GRAPH-DECOMPOSITIONS: TREES, TANGLES AND LOCALITY

DISSERTATION

zur Erlangung des

Doktorgrades der Naturwissenschaften

vorgelegt von

PAUL KNAPPE

im Fachbereich Mathematik

an der Fakultät

für Mathematik, Informatik und Naturwissenschaften der

UNIVERSITÄT HAMBURG

Hamburg, 2025

Betreuer
Reinhard Diestel, Universität Hamburg

Gutachter
Reinhard Diestel, Universität Hamburg
Louis Esperet, Laboratoire G-SCOP in Grenoble
Imre Leader, University of Cambridge

Prüfungskomission
Armin Iske, Universität Hamburg
Reinhard Diestel, Universität Hamburg
Nathan Bowler, Universität Hamburg
Ralf Holtkamp, Universität Hamburg
Matthias Hamann, Universität Hamburg

Datum der Disputation: 17. September 2025

Contents

Introduction **1**

I Local versus global aspects of graphs **11**

1 Graph decompositions via coverings **13**
1.1 Introduction . . . 13
1.2 Terminology . . . 20
1.3 Graph-decompositions . . . 22
1.4 Local coverings . . . 33
1.5 Displaying global structure . . . 43
1.6 Local covers are accessible . . . 49
1.7 Tools and lemmas for the proof of Theorem 1 . . . 49
1.8 Proof of the main result of this chapter . . . 57
1.9 Tree-Decompositions in quasi-transitive graphs . . . 60

2 Graph decompositions via local separations **63**
2.1 Introduction . . . 63
2.2 Review of graph-decompositions and coverings . . . 70
2.3 Bottlenecks . . . 75
2.4 Local separations . . . 82
2.5 High displacement: a guarantee for the existence of lifts . . . 89
2.6 Links and crossing . . . 95
2.7 Local bottlenecks . . . 101
2.8 Constructing graph-decompositions . . . 108
2.9 Proof of Theorems 4 and 5 . . . 118

3 Locally chordal graphs **121**
3.1 Introduction . . . 121
3.2 Basic characterizations of locally chordal graphs . . . 128
3.3 Graph-decompositions & region representations . . . 133
3.4 Region representations of (locally) chordal graphs . . . 137
3.5 Region representations of wheel-free graphs . . . 143
3.6 Characterizations via 3-acyclic region representations . . . 146
3.7 Graph-decompositions of locally chordal graphs . . . 148
3.8 Characterization via graph-decompositions . . . 154
3.9 Characterization via minimal local separators . . . 157

3.10 Computing clique graphs . . . 160
3.11 Application: Locally interval & circular arc graphs . . . 165
3.12 Future work . . . 170
3.13 Canonical subtree representations of chordal graphs . . . 172
3.14 Comparison: Canonical graph-decomposition via coverings . . . 177

II Coarse graph theory 179

4 Path-decompositions into small-radius parts 181
4.1 Introduction . . . 181
4.2 Preliminaries . . . 186
4.3 Graph-decompositions and quasi-isometries . . . 187
4.4 Radial path-width . . . 194
4.5 Radial cycle-width . . . 201
4.6 Radial star-width . . . 202
4.7 Counterexample . . . 210
4.8 Counterexamples to versions in the star and tree case . . . 212

5 Graphs quasi-isometric to K_4-minor-free graphs 217
5.1 Introduction . . . 217
5.2 Preliminaries . . . 220
5.3 Preparatory work: Sequences of graph-decompositions . . . 224
5.4 Proof of Theorem 26: Conjecture 24 for $X = K_4^-$ . . . 227
5.5 Proof of Theorem 25: Conjecture 24 for $X = K_4$ . . . 232

6 A Menger-type theorem for two induced paths 253
6.1 Introduction . . . 253
6.2 Interlaced systems of intervals . . . 255
6.3 Proof of Theorem 6.1.1 . . . 257

