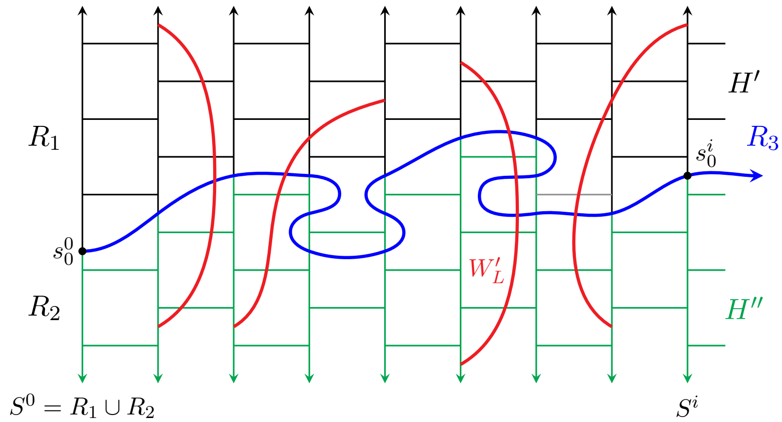

FIGURE 13.6: Sketch of Case 2 in the proof of Lemma 13.6.9: the hexagonal half-grid H and the ray R_3, which has distance $< i$ to every vertical double ray S^i of H by (a). By (b), once R_3 comes close to some S^i, it will never come close to some S^j with $j < i$ again. The paths W'_L have distance at least L from R_3 and 'jump over' R_3.

Hence, we may assume that $C_L = D_L$ for all $L \in \mathbb{N}$. We will now once again construct fat H-paths as in the premise of Lemma 13.6.5, which then yields that G contains $K_{\aleph_0}$ as an ultra fat minor, and which thus concludes the proof. Since $C_L = D_L$ and $H'_L \subseteq C_L$, $H''_L \subseteq D_L$ for all $L \geqslant K$, there are $(\bigcup_{i\in\mathbb{N}} B_G(T'_{iL}, \min\{i,L\}))$–$(\bigcup_{i\in\mathbb{N}} B_G(T''_{iL}, \min\{i,L\}))$ paths W_L that avoid $B_G(R_1 \cup R_3, 3L)$. We now modify W_L as follows. Let i'_L be such that W_L starts in $B_G(T'_{i'_L L}, \min\{i'_L, L\})$. If W_L meets $B_G(P_{ij}, \min\{i,L\})$ for some $P_{ij} \subseteq H'_{4L}$ with $i \notin \{i'_L, i'_L+1\}$, then we let P_{ij} be the last such path, and we shorten W_L so that it meets $B_G(P_{ij}, \min\{i,L\})$ precisely in its first vertex. Then we extend W_L by a shortest W_L–P_{ij} path and a suitable subpath of P_{ij} so that W_L ends in T'_{iL} (or in $T'_{(i-1)L}$ if the shortest W_L–P_{ij} path has distance $< \min\{i,L\}$ from $T'_{(i-1)L}$). Otherwise, we extend W_L by a shortest W_L–$T'_{i'_L L}$ path. Analogously, we modify the end of W_L. Let W'_L be the path which we obtain in this way from W_L and let $i_L, j_L \in \mathbb{N}$ be such that W'_L starts in $T'_{i_L L}$ and ends in $T''_{j_L L}$. Further, let w^0_L, w^1_L be the endvertices of W'_L on $T'_{i_L L}$ and $T''_{j_L L}$, respectively. Then, for all $L \in \mathbb{N}$,

(α) W'_L starts at $w^0_L \in T'_{i_L L}$ and ends at $w^1_L \in T''_{j_L L}$,

(β) W'_L avoids $B_G(R_1 \cup R_3, 2L)$, and

(γ) $d_G(W'_L, T'_{kL} \cup T''_{kL}) \geqslant \min\{k, L\}$ for all $k \neq i_L, j_L \in \mathbb{N}$.

Since W'_L avoids $B_G(R_1 \cup R_3, 2L)$ but is eventually contained in $B_G(R_1 \cup R_3, L')$ for some $L' > 2L$ because W'_L is finite, we may assume, by passing to a subsequence if necessary, that the W'_L are pairwise disjoint. Moreover, by Corollary 13.3.8 (i), every $S^i_{\geqslant 0}$ is contained in $B_G(R_1, L)$ for some

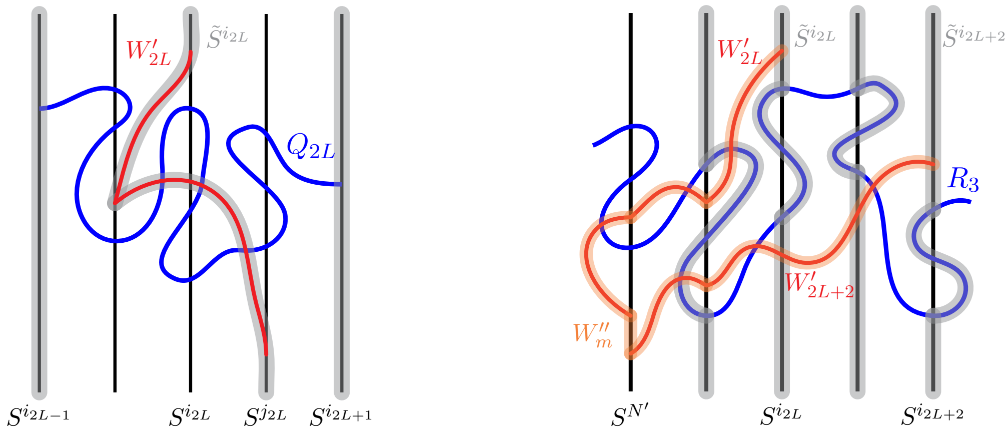

(a) Case 2a (b) Case 2b

FIGURE 13.7: The new double rays $\tilde{S}^{i_L}$ in Case 2a and 2b (indicated in grey), which again form the vertical double rays of an escaping subdivision $\tilde{H}$ of the hexagonal half-grid. In Case 2a, subpaths Q_{2L} of R_3 are fat $\tilde{H}$-paths, while in Case 2b, the paths W''_m are fat $\tilde{H}$-paths.

large enough $L \in \mathbb{N}$ and hence, by (β) and by once again passing to a subsequence, we may assume that

(δ) $i_1 < i_2 < \dots$.

Finally, by construction, W'_L can have distance $< \min\{i, L\}$ to some P_{ij} only if $i \in \{i_L, i_L + 1, j_L, j_L + 1\}$. So since G is locally finite and because of (α) and (δ), we may assume, by deleting at most finitely many P_{ij} for every $i \in \mathbb{N}$ and applying Proposition 13.3.6, that

(ε) $d_G(W'_L, P_{ij}) > \min\{i, L\}$ for all P_{ij} with $i \notin \{j_L, j_L + 1\}$.

We now distinguish two cases.

Case 2a: *There is a sequence $L_1 < L_2 < \dots \in \mathbb{N}$ such that $j_{L_1} < j_{L_2} < \dots$.*

Since also $i_1 < i_2 < \dots$ by (δ), we may assume, by passing to a subsequence of the W'_L, that $i_L, j_L < i_{L'}, j_{L'}$ for all $L < L' \in \mathbb{N}$. Then we obtain, similar as to (ε), that

(ζ) $d_G(W'_L, P_{ij}) > \min\{i, L\}$ also for all P_{ij} with $i \in \{j_L, j_L + 1\}$.