III Tree-like structures and tangles 267

7 Linked tree-decompositions into finite parts 269
7.1 Introduction . . . 269
7.2 Preliminaries . . . 275
7.3 Displaying ends . . . 281
7.4 A high-level proof of the main result . . . 283
7.5 Tree-decomposition along critical vertex sets . . . 289
7.6 Lifting paths and rays from torso . . . 292
7.7 Linked tree-decompositions into rayless parts . . . 295
7.8 Lean tree-decompositions . . . 309

7.9 Tree-decompositions distinguishing infinite tangles . . . 316
7.10 A closer look at Theorem Theorem 33' . . . 317

8 Counterexamples: linked tree-decompositions 321
8.1 Introduction . . . 321
8.2 Preliminaries . . . 325
8.3 Examples 35 and 36 . . . 326
8.4 Example 37 . . . 332
8.5 Upwards disjointness of adhesion sets . . . 333

9 Efficiently distinguishing all tangles 337
9.1 Introduction . . . 337
9.2 Preliminaries . . . 339
9.3 Limit separations . . . 340
9.4 The interlacing theorem . . . 346
9.5 Proofs of Theorems 38 and 39 . . . 352
9.6 Infinitely many rays beyond the limit . . . 356

10 On vertex sets inducing tangles 361
10.1 Introduction . . . 361
10.2 Preliminaries . . . 366
10.3 Definition of 'survive' and special cases of Theorem 43 . . . 369
10.4 Proof of Theorem 44 if G has a higher-order tangle . . . 372
10.5 Rainbow-cloud-decompositions . . . 377
10.6 RC-decompositions and separations . . . 385
10.7 Proof of Theorem 44 if G has no higher-order tangle . . . 389
10.8 The inductive proof method and its applications . . . 399

11 Grid theorem for strong immersions of walls 405
11.1 Introduction . . . 405
11.2 Almost α-thin graphs and 3-centres . . . 408
11.3 Strongly immersed walls are obstructions . . . 411
11.4 Excluding strongly immersed walls . . . 414

12 Lovász-Cherkassky theorem 423
12.1 Introduction . . . 423
12.2 Proof of the main result of this chapter . . . 425
12.3 A degree condition for vertices and ends . . . 428

Bibliography **431**

Appendices **445**

A Summary of the results **447**

B Zusammenfassung der Ergebnisse **451**

C Publications related to this dissertation **455**

D Declaration of my contributions **457**

E Danksagungen **461**

F Eidesstattliche Versicherung **463**

Introduction

In this thesis, we make the global structure of a graph visible, study the interplay of graph-theoretic concepts with the large-scale geometry of the underlying graph, and display the infinities, i.e. ends and critical vertex sets, of infinite graphs. The tools we use and the problems we study are inspired by Robertson and Seymour's 20-research-paper long series on graph minors [RS] and other related research. Their series culminated in the proof of one of the deepest theorems in graph theory:

Graph minor theorem. *For every infinite sequence $G_1, G_2, G_3, \ldots$ of finite graphs, there are two indices $i < j$ such that G_i is a minor of G_j.*

In their proof, they developed a connectivity theory that is specifically tailored to the central object of their study, *graph minors*. A new concept in their theory is that of a *tangle*, which captures high connectivity somewhere in a graph G without conforming to the usual graph-theoretic notions of high connectivity, such as highly connected subgraphs. This concept is based on the observation that, for every low-order separation of G, any highly connected minor in G lives mostly on one of the sides of the separation. By describing the orientation of the low-order separations induced in this way by a highly connected minor, Robertson and Seymour's tangles are an indirect way to capture just the essence of highly connected minors and other highly connected substructures of G without explicitly describing their structure.

Their tangle-tree theorem states that the tangles of a (finite) graph G are always arranged in a tree-like fashion [RS91]. They formalised the tangle-tree structure using the notion of *tree-decompositions*, that is a decomposition $\mathcal{V} = (V_t \mid t \in V(T))$, a family of vertex-subsets V_t of G, modelled on a tree T such that the induced subgraphs $G[V_t]$ together cover the entire graph G and, for each vertex v of G, the nodes $t \in T$ whose corresponding V_t contains v induce a connected subgraph of the tree T. The V_t are called the *bags* of the decomposition.