Since every $\tilde{T}_i := S^i \setminus (T'_{ii} \cup T''_{ii})$ is finite, it follows that $\tilde{T}_i \subseteq G[R_3, L']$ for some $L' \in \mathbb{N}$. So for every $i \in \mathbb{N}$ we have by (β) and (γ) that $d_G(S^i, W'_L) \geqslant i$ for all large enough L. Conversely, since every W'_L is finite and by (i), there exist for every W'_L at most finitely many S^i such that $d_G(W'_L, S^i) < i$. Hence, we may assume, by passing to a subsequence of the W'_L, that every S^i has distance $< i$ to at most one W'_L. Indeed, we may pick some W'_{L_1} that has distance < 1 to S^1 (if such a W'_{L_1} exists; otherwise we choose W'_{L_1} arbitrarily) and delete all other W'_L that have distance < 1 to S^1. Then there are still infinitely many W'_L left, and also there are infinitely many S^i left that do not have distance $< i$ to W'_{L_1}; so we may pick a new path W'_L for the next such S^i, and so on. We enumerate these W'_{L_i} again by W'_i. In particular, it follows that $d_G(W'_L, S^{i_{L'}}) \geqslant i_{L'}$ and $d_G(W'_L, S^{j_{L'}}) \geqslant j_{L'}$ for all $L \neq L' \in \mathbb{N}$.

For every $2L \in \mathbb{N}$, we set $\tilde{S}^{i_{2L}} := S^{j_{2L}} w^1_{2L} W'_{2L} w^0_{2L} S^{i_{2L}}$ (see Figure 13.7 a). Since every $\tilde{S}^{i_{2L}}$ has a tail in $S^{i_{2L}}_{\geqslant 0}$ and in $S^{j_{2L}}_{\leqslant 0}$ and because of (ε) and (ζ), we can find in H infinitely many $S^{i_{2L-1}}$–$\tilde{S}^{i_{2L}}$ paths and infinitely many $\tilde{S}^{i_{2L}}$–$S^{i_{2L+1}}$ paths that make the $\tilde{S}^{i_{2L}}$ and the $S^{i_{2L+1}}$ into an escaping subdivision $\tilde{H}$ of the hexagonal half-grid. By (a) and (b), R_3 contains for every $2L$ a subpath Q_{2L} that starts in $B_G(S^{i_{2L-1}}, i_{2L-1})$, ends in $B_G(S^{i_{2L+1}}, i_{2L+1})$ and is otherwise disjoint from all $B_G(S^{i_{L'}}, i_{L'})$ and $B_G(S^{j_{L'}}, j_{L'})$ with $L' \neq 2L \in \mathbb{N}$. Since $d_G(R_3, W'_{2L'}) \geqslant 2L'$ by (β), it follows that also $d_G(Q_{2L}, \tilde{S}^{i_{2L'}}) \geqslant 2L'$ for all $L' \neq L \in \mathbb{N}$. Moreover, by the definition of $\tilde{S}^{i_{2L}}$ and because of (β) and (c), we also have $d_G(Q_{2L}, \tilde{S}^{i_{2L}}) \geqslant 2L$. Hence, extending the Q_{2L} by shortest paths to $S^{i_{2L-1}}$–$S^{i_{2L+1}}$ paths yields fat $\tilde{H}$-paths as in Lemma 13.6.5. This implies that G contains $K_{\aleph_0}$ as an ultra fat minor, and hence concludes this case of the proof.

Case 2b: *There is some $N \in \mathbb{N}$ such that $j_L \leqslant N$ for infinitely many $L \in \mathbb{N}$.*

By passing to a subsequence of the W'_L, we may assume that $j_L = N'$ for some $N' \in \mathbb{N}$ and all $L \in \mathbb{N}$. We again modify H as follows. For every $i \in \mathbb{N}$, let $\tilde{S}^i$ be obtained from S^i by replacing the subpath $\tilde{T}_i := S^i \setminus (T'_{ii} \cup T''_{ii})$ of S^i by a path Q_i that consists of a suitable subpath of R_3 together with shortest T'_{ii}–R_3 and R_3–T''_{ii} paths (see Figure 13.7 b). Then the $\tilde{S}^i$ are again double rays. By (i) and (b) and because G is locally finite and the Q_i are finite and pairwise disjoint, there are $N' < i_{L_1} < i_{L_2} < \ldots \in \mathbb{N}$ such that the $\tilde{S}^{i_{L_j}}$ again satisfy (i) with $M_{i_{L_j}}$ updated to $M_{i_{L(j+1)}-1}$. Since every $\tilde{S}^i$ still has a tail in $S^i_{\geqslant 0}$ and in $S^i_{\leqslant 0}$ and because of (iv), we can find in H, for every $j \in \mathbb{N}$, infinitely many $\tilde{S}^{i_{L_j}}$–$\tilde{S}^{i_{L_{j+1}}}$ paths that make the $\tilde{S}^{i_{L_j}}$ into an escaping subdivision $\tilde{H}$ of the hexagonal half-grid. Since all W'_{L_j} end in $S^{N'}$, we can connect the $W'_{L_{2j}}$ pairwise by suitable subpaths of $S^{N'}$, to obtain infinitely many pairwise disjoint $\tilde{H}$-paths W''_m. Note that a path W''_m obtained from W'_{L_i} and W'_{L_j} for $i, j \in 2\mathbb{N}$ is an $\tilde{H}$-path that starts in $\tilde{S}^{i_{L_i}}$ and ends in $\tilde{S}^{i_{L_j}}$. Moreover, by (β), (γ) and (ε) and because $\tilde{H}$ is escaping, W''_m is $\min\{L_i, L_j\}$-fat. Hence, by (δ) and because we only used the paths W'_{L_i} with $i \in 2\mathbb{N}$ to construct the W''_m, infinitely many of the W''_m yield fat $\tilde{H}$-paths as in Lemma 13.6.5, which implies that G contains $K_{\aleph_0}$ as an ultra fat minor, and which thus concludes the proof. □

We are now in a position to prove Theorem 13.6.1.

Proof of Theorem 13.6.1. The assertion follows immediately by first applying Lemmas 13.6.9 and 13.6.8 and then Lemma 13.6.2. □

13.7 Further comments

Before we discuss other related topics to this chapter, let us note that our main results give only partial answers to the problems that we mentioned in the introduction, that is to [75, Problem 7.3] and [74, Problems 4.1 and 4.2]. Hence, these problems are still open for arbitrary finitely generated groups that need not be finitely presented and, more generally, for locally finite, quasi-transitive graphs without any restrictions on their cycle spaces.

13.7.1 Coarse embeddings

Theorem 36' asserts that we can find a diverging subdivision of the hexagonal full-grid in every locally finite, quasi-transitive graph whose cycle space is generated by cycles of bounded length and that has a thick end. The advantage of diverging subdivisions over arbitrary subdivisions is that they preserve some of the geometry of the original graph. One might wish to strengthen Theorem 36', by asking for a subdivision of the hexagonal full-grid whose geometry is even closer related to the geometry of G.

For two graphs G and H, a map $f\colon V(H) \to V(G)$ is a *coarse embedding* if there exist functions $\varrho^-\colon [0,\infty) \to [0,\infty)$ and $\varrho^+\colon [0,\infty) \to [0,\infty)$ such that $\varrho^-(a) \to \infty$ for $a \to \infty$ and

$$\varrho^-(d_H(u,v)) \leqslant d_G(f(u),f(v)) \leqslant \varrho^+(d_H(u,v))$$

for all $u, v \in V(H)$. It is easy to check that a coarse embedding of the hexagonal full-grid always yields a diverging subdivision; however, conversely, a diverging subdivision is in general much weaker than a coarse embedding. One may thus ask whether we can always find a coarse embedding of the hexagonal full-grid in a locally finite, quasi-transitive graph with a thick end.

However, it was already discussed in [74] that for arbitrary locally finite, quasi-transitive graphs (without any condition on their cycle spaces) we cannot ask for coarse embeddings of the hexagonal full-grid. Indeed, coarse embeddings preserve the asymptotic dimension[3], that is, if H has asymptotic dimension at least n and H is coarsely embeddable into G, then the asymptotic dimension of G is at least n, too. Since every locally finite Cayley graph of the lamplighter group has asymptotic dimension 1, see Gentimis [73], and has a thick end, but the full-grid has asymptotic dimension 2, we cannot ask for coarse embeddings of the hexagonal full-grid into all locally finite, quasi-transitive graphs with thick ends.

However, a special case of a theorem by Fujiwara and Whyte [71] states that every locally finite, quasi-transitive graph with a thick end whose cycle space is generated by cycles of bounded length has asymptotic dimension at least 2. Thus, the asymptotic dimension of the full-grid does

[3] See e.g. [23] for a definition of the asymptotic dimension.

not prevent it from being coarsely embeddable into such graphs. This motivates the following problem.

Problem 13.7.1. *Let G be a locally finite, quasi-transitive graph whose cycle space is generated by cycles of bounded length and that has a thick end. Is the hexagonal full-grid coarsely embeddable into G?*

Note that a positive answer to this question would also yield a positive answer to [74, Problem 4.5].

13.7.2 Quasi-isometries to trees

As we have discussed in the introduction, a result by Krön and Möller [96, Theorem 5.5] asserts that a locally finite, quasi-transitive, connected graph is quasi-isometric to a tree if and only if it has no thick end. Hence, we obtain the following corollary from Theorems 35 and 36, which yields two new characterisations of quasi-transitive, locally finite, connected graphs that are quasi-isometric to trees for the special case that the cycle space is generated by cycles of bounded length.

Corollary 13.7.2. *Let G be a locally finite, quasi-transitive, connected graph whose cycle space is generated by cycles of bounded length. Then the following are equivalent:*

(i) *G has a thick end.*

(ii) *G contains the full-grid as an asymptotic minor.*

(iii) *G contains the full-grid as a diverging minor.*

(iv) *G is not quasi-isometric to a tree.* □

For further characterisations of quasi-transitive, locally finite, connected graphs that are quasi-isometric to trees, we refer the reader to [14, 87, 96, 124].

13.7.3 Quasi-isometries to planar graphs

Finally, we would like to draw the reader's attention to another related problem, which is still open. As we discussed in the previous subsection, the (global) geometry of locally finite, quasi-transitive graphs without a thick end is well understood as they are quasi-isometric to forests. In that sense, our results, Theorems 35 and 36, can be seen as a step towards understanding the (global) geometry of the remaining locally finite, quasi-transitive graphs – those with a thick end. In the case where the cycle space of a locally finite, quasi-transitive graph G is generated by cycles of bounded length, we showed that G contains the full-grid as an asymptotic minor. Since asymptotic minors cannot hide in a ball of small radius, they will appear in the global structure of G. However, even if G is one-ended, this does not mean that the geometry of G resembles that of the full-grid or, more generally, of a one-ended, planar graph. The reason for this is simple: the global structure of G may be far more involved, and may contain the full-grid only as a substructure. Indeed, even in our proof, we might have found the asymptotic full-grid inside an asymptotic minor of the infinite complete graph. However, Georgakopoulos and Papasoglu conjectured that this is in fact the only thing that can happen.

Conjecture 13.7.3. [75, Conjecture 9.3] *Let G be a locally finite, transitive graph. Then G either is quasi-isometric to a planar graph or contains every finite graph as an asymptotic minor.*

Note that this can be seen as a coarse version of Thomassen's [122] result that every locally finite, one-ended, transitive graph is either planar or can be contracted into the infinite complete graph.

MacManus [99] proved Conjecture 13.7.3 in the special case where G is a locally finite Cayley graph of a finitely presented group. Recall that every Cayley graph of a finitely presented group is transitive and has a cycle space which is generated by cycles of bounded length. So the assumption on the cycle space, which was already crucial for our proofs of Theorems 35 to 38, reappears here.

We remark that MacManus's proof uses deep group-theoretic results that have no counterpart for quasi-transitive graphs. Hence, Conjecture 13.7.3 is still open for arbitrary locally finite, quasi-transitive graphs whose cycle space is generated by cycles of bounded length.

References

[1] T. Abrishami, R. W. Jacobs, P. Knappe, and J. Kobler, *Locally chordal graphs*, arXiv preprint (2025), available at arXiv:2501.17320.

[2] R. Aharoni and E. Berger, *Menger's theorem for infinite graphs*, Inventiones mathematicae **176** (2009), no. 1, 1–62.

[3] S. Albrechtsen, *Optimal trees of tangles: refining the essential parts*, arXiv preprint (2023), available at arXiv:2304.12078.

[4] ———, *Refining trees of tangles in abstract separation systems: inessential parts*, Combinatorial Theory **4(1)** (2024).

[5] ———, *Tangle-tree duality in infinite graphs*, arXiv preprint (2024), available at arXiv:2405.06756.

[6] ———, *Refining tree-decompositions so that they display the k-blocks*, Journal of Graph Theory **109** (2025), no. 3, 310–314.

[7] S. Albrechtsen, H. von Bergen, R. W. Jacobs, P. Knappe, and P. Wollan, *On vertex sets inducing tangles*, arXiv preprint (2024), available at arXiv:2411.13656.

[8] S. Albrechtsen, R. Diestel, A.-K. Elm, E. Fluck, R. W. Jacobs, P. Knappe, and P. Wollan, *A structural duality for path-decompositions into parts of small radius*, arXiv preprint (2023), available at arXiv:2307.08497v1.

[9] S. Albrechtsen and M. Hamann, *Asymptotic half-grid and full-grid minors*, arXiv preprint (2024), available at arXiv:2412.15675.

[10] S. Albrechtsen, T. Huynh, R. W. Jacobs, P. Knappe, and P. Wollan, *A Menger-type theorem for two induced paths*, SIAM Journal on Discrete Mathematics **38** (2024), no. 2, 1438–1450.

[11] S. Albrechtsen, R. W. Jacobs, P. Knappe, and M. Pitz, *Counterexamples regarding linked and lean tree-decompositions of infinite graphs*, arXiv preprint (2024), available at arXiv:2405.06755.

[12] ———, *Linked tree-decompositions into finite parts*, arXiv preprint (2024), available at arXiv:2405.06753.

[13] S. Albrechtsen, R. W. Jacobs, P. Knappe, and P. Wollan, *A characterisation of graphs quasi-isometric to K_4-minor-free graphs*, arXiv preprint (2024), available at arXiv:2408.15335.

[14] Y. Antolín, *On Cayley graphs of virtually free groups*, Groups. Complexity. Cryptology **3** (2011), no. 2, 301–327.

[15] G. Bell and A. Dranishnikov, *Asymptotic dimension*, Topology and its Applications **155** (2008), 1265–1296.

[16] P. Bellenbaum and R. Diestel, *Two short proofs concerning tree-decompositions*, Combinatorics, Probability and Computing **11** (2002), no. 6, 541–547.

[17] R. Belmonte, F. V. Fomin, P. A. Golovach, and M. S. Ramanujan, *Metric dimension of bounded tree-length graphs*, SIAM Journal on Discrete Mathematics **31** (2017), no. 2, 1217–1243.

[18] O. Bendele and D. Rautenbach, *Additive tree $O(\varrho \log n)$-spanners from tree breadth ϱ*, Theoretical Computer Science **914** (2022), 39–46.

[19] E. Berger and P. D. Seymour, *Bounded-diameter tree-decompositions*, Combinatorica (2024), 1–16.

[20] D. Bienstock, *On the complexity of testing for odd holes and induced odd paths*, Discrete Mathematics **90** (1991), no. 1, 85–92.