The first part of this thesis deals with the most fundamental graph-theoretical problem in this thesis, and develops a novel approach to tackle it. Few graphs are totally homogeneous at all levels of local or global focus: most have some clusters that are particularly highly connected. These clusters then form a global structure between them. Robertson and Seymour's tangle-tree structure displays the global structure of a graph when this is tree-like. However, not all graphs are globally tree-like. For those that are not, the tangle-tree structure is a poor fit to display their global structure. We make the global structure of any given finite graph G visible in the form of a decomposition modelled on a graph H that need not be a tree; see Figure 1 for an example of this visualisation of the global structure.

In the second part, we solve some previously unkown cases of Georgakopoulos and Papasoglu's *fat minor conjecture* in the recently emerged field of *coarse graph theory*. A solution to this conjecture for a given minor-closed graph class $\mathcal{G}$ yields a description of the large-scale geometry of any given graph $G \in \mathcal{G}$ by another simpler graph H.

In the third part, we prove that, if an infinite graph admits a tree-decomposition into finite parts, then there is such a tree-decomposition which displays its infinities, i.e. ends and critical vertex sets. Additionally, we consider several other problems regarding tree-like structures and also obstructions to their existence in finite and also infinite graphs.

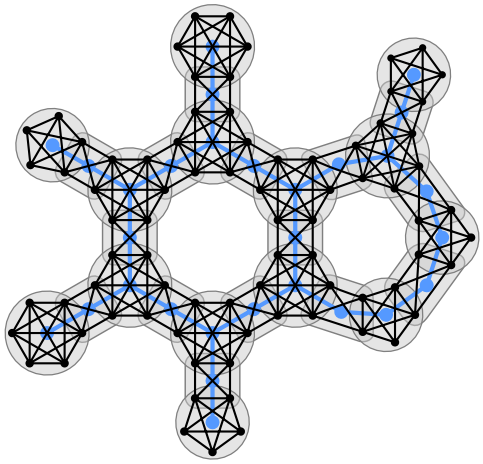

Figure 1: A black graph with non-tree-like global structure that is displayed by our canonical decomposition into the grey bags modelled on the blue graph.

The remainder of this chapter contains a brief overview of the parts and also their respective chapters. For more details, especially regarding terminology, we refer the reader to the chapters themselves, as each chapter has its own independent and more detailed introduction.

Part I: Local versus global aspects of graphs

In this part, we assume that we are given any integer $r \geqslant 0$, the *locality parameter*, which describes the local focus: it tells apart local from global structure, in that we consider the structure in the balls of radius $r/2$ as local and everything else as global.

Chapter 1: Canonical graph decompositions via coverings

In this chapter, we make the global structure of graphs visible. Robertson and Seymour's tangle-tree theorem did this for graphs which are globally tree-like, but we do it even when the global structure is not tree-like. To allow this, we generalise tree-decompositions to graph-decompositions by allowing not only trees T but any graph H as model for the decomposition. Then we canonically, i.e. with reference only to the locality parameter r, construct a decomposition $\mathcal{V} = (V_h \mid h \in V(H))$ of any given finite graph G modelled on a graph H that displays the global structure of G. In particular, the graph H is solely determined by the graph G being decomposed and the locality parameter r: the global structure of G is *found*, not imposed to be in a specific class of graphs, such as the class of trees in the case of the tangle-tree theorem. We depicted our canonical graph-decomposition of an example for $r = 4$ in Figure 1.

To construct our desired graph-decomposition, we first develop a new framework to discriminate between the local structure of a graph and its global structure. We encode this by relying on coverings from topology: we consider a (possibly infinite) graph G_r, the *r-local cover* of G, that reflects only the local structure of G, that is the $r/2$-balls, and thus unfolds its remaining (global) structure (see Section 1.4 for the formal definition). Due to the nature of coverings,

the global structure of G is unfolded in a tree-like way in G_r. Therefore, the global structure of G_r is tree-like and is captured by its tangle-tree structure. Our key technical result, then, is an extension of Robertson and Seymour's tangle-tree theorem to locally finite, quasi-transitive graphs, which include in particular the cover G_r we consider. Thus, the tangle-tree structure of G_r can again be formalised by a tree-decomposition $(T, \hat{\mathcal{V}})$ (Lemma 1.8.5), which, in addition, is canonical. In particular, the group Γ_r of all deck transformations acts naturally on this tree-decomposition.