[21] ———, *Corrigendum: "On the complexity of testing for odd holes and induced odd paths"*, Discrete Mathematics **102** (1992), no. 1, 109.

[22] G. Birkhoff, *Lattice theory*, revised edition, Colloquium publications, American Mathematical Society, 1948.

[23] M. Bonamy, N. Bousquet, L. Esperet, C. Groenland, C.-H. Liu, F. Pirot, and A. Scott, *Asymptotic dimension of minor-closed families and Assouad–Nagata dimension of surfaces*, Journal of the European Mathematical Society **26** (2023), no. 10, 3739–3791.

[24] M. Bonamy, J. Davies, L. Esperet, and A. Wesolek, *Personal communication*, 2024.

[25] J. P. Burling, *On coloring problems of families of prototypes*, PhD dissertation, University of Colorado, Boulder (1965).

[26] J. Carmesin, *All graphs have tree-decompositions displaying their topological ends*, Combinatorica **39** (2019), no. 3, 545–596.

[27] 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** (2016), 1–24.

[28] ———, *Canonical tree-decompositions of finite graphs II. Essential parts*, Journal of Combinatorial Theory, Series B **118** (2016), 268–283.

[29] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein, *Connectivity and tree structure in finite graphs*, Combinatorica **34** (2014), no. 1, 1–35.

[30] J. Carmesin and J. P. Gollin, *Canonical tree-decompositions of a graph that display its k-blocks*, Journal of Combinatorial Theory, Series B **122** (2017), 1–20.

[31] J. Carmesin, M. Hamann, and B. Miraftab, *Canonical trees of tree-decompositions*, Journal of Combinatorial Theory, Series B **152** (2022), 1–26.

[32] J. Carmesin, R. W. Jacobs, P. Knappe, and J. Kurkofka, *Canonical graph decompositions and local separations: From infinite coverings to a finite combinatorial theory*, arXiv preprint (2025), available at arXiv:2501.16170.

[33] J. Carmesin and J. Kurkofka, *Characterising 4-tangles through a connectivity property*, arXiv preprint (2023), available at arXiv:2309.00902.

[34] ———, *Entanglements*, Journal of Combinatorial Theory, Series B **164** (2024), 17–28.

[35] J. Carmesin, F. Lehner, and R. G. Möller, *On tree-decompositions of one-ended graphs*, Mathematische Nachrichten **292** (2019), no. 3, 524–539.

[36] V. Chepoi, F. Dragan, B. Estellon, M. Habib, and Y. Vaxès, *Diameters, centers, and approximating trees of delta-hyperbolic geodesic spaces and graphs*, Proceedings of the twenty-fourth annual symposium on computational geometry, 2008, pp. 59–68.

[37] V. Chepoi, F. Dragan, I. Newman, Y. Rabinovich, and Y. Vaxes, *Constant approximation algorithms for embedding graph metrics into trees and outerplanar graphs*, Discrete & Computational Geometry **47** (2012), 187–214.

[38] J. Davies, R. Hickingbotham, F. Illingworth, and R. McCarty, *Fat minors cannot be thinned (by quasi-isometries)*, arXiv preprint (2024), available at arXiv:2405.09383.

[39] R. Diestel, *The depth-first search tree structure of $TK^{\aleph_0}$-free graphs*, Journal of Combinatorial Theory, Series B **61** (1994), no. 2, 260–262.

[40] ———, *Ends and tangles*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **87**, Special issue in memory of Rudolf Halin (2017), 223–244.

[41] ———, *Graph Theory*, 5th ed., Springer, 2017.

[42] ———, *Abstract separation systems*, Order **35** (2018), 157–170.

[43] ———, *Tree sets*, Order **35** (2018), 171–192.

[44] ———, *Tangles: a structural approach to artificial intelligence in the empirical sciences*, Cambridge University Press, 2024.

[45] R. Diestel, P. Eberenz, and J. Erde, *Duality theorems for blocks and tangles in graphs*, SIAM Journal on Discrete Mathematics **31** (2017), no. 3, 1514–1528.

[46] R. Diestel, C. Elbracht, and R. W. Jacobs, *Point sets and functions inducing tangles of set separations*, Journal of Combinatorics **15** (2024), no. 3, 283–306.

[47] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39** (2019), 37–75.

[48] R. Diestel, R. W. Jacobs, P. Knappe, and J. Kurkofka, *Canonical graph decompositions via coverings*, arXiv preprint (2022), available at arXiv:2207.04855.

[49] ———, *Canonical graph decompositions via coverings, extended version*, arXiv preprint (2022), available at arXiv:2207.04855v5.

[50] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, Journal of Combinatorial Theory, Series B **87** (2003), 197–206.

[51] ———, *Graph minor hierarchies*, Discrete Applied Mathematics **145** (2005), 167–182.

[52] R. Diestel and I. Leader, *A conjecture concerning a limit of non-Cayley graphs*, Journal of Algebraic Combinatorics **14** (2001), no. 1, 17–25.

[53] ———, *Normal spanning trees, Aronszajn trees and excluded minors*, Journal of the London Mathematical Society **63** (2001), no. 1, 16–32.

[54] R. Diestel and S.-i. Oum, *Tangle-tree duality: in graphs, matroids and beyond*, Combinatorica **39** (2019), no. 4, 879–910.

[55] ———, *Tangle-tree duality in abstract separation systems*, Advances in Mathematics **377** (2021), 107470.

[56] Y. Dourisboure and C. Gavoille, *Tree-decompositions with bags of small diameter*, Discrete Mathematics **307** (2007), no. 16, 2008–2029.

[57] F. F. Dragan and E. Köhler, *An approximation algorithm for the tree t-spanner problem on unweighted graphs via generalized chordal graphs*, Algorithmica **69** (2014), 884–905.

[58] M. J. Dunwoody, *An inaccessible group*, Geometric Group Theory (London Mathematical Society Lecture Note Series), 1993, pp. 75–78.

[59] ———, *An inaccessible graph*, Random walks, boundaries and spectra, 2011, pp. 1–14.

[60] P. Eberenz, *Characteristics of profiles*, MSc dissertation, Hamburg (2015). Available online.

[61] C. Elbracht, *Tangles determined by majority vote*, MSc dissertation, Hamburg (2017). Available online.

[62] C. Elbracht, J. Kneip, and M. Teegen, *Tangles are decided by weighted vertex sets*, Advances in Combinatorics (2020).

[63] ———, *Trees of tangles in abstract separation systems*, Journal of Combinatorial Theory, Series A **180** (2021), 105425.

[64] ———, *Trees of tangles in infinite separation systems*, Mathematical Proceedings of the Cambridge Philosophical Society **173** (2022), no. 2, 297–327.

[65] A.-K. Elm and J. Kurkofka, *A tree-of-tangles theorem for infinite tangles*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **92** (2022), 139–178.

[66] J. Erde, *Refining a tree-decomposition which distinguishes tangles*, SIAM Journal on Discrete Mathematics **31** (2017), 1529–1551.

[67] ———, *A unified treatment of linked and lean tree-decompositions*, Journal of Combinatorial Theory, Series B **130** (2018), 114–143.

[68] A. Eskin, D. Fisher, and K. Whyte, *Coarse differentiation of quasi-isometries I: Spaces not quasi-isometric to Cayley graphs*, Annals of Mathematics. Second Series **176** (2012), no. 1, 221–260.

[69] L. Esperet and U. Giocanti, *Coarse geometry of quasi-transitive graphs beyond planarity*, Electronic Journal of Combinatorics **31(2)** (2024).

[70] K. Fujiwara and P. Papasoglu, *A coarse-geometry characterization of cacti*, arXiv preprint (2023), available at arXiv:2305.08512.