Now the decomposition $\mathcal{V} = (V_h \mid h \in V(H))$ and its model H are obtained by folding the tangle-tree structure of G_r onto G, in that H is the orbit graph T/Γ_r and, for each node $h \in H$, the bag V_h is the projection of the bag $\hat{V}_t$ corresponding to any node t of T in the Γ_r-orbit h. We denote this graph-decomposition $(H, \mathcal{V})$ by $\mathscr{D}_r(G)$. In this way, the tangle theory from graph minors is brought to bear canonically on arbitrary graphs, which need not be globally tree-like.

Chapter 2: Canonical graph decompositions via local separations

The r-local cover G_r is often infinite. Hence, following the above construction of the canonical decomposition $\mathscr{D}_r(G)$ of G, it is unclear whether $\mathscr{D}_r(G)$ is computable. In this chapter, we reconstruct the canonical decomposition $\mathscr{D}_r(G)$ from local information in G itself. Let us recall, e.g. from [CK24; DHL19; EKT22], that the tangle-tree structure of a graph may be defined entirely in terms of its low-order separations. In order to reconstruct the canonical decomposition of G, we develop a new theory of *local* separations that we want to satisfy the following properties:

(i) the local separations of G_r correspond to the usual (global) separations of G_r, and

(ii) the local separations of G lift to the usual (global) separations of G_r, and

(iii) vice versa, the usual (global) separations of G_r project to the local separations of G.

However, while our theory of local separations always satisfies (i), unfortunately it need not in general satisfy (ii) and (iii). Yet, it is guaranteed to satisfy (ii) and (iii) whenever some computable graph invariant of G, the *r-displacement* of G (see Section 2.5 for the definition), is sufficiently large depending on r and the order of the (local) separations considered. In turn, this will also guarantee the computability of $\mathscr{D}_r(G)$ for those graphs G.

Chapter 3: Locally chordal graphs

In this chapter, we investigate r-locally chordal graphs. A graph is *r-locally chordal* if all its balls of radius $r/2$ are chordal. See Figure 1 for an example. First, we characterise them as graphs whose r-local cover is chordal. Second, we characterise them as graphs that contain no induced subgraphs that are wheels or cycles of lengths between 4 and r. In this context, one of the greatest hopes is to develop a common global description of these local constraints. We achieve this by proving a structural characterisation of locally chordal graphs in terms of

graph-decompositions into cliques. This extends the well-known structural characterisation of chordal graphs as graphs admitting tree-decompositions into cliques. For this, we heavily rely on the general theory introduced in the previous two chapters that is designed to facilitate local-to-global analysis. Moreover, we will show that the r-locally chordal graphs are precisely those graphs G whose canonical graph-decomposition $\mathscr{D}_r(G)$, which displays the global structure, is into cliques.

Part II: Coarse graph theory

In this part, we present a width notion such that a graph G admits an H-decomposition of small 'width' if and only if H resembles the large-scale structure of G. As a first application, this yields a new approach to solve Georgakopoulos and Papasoglu's *fat minor conjecture* in the currently emerging field of *coarse graph theory*. Coarse graph theory is of particular interest as it provides a benchmark for the type of theorems we are looking for: they must be invariant under quasi-isometries with any given constants. Using our tools, we solve some previously unknown cases of the fat minor conjecture.

Following the naive approach to define a 'width' of a graph-decomposition, one may define it analogously to the width of a tree-decomposition, i.e. as the maximum size of a bag of the decomposition (minus 1). However, given a class $\mathcal{H}$ of finite graphs, the respective '$\mathcal{H}$-width' of a graph G, the minimum width among all $\mathcal{H}$-decompositions of G, i.e. H-decompositions with $H \in \mathcal{H}$, does not yield a meaningful extension of tree-width: if $\mathcal{H}$ has bounded tree-width, then every graph of small $\mathcal{H}$-width has small tree-width too, and if $\mathcal{H}$ has unbounded tree-width, then the $\mathcal{H}$-width of every graph is at most 2 (minus 1) [DK05].

This inspired us to take a completely different approach. To define our desired 'width', we adopted a metric perspective. Rather than measuring a bag of the decomposition by its cardinality, we consider the radius of its induced subgraph: the *(inner-)radial width* of a decomposition $(H, \mathcal{V})$ is the smallest radius of the induced subgraphs $G[V_h]$ on the bags of the decomposition. The *(inner-)radial $\mathcal{H}$-width* of a graph G is the smallest radial width among the $\mathcal{H}$-decompositions of G.

In this context, one of the first questions of interest is: what are the obstruction to small radial $\mathcal{H}$-width for a given class $\mathcal{H}$ of finite graphs? A moment's thought shows that it entails no loss of generality to restrict to graph classes $\mathcal{H}$ closed under minors. Then, by Robertson and Seymour's graph minor theorem, there is a finite set $\mathcal{X}$ of finite graphs such that $\mathcal{H}$ consists precisely of the graphs with no X minor for any $X \in \mathcal{X}$. 'Coarse' X minors for any $X \in \mathcal{X}$ in G force G to have large radial $\mathcal{H}$-width (Lemma 4.3.15). A coarse, so called *K-fat* for some $K \in \mathbb{N}$, X minor of any $X \in \mathcal{X}$ in G is a minor whose branch sets and branch paths are far, i.e. $> K$, apart in G whenever it is possible, that is except for incident pairs of vertices and edges (see Section 5.2.2 for a definition). We conjectured that fat X-minors with $X \in \mathcal{X}$ are the only obstructions:

Conjecture. *Let $\mathcal{X}$ be a finite set of finite graphs. Then there exists a function $f\colon \mathbb{N} \to \mathbb{N}$ such that, for all $K \in \mathbb{N}$, every graph G with no K-fat X minor for any $X \in \mathcal{X}$ admits a decomposition of radial width at most $f(K)$ modelled on a graph H with no X minor for any $X \in \mathcal{X}$.*

Independent from our research, Georgakopoulos and Papasoglu [GP23] followed Gromov's ideas on coarse geometry [Gro93] into the realm of graphs: they suggested a study of graphs from a coarse or metric perspective, which revolves around the concept of quasi-isometry, which is a generalisation of bi-Lipschitz maps that additionally allows for an additive error. Roughly speaking, two metric spaces are quasi-isometric if their large-scale geometry coincides. As it turns out, whether G admits an H-decomposition of small radial width is closely related to whether G has a quasi-isometry to H. In fact, both are the same question if we ask the graph-decomposition to not only have small radial width but additionally be 'honest' and have small 'radial spread' (see Section 5.2.5 for more details). Georgakopoulos and Papasoglu [GP23, Conjecture 1.1] conjectured a statement similar to our above-mentioned conjecture but in the setting of quasi-isometries:

Fat minor conjecture. *Let $\mathcal{X}$ be a finite set of finite graphs. Then there exists a function $f\colon \mathbb{N} \to \mathbb{N}^2$ such that, for all $K \in \mathbb{N}$, every graph G 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 $X \in \mathcal{X}$.*

Recently, the fat minor conjecture has been disproven by Davies, Hickingbotham, Illingworth and McCarty for general sets $\mathcal{X}$ [JRFR24, Theorem 1]. However, the fat minor conjecture is known to be true in some small cases (see [CDNRV12; FP23; GP23] and Chapters 4 and 5).

The known counterexamples to the fat minor conjecture are also counterexamples to our conjecture. We ask the question of whether Georgakopoulos and Papasoglu's fat minor conjecture and our conjecture are basically equivalent, in that quasi-isometries and decomposition of small radial width are the same for the relevant graph classes (Question 4.3.11).