[71] K. Fujiwara and K. Whyte, *A note on spaces of asymptotic dimension one*, Algebraic & Geometric Topology **7** (2007), 1063–1070.

[72] P. Gartland, T. Korhonen, and D. Lokshtanov, *On induced versions of Menger's theorem on sparse graphs*, arXiv preprint (2023), available at arXiv:2309.08169.

[73] T. Gentimis, *Asymptotic dimension of finitely presented groups*, Proceedings of the American Mathematical Society **136** (2008), no. 12, 4103–4110.

[74] A. Georgakopoulos and M. Hamann, *A full Halin grid theorem*, Discrete & Computational Geometry (2024).

[75] A. Georgakopoulos and P. Papasoglu, *Graph minors and metric spaces*, to appear in Combinatorica (2023), available at arXiv:2305.07456.

[76] J. P. Gollin, M. Hatzel, and S. Wiederrecht, *On graphs with a simple structure of maximal cliques*, arXiv preprint (2025), available at arXiv:2504.16863.

[77] J. P. Gollin and K. Heuer, *Characterising k-connected sets in infinite graphs*, Journal of Combinatorial Theory, Series B **157** (2022), 451–499.

[78] M. Grohe, *Tangles and connectivity in graphs*, Language and Automata Theory and Applications (2016), 24–41, available at arXiv:1602.04727.

[79] M. Grohe and P. Schweitzer, *Isomorphism testing for graphs of bounded rank width*, IEEE 56th Annual Symposium on Foundations of Computer Science (2015), 1010–1029.

[80] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric Group Theory: Proceedings of the Symposium held in Sussex, 1991, vol. 2, Cambridge University Press, 1993.

[81] R. Halin, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Mathematische Nachrichten **30** (1965), 63–85.

[82] ———, *Some path problems in graph theory*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **44** (1975), no. 1, 175–186.

[83] ———, *Systeme disjunkter unendlicher Wege in Graphen*, Numerische Methoden bei Optimierungsaufgaben Band 3: Optimierung bei graphentheoretischen und ganzzahligen Problemen (1977), 55–67.

[84] ______, *Simplicial decompositions of infinite graphs*, Annals of Discrete Mathematics **3** (1978), 93–109.

[85] ______, *Tree-partitions of infinite graphs*, Discrete Mathematics **97** (1991), no. 1-3, 203–217.

[86] M. Hamann, *Accessibility in transitive graphs*, Combinatorica **38** (2018), 847–859.

[87] M. Hamann, F. Lehner, B. Miraftab, and T. Rühmann, *A Stallings type theorem for quasi-transitive graphs*, Journal of Combinatorial Theory, Series B **157** (2022), 40–69.

[88] K. Hendrey, S. Norin, R. Steiner, and J. Turcotte, *On an induced version of Menger's theorem*, Electronic Journal of Combinatorics **31** (2024), no. 4.

[89] K. Heuer, *Excluding a full grid minor*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **87** (2017), no. 2, 265–274.

[90] R. W. Jacobs and P. Knappe, *Efficiently distinguishing all tangles in locally finite graphs*, Journal of Combinatorial Theory, Series B **167** (2024), 189–214.

[91] H. A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, Mathematische Nachrichten **41** (1969), 1–22.

[92] K.-i. Kawarabayashi, S. Norine, R. Thomas, and P. Wollan, *K_6 minors in 6-connected graphs of bounded tree-width*, Journal of Combinatorial Theory, Series B **136** (2012).

[93] M. Koloschin, T. Krill, and M. Pitz, *End spaces and tree-decompositions*, Journal of Combinatorial Theory, Series B **161** (2023), 147–179.

[94] P. Komjáth, *A note on minors of uncountable graphs*, Mathematical Proceedings of the Cambridge Philosophical Society **117** (1995), no. 1, 7–9.

[95] I. Kříž and R. Thomas, *The menger-like property of the tree-width of infinite graphs*, Journal of Combinatorial Theory, Series B **52** (1991), no. 1, 86–91.

[96] B. Krön and R. G. Möller, *Quasi-isometries between graphs and trees*, Journal of Combinatorial Theory, Series B **98** (2008), no. 5, 994–1013.

[97] J. Kurkofka, R. Melcher, and M. Pitz, *A strengthening of Halin's grid theorem*, Mathematika **68** (2022), no. 4, 1009–1013.

[98] A. Leitert, *Tree-breadth of graphs with variants and applications*, Ph.D. Thesis, 2017.

[99] J. P. MacManus, *Fat minors in finitely presented groups*, arXiv preprint (2024), available at arXiv:2408.10748.

[100] J. F. Manning, *Geometry of pseudocharacters*, Geometry & Topology **9** (2005), no. 2, 1147–1185.

[101] K. Menger, *Zur allgemeinen Kurventheorie*, Fundamenta Mathematicae **10** (1927), 96–115.

[102] T. Nguyen, A. Scott, and P. D. Seymour, *Coarse tree-width*, arXiv preprint (2025), available at arXiv:2501.09839.

[103] T. H. Nguyen, A. Scott, and P. D. Seymour, *A counterexample to the coarse Menger conjecture*, Journal of Combinatorial Theory, Series B **173** (2025), 68–82.

[104] M. Pitz, *A unified existence theorem for normal spanning trees*, Journal of Combinatorial Theory, Series B **145** (2020), 466–469.

[105] ______, *Proof of Halin's normal spanning tree conjecture*, Israel Journal of Mathematics **246** (2021), no. 1, 353–370.

[106] ______, *A note on minor antichains of uncountable graphs*, Journal of Graph Theory **102** (2023), no. 3, 552–555.

[107] N. Polat, *Ends and multi-endings I*, Journal of Combinatorial Theory, Series B **67** (1996), 86–110.

[108] P. Pournajafi and N. Trotignon, *Burling graphs revisited, part I: New characterizations*, European Journal of Combinatorics **110** (2023), 103686.

[109] ———, *Burling graphs revisited, part II: Structure*, European Journal of Combinatorics **116** (2024), 103849.

[110] N. Robertson and P. D. Seymour, *Graph minors I–XX*, Journal of Combinatorial Theory, Series B (1983–2004).

[111] ———, *Graph minors. II. Algorithmic aspects of tree-width*, Journal of Algorithms **7** (1986), 309–322.

[112] ———, *Graph minors. V. Excluding a planar graph*, Journal of Combinatorial Theory, Series B **41** (1986), 92–114.

[113] ———, *Graph minors. IV. Tree-width and well-quasi-ordering*, Journal of Combinatorial Theory, Series B **48** (1990), 227–254.

[114] ———, *Graph minors. X. Obstructions to tree-decomposition*, Journal of Combinatorial Theory, Series B **52** (1991), 153–190.

[115] N. Robertson, P. D. Seymour, and R. Thomas, *Excluding infinite clique minors*, Memoirs of the American Mathematical Society **118** (1995), no. 566, vi+103.

[116] P. D. Seymour and R. Thomas, *Excluding infinite trees*, Transactions of the American Mathematical Society **335** (1993), no. 2, 597–630.

[117] ———, *Graph searching and a min-max theorem for tree-width*, Journal of Combinatorial Theory, Series B **58** (1993), no. 1, 22–33.

[118] R. Thomas, *A counter-example to 'Wagner's conjecture' for infinite graphs*, Mathematical Proceedings of the Cambridge Philosophical Society **103** (1988), no. 1, 55–57.

[119] ———, *Well-quasi-ordering infinite graphs with forbidden finite planar minor*, Transactions of the American Mathematical Society **312** (1989), no. 1, 279–313.

[120] ———, *A menger-like property of tree-width: The finite case*, Journal of Combinatorial Theory, Series B **48** (1990), 67–76.

[121] C. Thomassen, *Configurations in graphs of large minimum degree, connectivity, or chromatic number*, Annals of the New York Academy of Sciences **555** (1989).