Chapter 4: A structural duality for path-decompositions into parts of small radius

In this chapter, we prove that the fat minor conjecture holds for $\mathcal{X} = \{K_3, K_{1,3}\}$ (which describes 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 depicted in Figure 4.1. Moreover, whenever we cannot find a small radial-width decomposition of the desired form, we instead find a graph $X \in \mathcal{X}$ not just as a fat minor but as a quasi-geodesic topological minor in G. These are topological X minors in G, whose minor-models are quasi-geodesic in G in that their metrics agrees with the one induced by G up to some multiplicative error (see Section 4.3.4).

Chapter 5: A characterisation of graphs quasi-isometric to K_4-minor-free graph

In this chapter, we prove that the fat minor conjecture holds for $\mathcal{X} = \{K_4\}$ by proving its corresponding version for graph-decompositions. The advantage of this approach is that, given a graph G with no K-fat K_4 minor, we obtain the desired graph-decomposition as the limit of a (possibly infinite if G is infinite) sequence of suitable increasing graph-decompositions of larger and larger subgraphs of G which ultimately exhaust G (Section 5.3). Our proof technique also yields a new short proof of the fat minor conjecture for $\mathcal{X} = \{K_4^-\}$, which was first established by Fujiwara and Papasoglu [FP23].

Chapter 6: A Menger-type theorem for two induced paths

Any proof of the fat minor conjecture for a given class $\mathcal{X}$ of graphs must at some point find fat minors of the graphs $X \in \mathcal{X}$. A tool which is often very helpful to construct such objects is Menger's theorem. Thus, we conjectured a Menger-type theorem that fits in the setting of coarse graph theory:

Coarse Menger conjecture. *Let A, B be two sets of vertices in a graph G. Then there exists a function $f\colon \mathbb{N}^2 \to \mathbb{N}$ such that, for every $k, d \in \mathbb{N}$, there are either k many A–B paths in G that are at least d apart or there is a set S of at most $k-1$ vertices of G such that each A–B paths meets some vertex which is at most $f(k,d)$ close to some $s \in S$.*

In this chapter, we prove the coarse Menger conjecture for $k = 2$, which was sufficient to our proof of case $\mathcal{X} = \{K_4\}$ of the fat minor conjecture. The coarse Menger conjecture immediately generated a lot of attention [GKL23; GHW25; HNST23; NSS24]. Recently, Nguyen, Scott and Seymour [NSS24] disproved the (exact) coarse Menger conjecture for $k, d \geqslant 3$. The more general qualitative coarse Menger conjecture which asks just for a set of vertices whose size is bounded in k is still open.

Part III: Tree-like structures and tangles

In this part, we discuss various problems regarding tree-like structures and also obstructions to their existence in finite and infinite graphs. In Chapters 7 to 9, we aim to display the infinities, i.e. ends and *critical vertex sets*, as well as possible by relying on the well-known concepts of *linked* and *lean* tree-decompositions and those which *efficiently distinguish all the tangles*. Chapter 10 deals with the structure of tangles, the main result in Chapter 11 describes the structure of graphs with no immersions of walls, and in Chapter 12 we extend the Lovász-Cherkassky theorem to locally finite graphs and their ends.

Chapter 7: Linked tree-decompositions into finite parts

Kříž and Thomas's result on the existence of linked tree-decompositions forms a cornerstone both in Robertson and Seymour's work [RS90a] on well-quasi-ordering finite graphs, and in

Thomas's result [Tho89] that the class of infinite graphs of tree-width $< k$ is well-quasi-ordered under the minor relation, for all $k \in \mathbb{N}$. It reads as follows:

Theorem (Thomas 1990 [Tho90], Kříž and Thomas 1991 [KT91])**.** *For any integer $k \geqslant 0$, any (finite or infinite) graph of tree-width $< k$ has a linked rooted tree-decomposition*[1] *of width $< k$.*

The main result of this chapter is that their result also holds for $k = \aleph_0$. This result required the deepest insights in this part of the thesis. In contrast to their result, the linked rooted tree-decomposition in our result is additionally tight and componental. Moreover, it displays all the infinities, i.e. ends and *critical vertex sets*, of the underlying graph.

Our result has several applications: we extend the stronger version of Kříž and Thomas's result on the existence of lean tree-decompositions to graphs without half-grid minor, we resolve a question of Halin from 1977, and we obtain short, unified proofs of the characterisations by Robertson, Seymour and Thomas of graphs without half-grid minor, and of graphs without binary tree subdivision.

Chapter 8: Counterexamples regarding linked and lean tree-decompositions of infinite graphs

In this chapter, we present examples which suggest that our main result from the previous Chapter 7 is at the frontier of what is could still be true. Furthermore, we construct several examples to demonstrate the limitations of generalising Kříž and Thomas's result to arbitrary infinite tree-width graphs. To achieve this, we in particular construct a locally finite, planar, connected graph that has no lean tree-decomposition at all.

Chapter 9: Efficiently distinguishing all tangles in locally finite graphs

One of the highlights of Robertson and Seymour's graph minors series is the tangle-tree theorem: any finite graph admits a tree-decomposition that efficiently distinguishes its tangles. While it was known that the theorem could not be extended verbatim to infinite, even locally finite, graphs [EKT22, Example 4.9], the obstructions were not understood. While previous work, e.g. [EKT22, §6.1], achieved extensions through considerable relaxation of the conclusion, in this chapter we identify the obstructions in the case of locally finite graphs directly and prove a full-strength extension when these are absent. We show that there is only one type of possible obstruction in locally finite graphs: a thick end (equivalently, an infinite half-grid minor). The (canonical) tangle-tree theorem therefore holds for locally finite graphs without thick ends.

Chapter 10: On vertex sets inducing tangles

Tangles are designed to capture highly connected substructures of a graph: orienting every low-order separation $\{A, B\}$ towards its side that contains most of the substructure forms a

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

tangle. Diestel, Hundertmark and Lemanczyk [DHL19] asked whether every tangle is induced by 'majority vote' of some set of vertices, as follows: given a tangle τ in a graph G, is there a set $U \subseteq V(G)$ such that $|U \cap A| < |U \cap B|$ for every $(A, B) \in \tau$? It seems difficult to answer their question; only some special cases [Elb15] and the fractional version [EKT20] are known.

In this chapter, we reduce this question for tangles of order k to graphs whose size is bounded by a function in k. The key step of our proof, which may be of independent interest, is to 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.

Let us recall Elbracht, Kneip and Teegen's affirmative answer to the fractional version [EKT20] of Diestel, Hundertmark and Lemanczyk's question: every tangle in any graph G is induced by a weight function $w\colon V(G) \to \mathbb{N}$, i.e. $\sum_{a\in A} w(a) < \sum_{b\in B} w(b)$ for every $(A, B) \in \tau$. Together with our above described result, we strengthen their result: every tangle is induced by some weight function w whose total weight $\sum_{v\in V(G)} w(v)$ is bounded in the order of the tangle.

Chapter 11: A grid theorem for strong immersions of walls

Like minors, immersion minors generalise the notion of topological minors, but they are unrelated to minors in that neither the existence of an immersion minor implies the existence of a minor nor vice-versa. Immersion minors have attracted attention from various algorithmic and structural viewpoints in recent years, often finding results analogous to those for minors [AL03; DDFMMS14; DK14b; DW16; FGTW08; GKMW11; KK24; RS10; Wol15].

In this chapter, we are concerned with the structure of graphs not containing a large wall as a strong immersion minor. Dvořák and Wollan's structure theorem for strong immersions [DW16, Theorem 4] has in the case of walls no (qualitative) converse. This motivated us to prove a structure theorem for graphs with no strong immersions of walls in this chapter: a graph contains a large wall as a strong immersion minor if and only if the graph has a tree-cut decomposition of whose torsos admit a certain path-like structure.