[122] ———, *The Hadwiger number of infinite vertex-transitive graphs*, Combinatorica **12** (1992), no. 4, 481–491.

[123] C. Thomassen and W. Woess, *Vertex-transitive graphs and accessibility*, Journal of Combinatorial Theory, Series B **58** (1993), no. 2, 248–268.

[124] W. Woess, *Graphs and groups with tree-like properties*, Journal of Combinatorial Theory, Series B **47** (1989), no. 3, 361–371.

Appendix

A Summary

This thesis deals with questions and problems from structural graph theory, a branch of discrete mathematics. Reoccurring objects are tangles and tree-decompositions and the more general graph-decompositions.

The thesis consists of three parts. The first two parts consider tangles and tree-decompositions – first in infinite and then in finite graphs. The third part deals with questions from coarse graph theory.

Part I:

Kříž and Thomas [95, 120] showed that every (finite or infinite) graph of tree-width $k \in \mathbb{N}$ admits a lean tree-decomposition of width k. In Chapter 4 we discuss a number of counterexamples demonstrating the limits of possible generalisations of their result to arbitrary infinite tree-width. In particular, we construct a locally finite, planar, connected graph that has no lean tree-decomposition.

On the positive side, we obtain in Chapter 3 a version of Kříž and Thomas's result for graphs of finite tree-width: *Every graph which admits a tree-decomposition into finite parts has a rooted tree-decomposition into finite parts that is linked, tight and componental.*

As an application, we obtain that every graph without half-grid minor has a lean tree-decomposition into finite parts. In particular, it follows that every graph without half-grid minor has a tree-decomposition which efficiently distinguishes all ends and critical vertex sets, strengthening results by Carmesin [26] and by Elm and Kurkofka [65] for this graph class.

As a second application, it follows that every graph G which admits a tree-decomposition into finite parts has a tree-decomposition into finite parts that displays all the ends of G and their combined degrees, resolving a question of Halin [83] from 1977. This latter tree-decomposition yields short, unified proofs of the characterisations due to Robertson, Seymour and Thomas [115, 116] of graphs without half-grid minor, and of graphs without binary tree subdivision.

A third application is presented in Chapter 5. There, we extend Robertson and Seymour's [114] tangle-tree duality theorem to infinite graphs.

Part II:

Chapter 6: We combine the two fundamental fixed-order tangle theorems of Robertson and Seymour [114] into a single theorem that implies both, in a best possible way. We show that, for

every $k \in \mathbb{N}$, every tree-decomposition of a graph G which efficiently distinguishes all its k-tangles can be refined to a tree-decomposition whose parts are either too small to be home to a k-tangle, or as small as possible while being home to a k-tangle.

Chapter 7: Carmesin and Gollin [30] proved that, for every $k \in \mathbb{N}$, every finite graph has a canonical tree-decomposition $(T, \mathcal{V})$ of adhesion less than k that efficiently distinguishes every two distinct k-profiles, and which has the further property that every separable k-block is equal to the unique part of $(T, \mathcal{V})$ in which it is contained.

We give a shorter proof of this result by showing that such a tree-decomposition can in fact be obtained from any canonical tight tree-decomposition of adhesion less than k. For this, we decompose the parts of such a tree-decomposition by further tree-decompositions. As an application, we also obtain a generalization of Carmesin and Gollin's result to locally finite graphs.

Chapter 8: Diestel, Hundertmark and Lemanczyk [47] asked whether every k-tangle in a graph is induced by a set of vertices by majority vote. We reduce their question to graphs whose size is bounded by a function in k. Additionally, we show that if for any fixed k this problem has a positive answer, then every k-tangle is induced by a vertex set whose size is bounded in k. More generally, we prove for all k that every k-tangle in a graph G is induced by a weight function $V(G) \to \mathbb{N}$ whose total weight is bounded in k. As the key step of our proofs, we show that any given k-tangle in a graph G is the lift of a k-tangle in some topological minor of G whose size is bounded in k.

Part III:

Chapter 10: We give an approximate Menger-type theorem for when a graph G contains two X–Y paths P_1 and P_2 that are far apart in G. More precisely, we prove that there exists a function $f(d) \in O(d)$, such that for every graph G and $X, Y \subseteq V(G)$, either there exist two X–Y paths P_1 and P_2 such that $d_G(P_1, P_2) \geqslant d$, or there exists $v \in V(G)$ such that the ball of radius $f(d)$ around v intersects every X–Y path.

Chapters 11 and 12: We prove several special cases of the following conjecture by Georgakopoulos and Papasoglu [75]: *Let $\mathcal{X}$ be a finite set of finite graphs. Then there exists a function $f : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that, for all $K \in \mathbb{N}$, every graph with no K-fat X minor for any $X \in \mathcal{X}$ is $f(K)$-quasi-isometric to a graph with no X minor for any $X \in \mathcal{X}$.*

In Chapter 11 we show that their conjecture is true (for finite graphs) for $\mathcal{X} = \{K_3, K_{1,3}\}$ (which characterises the graphs quasi-isometric to a disjoint union of paths), $\mathcal{X} = \{K_{1,3}\}$ (cycles and paths) and $\mathcal{X} = \{K_3, W\}$ (subdivided stars), where W is the graph obtained from the disjoint union of two paths of length two by adding an edge between their inner vertices. In fact, we prove a stronger result that finds a graph in $\mathcal{X}$ even as a quasi-geodesic topological minor.

We also show in Chapter 11 that two graphs are quasi-isometric if and only if one has a decomposition modelled on the other of bounded radial width and spread. It follows that our results imply characterisations of the graphs of bounded radial path-width, cycle-width and star-width, respectively.

In Chapter 12 we verify the conjecture for $\mathcal{X} = \{K_4\}$. Our proof technique also yields a new short proof of the case $\mathcal{X} = \{K_4^-\}$ (cacti), which was first established by Fujiwara and Papasoglu [70].

Chapter 13: We prove that every locally finite, quasi-transitive graph with a thick end whose cycle space is generated by cycles of bounded length contains the full-grid as an asymptotic minor and as a diverging minor. This in particular includes all locally finite Cayley graphs of finitely presented groups, and partially solves problems of Georgakopoulos and Papasoglu [75] and of Georgakopoulos and Hamann [74].

Additionally, we show that every (not necessarily quasi-transitive) graph of finite maximum degree which has a thick end and whose cycle space is generated by cycles of bounded length contains the half-grid as an asymptotic minor and as a diverging minor.

B Deutschsprachige Zusammenfassung

Diese Arbeit beschäftigt sich mit Fragen und Problemen aus der strukturellen Graphentheorie, einem Teilgebiet der diskreten Mathematik. Wiederkehrende Objekte sind dabei vor allem Knäuel, Baumzerlegungen und allgemeinere Graphenzerlegungen.

Diese Arbeit besteht aus drei Teilen. Die ersten beiden Teile beschäftigen sich mit Knäuel und Baumzerlegungen – zunächst in unendlichen und dann in endlichen Graphen. Der dritte Teil beschäftigt sich mit Fragen aus der „groben Graphentheorie“.

Teil I:

Křížund Thomas [95, 120] haben gezeigt, dass jeder (endliche oder unendliche) Graph mit Baumweite $k \in \mathbb{N}$ eine schlanke Baumzerlegung der Weite k besitzt. In Kapitel 4 diskutieren wir einige Gegenbeispiele, welche die Grenzen möglicher Verallgemeinerungen von diesem Ergebnis auf unendliche Baumweite aufzeigen. Insbesondere konstruieren wir einen lokal endlichen, planaren, zusammenhängenden Graphen, der keine schlanke Baumzerlegung hat.

Im Gegensatz zu diesen Gegenbeispielen beweisen wir in Kapitel 3 eine Version von Kříž und Thomas' Resultat für Graphen von endlicher Baumweite: *Jeder Graph von endlicher Baumweite hat eine gewurzelte Baumzerlegung, welche verbunden, eng und komponententreu ist.*

Als eine Anwendung von diesem Resultat erhalten wir, dass jeder Graph ohne Halbgitterminor eine schlanke Baumzerlegung in endliche Teile hat. Daraus folgt insbesondere, dass jeder Graph ohne Halbgitterminor eine Baumzerlegung hat, welche alle Enden und kritischen Eckenmengen effizient unterscheidet. Dies verstärkt Resultate von Carmesin [26] und von Elm und Kurkofka [65] für diese Graphenklasse.

Als zweite Anwendung zeigen wir, dass jeder Graph G, der eine Baumzerlegung in endliche Teile hat, auch eine solche Baumzerlegung besitzt, die zusätzlich alle Enden von G und deren kombinierte Grade darstellt. Dies beantwortet nicht nur eine Frage von Halin [83] aus 1977, sondern ergibt auch kurze, einheitliche Beweise von Charakterisierungen von Robertson, Seymour und Thomas [115, 116] von Graphen ohne Halbgitterminor und von Graphen ohne Unterteilungen des unendlichen Binärbaums.

Als dritte Anwendung erweitern wir in Kapitel 5 Robertson und Seymours Knäuel-Baum-Satz auf unendliche Graphen.

Teil II:

Kapitel 6: Wir kombinieren die beiden fundamentalen Sätze über k-Knäuel von Robertson und Seymour [114] zu einem einzigen Satz, der beide impliziert – auf bestmögliche Weise. Wir zeigen, dass für jedes $k \in \mathbb{N}$, jede Baumzerlegung eines Graphen G, die alle seine k-Knäuel effizient unterscheidet, zu einer Baumzerlegung verfeinert werden kann, deren Teile alle entweder zu klein sind, um ein k-Knäuel zu enthalten, oder so klein wie möglich sind, während sie ein k-Knäuel enthalten.

Kapitel 7: Carmesin und Gollin [30] haben bewiesen, dass für jedes $k \in \mathbb{N}$ jeder endliche Graph eine kanonische Baumzerlegung $(T, \mathcal{V})$ von Adhäsion kleiner k besitzt, die alle k-Profile effizient unterscheidet und die die weitere Eigenschaft hat, dass jeder trennbare k-Block gleich dem eindeutigen Teil von $(T, \mathcal{V})$ ist, in dem er enthalten ist.

Wir liefern einen kürzeren Beweis für dieses Resultat, indem wir zeigen, dass man eine solche Baumzerlegung aus jeder kanonischen, engen Baumzerlegung von Adhäsion kleiner k erhalten kann. Hierzu zerlegen wir die Teile einer solchen Baumzerlegung mit weiteren Baumzerlegungen. Als Anwendung hiervon erhalten wir eine Verallgemeinerung von Carmesin und Gollins Resultat auf lokal endliche Graphen.

Kapitel 8: Diestel, Hundertmark und Lemanczyk [45] haben gefragt, ob jedes k-Knäuel in einem Graphen von einer Menge von Ecken durch Mehrheitswahl induziert ist. Wir reduzieren diese Frage auf Graphen, deren Größe durch eine Funktion in k beschränkt ist. Zusätzlich zeigen wir, dass, falls dieses Problem für ein festes k eine positive Antwort hat, jedes k-Knäuel von einer Menge von Ecken induziert ist, deren Größe in k beschränkt ist. Allgemeiner zeigen wir, dass für jedes k jedes k-Knäuel von einer Gewichtsfunktion $V(G) \to \mathbb{N}$ induziert wird, deren Gesamtgewicht in k beschränkt ist. Als entscheidenden Schritt in unserem Beweis zeigen wir, dass jedes k-Knäuel in einem Graphen G die Hochhebung von einem k-Knäuel in einem topologischen Minor von G ist, dessen Größe in k beschränkt ist.

Teil III:

Kapitel 10: Wir beweisen einen Menger-artigen Satz, welcher beschreibt, wann ein Graph zwei X–Y-Wege P_1 und P_2 enthält, die großen Abstand voneinander haben. Genauer gesagt zeigen wir, dass es eine Funktion $f(d) \in O(d)$ gibt, sodass für jeden Graphen G und für je zwei Mengen $X, Y \subseteq V(G)$ mindestens eine der beiden folgenden Aussagen wahr ist: Es gibt in G zwei X–Y-Wege P_1 und P_2, sodass $d_G(P_1, P_2) \geqslant d$, oder es gibt ein $v \in V(G)$, sodass der Ball von Radius $f(d)$ um v jeden X–Y-Weg trifft.

Kapitel 11 und 12: Wir beweisen einige Spezialfälle der folgenden Vermutung von Georgakopoulos und Papasoglu [75]: *Sei $\mathcal{X}$ eine endliche Menge von endlichen Graphen. Dann gibt es eine Funktion $f : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$, sodass für alle $k \in \mathbb{N}$ und jeden Graph G die folgende Aussage gilt: Wenn G kein $X \in \mathcal{X}$ als K-fetten Minor enthält, dann ist G quasi-isometrisch zu einem Graphen, der kein $X \in \mathcal{X}$ als Minor enthält.*

In Kapitel 11 zeigen wir, dass diese Vermutung (für endliche Graphen) in den folgenden Fälle wahr ist: $\mathcal{X} = \{K_3, K_{1,3}\}$ (dies charakterisiert die Graphen, welche quasi-isometrisch zu einer disjunkten Vereinigungen von Wegen sind), $\mathcal{X} = \{K_{1,3}\}$ (Kreise und Wege) und $\mathcal{X} = \{K_3, W\}$ (unterteilte Sterne). Tatsächlich zeigen wir ein stärkeres Resultat, welches einen Graphen in $\mathcal{X}$ sogar als quasi-geodätischen topologischen Minor findet.

Außerdem zeigen wir in Kapitel 11, dass zwei Graphen genau dann quasi-isometrisch zueinander sind, wenn jeder eine Graphenzerlegung entlang des anderen hat von beschränkter radialer Weite und Ausdehnung. Hieraus folgt, dass unsere Resultate Charakterisierungen von denjenigen Graphen implizieren, die beschränkte Wegweite, Kreisweite beziehungsweise Sternweite haben.

In Kapitel 12 verifizieren wir die Vermutung im Fall $\mathcal{X} = \{K_4\}$. Unsere Beweistechnik liefert auch einen neuen, kurzen Beweis des Falls $\mathcal{X} = \{K_4^-\}$ (Kakteen), welcher zuerst von Fujiwara und Papasoglu [70] bewiesen wurde.

Kapitel 13: Wir beweisen das jeder lokal endliche, quasi-transitive Graph, der ein dickes Ende hat und dessen Zyklenraum von Kreisen beschränkter Länge erzeugt wird, das Vollgitter als asymptotischen Minor und als divergierenden topologischen Minor enthält. Dies schließt insbesondere alle lokal endlichen Cayleygraphen von endlich präsentierten Gruppen mit ein und löst teilweise Probleme von Georgakopoulos und Papasoglu [75] und von Georgakopoulos und Hamann [74].

Zusätzlich zeigen wir, dass jeder (nicht notwendigerweise quasi-transitive) Graph mit endlichen Maximalgrad, der ein dickes Ende hat und dessen Zyklenraum von Kreisen beschränkter Länge erzeugt wird, das Halbgitter als asymptotischen Minor und als divergierenden topologischen Minor enthält.

C Publications related to this thesis

The following papers are related to this dissertation:

Part I:

Chapter 3 is based on [12].

Chapter 4 is based on [11].

Chapter 5 is based on [5].

Part II:

Chapter 6 is based on [3, §1-4].

Chapter 7 is based on [6].

Chapter 8 is based on [7].

Part III:

Chapter 10 is based on [10].

Chapter 11 is based on [8].

Chapter 12 is based on [13].

Chapter 13 is based on [9].

D Declaration of contributions

The research in this thesis is based on work that I have done with various co-authors (see also Chapter C). In general, the involved authors carried out the work in close co-operation and often share an equal amount of the work on the research.

In what follows, I will go into more detail on the work that went into the respective chapters of this thesis to point out where contributions differ and emphasise some of the ideas I contributed myself. In cases where I do not emphasise who did which work, it should be assumed that all involved authors contributed equally to the research.

Part I:

Chapter 3: This chapter is based on the paper [12], which is joint work with Raphael Jacobs, Paul Knappe and Max Pitz. Max Pitz started the project by asking Raphael Jacobs and Paul Knappe a question about efficiently distinguishing ends in a graph. Based on their input, Max Pitz developed an idea for an algorithm constructing linked, componental, tight tree-decompositions into finite parts of locally finite graphs. Together, Raphael Jacobs, Paul Knappe and Max Pitz worked out the details. Paul Knappe generalised the algorithm to produce, for any (not necessarily locally finite) graph of finite tree-width, a linked, componental, tight tree-decomposition into rayless parts (Section 3.7). They presented this result to the Discrete Mathematics group in Hamburg (including me), and asked whether one can improve the tree-decomposition to be even into finite parts.

I suggested considering critical vertex sets. I found a tree-decomposition displaying the critical vertex sets (Section 3.5), which I suggested to refine with the previous linked, componental, tight tree-decomposition into rayless parts to obtain the final tree-decomposition into finite parts (Section 3.4.3). Theorem 3 and Corollary 5 are due to me and Corollaries 3.1.2 and 3.1.3 are due to Max Pitz. Raphael Jacobs, Paul Knappe and I worked out the details of all proofs in this chapter. The final presentation, which included a complete overhaul of the story, was then developed by Paul Knappe and myself.

Chapter 4: This chapter is based on the paper [11], which is joint work with Raphael Jacobs, Paul Knappe and Max Pitz. I developed Construction 4.2.1 and proved Examples 8 and 9. For Example 10, Max Pitz suggested considering the graph from [26, Example 3.7], and I proved that it indeed works. Raphael Jacobs, Paul Knappe and Max Pitz developed Theorem 4.4.1. Paul Knappe and I worked out the details of the proofs.

Chapter 5: This chapter is based on the paper [5], which I created entirely on my own.

Part II:

Chapter 6: This chapter is based on the paper [3], which I created entirely on my own.

Chapter 7: This chapter is based on the paper [6], which I created entirely on my own.

Chapter 8: This chapter is based on the paper [7], which is joint work with Hanno von Bergen, Raphael Jacobs, Paul Knappe and Paul Wollan. We have in close collaboration developed a rough outline of the proof during a 3-weeks research visit of Paul Wollan in Hamburg in October 2021. Hanno von Bergen and I developed a more detailed outline of the proof. Afterwards, Raphael Jacobs, Paul Knappe and I worked out the full details of the proofs and developed the final presentation.

Part III:

Chapter 10: This chapter is based on the paper [10], which is joint work with Tony Huynh, Raphael Jacobs, Paul Knappe and Paul Wollan. This project was started during a 3-weeks research visit of Paul Wollan in Hamburg in June and July 2022. Raphael Jacobs, Paul Knappe, Paul Wollan and I have in close collaboration developed Conjecture 28 and tried various proof ideas for Theorem 27, but were unsuccessful in our attempts. Tony Huynh and Paul Wollan then proved Theorem 27 building on some of our ideas. I spotted and repaired a major error in the construction of the two paths P_1, P_2 starting on page 197. Paul Knappe, Raphael Jacobs and I worked out the details of this fix and rewrote the construction of the two paths P_1, P_2 starting on page 197.

Chapter 11: This chapter is based on the paper [8], which is joint work with Reinhard Diestel, Ann-Kathrin Elm, Eva Fluck, Raphael Jacobs, Paul Knappe and Paul Wollan. During a 2-weeks research visit of Paul Wollan and Eva Fluck in Hamburg in April 2022, we have developed in close collaboration a detailed outline of a proof of a version of Theorems 30 and 31, where we excluded 3-geodesic subdivision of K_3. Ann-Kathrin Elm came up with and proved Lemma 11.4.4, which improved our earlier version of Theorem 30 to the one presented in this thesis. I proved Theorem 32. Ann-Kathrin Elm, Eva Fluck, Raphael Jacobs, Paul Knappe and I worked out the details of the proofs and finalised a first version ([8]). The final presentation of the results in Sections 11.1–11.3 in this thesis, which takes into account several recent developments in coarse graph theory, is due to Raphael Jacobs, Paul Knappe and myself.

Chapter 12: This chapter is based on the paper [13], which is joint work with Raphael Jacobs, Paul Knappe and Paul Wollan. During a 3-weeks research visit of Paul Wollan in Hamburg in June and July 2022, we have in close collaboration gathered proof ideas for Theorems 33 and 34, but were unsuccessful in our attempts. I then proved Theorems 33 and 34 building on some of these ideas; in particular, the proof of Lemma 12.4.2 is completely my own. Paul Knappe has formalised the inductive proof method described in Section 12.2. Paul Knappe and I worked out the details of the proofs of Theorems 33 and 34 and developed the final presentation.

Chapter 13: This chapter is based on the paper [9], which is joint work with Matthias Hamann. We have in close collaboration developed a detailed outline for Theorem 37 and Section 13.4. I then proved Theorem 35 and modified the proofs of Theorems 35 and 37 to also yield Theorems 36 and 38. Matthias Hamann and I polished the proofs in this chapter and developed the final presentation.

Acknowledgements

I would like to thank everyone without whom this thesis would not have been possible. First, I would like to thank my supervisor, Reinhard Diestel, for giving me the opportunity to do a PhD, and for his advice throughout my studies.

Next, I would like to thank my co-authors, especially Paul and Raphael, for our enjoyable and exciting work together. Together, we climbed the highest *trees*, reaching up to the *clouds*, the *sun*, the *stars*, to the *infinities* and far beyond. We walked any *distance* to discover new *paths* leading to far-away *ends*, we came across some prickly *cacti* and saw some colourful *rainbows*. Research would not have been half as much fun without you!

I would also like to thank my office mates Paul and Raphael (again); it was a lot of fun sharing an office with you! :D Moreover, I would like to thank everyone in the DM group at Hamburg for creating a pleasant working environment and for inspiring and fun conversations. Our game nights were always a lot of fun; I enjoyed every single game of 'Tempel des Schreckens'! :D

I would also like to thank Anusch Taraz for introducing me to the joys of graph theory. Moreover, I would like to thank Helge Wilms and Jens Wohlenberg, two of my former high school teachers, and the William-Stern-Gesellschaft for introducing me to the joys of mathematics.

Last but not least, I would like to thank my mum and dad for their unwavering love and support, not only during the four years of my PhD studies, but throughout my entire life!

An honourable mention goes to the person who gave a short presentation about the Technomathematik degree programme at the Matheolympiade awards ceremony at the TUHH around nine years ago. Without them, I would probably have studied engineering instead.

So long, and thanks for all the fun!

Eidesstattliche Erklärung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Dissertation selbst verfasst und keine anderen als die angegebenen Hilfsmittel benutzt habe. Darüber hinaus versichere ich, dass diese Dissertation nicht in einem früheren Promotionsverfahren eingereicht wurde.

Elmshorn, 05.06.2025

Sandra Albrechtsen