Chapter 12: The Lovász-Cherkassky theorem for locally finite graphs with ends

Let T be any set of vertices in a finite graph G. In a set $\mathcal{P}$ of edge-disjoint T-paths, there are at most $\lambda_G(t, T \setminus \{t\})$ many paths that link a vertex $t \in T$ to $T \setminus \{t\}$, where $\lambda_G(t, T \setminus \{t\})$ denotes the minimum size of a cut separating t and $T \setminus \{t\}$. It follows that $|\mathcal{P}| \leqslant \frac{1}{2} \sum_{t\in T} \lambda_G(t, T \setminus \{t\})$. This upper bound cannot be achieved in general. Perhaps surprisingly, Lovász [Lov76] and Cherkassky [Che77] independently showed that it can be achieved under the additional assumption that the graph G is *inner-Eulerian* for T in that every vertex of G which is not in T has even degree in G. Diestel asked whether the theorem of Lovász and Cherkassky may be extended to all locally-finite infinite graphs and their ends [Die20]. In this chapter, we answer his question in the affirmative.

PART I

Local versus global aspects of graphs

CHAPTER 1

Canonical graph decompositions via coverings

This chapter is based on [DJKK22], which is joint work with Reinhard Diestel, Raphael W. Jacobs and Jan Kurkofka.

1.1 Introduction

The question to what extent graph invariants—the chromatic number, say, or connectivity—are of local or global character, and how their local and global aspects interact, drives much of the research in graph theory both structural and extremal [Sco22]. In this chapter we offer such studies a possible formal basis.

We show that every finite graph G decomposes canonically into local parts which, between them, form a global structure displayed by a simpler graph H. Both H and the decomposition $(H, (G_h)_{h\in H})$ of G into parts indexed by the nodes of H are unique once we have set an integer parameter $r > 0$ to define our desired threshold between 'local' and 'global': cycles in G of length $\leqslant r$ are considered as local, and are reflected inside the parts G_h of the decomposition, while the cycles of H reflect only global cycles of G, longer cycles that are not generated by the short ones.

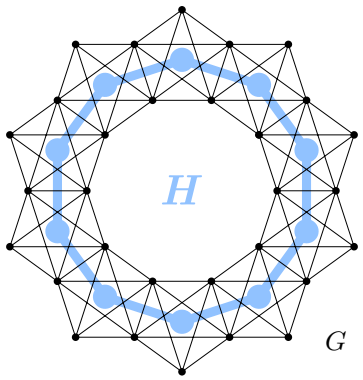

Figure 1.1: The global structure of G is displayed by a cycle H. Its local parts are K^5s.

Our main result reads as follows:

Theorem 1. *Let G be any finite graph, and $r > 0$ an integer. Then G has a unique canonical decomposition modelled on another finite graph $H = H(G, r)$ that displays its r-global structure.*

See Sections 1.1.1–1.1.3 below for a summary of the definitions needed here. More details are given in Definition 1.5.4 and the text following it.

Since $H = H(G, r)$ and the associated decomposition of G are canonical, they are invariants of G, and we can study how they interact with other graph invariants as r ranges between 1

and $|G|$. For example, we may study how the chromatic number, connectivity, or edge density of G is reflected in H, and is in this sense global, or is reflected in the parts G_h of the decomposition of G, and is in that sense local.

Theorem 1 extends to connected Cayley graphs even when these are only locally finite. It thus implies a canonical decomposition theorem for finitely generated groups into local parts, whose relative structure is displayed by a graph:

Theorem 2. *Let Γ be a group given with a finite set S of generators. For every integer $r > 0$, the Cayley graph $\mathrm{Cay}(\Gamma, S)$ has a unique canonical decomposition modelled on a graph $H = H(\Gamma, S, r)$ that displays the r-global structure of the group Γ as presented by S.*

Our results have further group-theoretical implications. We shall see that finitely generated groups Γ extend to finitely presented groups Γ_r that describe their *r-local* structure, the aspect of their structure that complements their r-global structure given by Theorem 2. Similarly, we shall see that the r-local structure of any finite graph G that complements its r-global structure as described in Theorem 1 is displayed by a graph G_r modelled on the Cayley graph of a finitely presented group. This graph G_r is typically infinite. It describes the local structure of G by way of a canonical covering $G_r \to G$. See Section 1.4 for details. Just recently, similar graph-decomposition methods have been used by Carmesin, Kontogeorgiou, Kurkofka and Turner to prove a low-order Stallings-type theorem for finite nilpotent groups [CKKT24]. More group-theoretical applications of this approach have been obtained by Carmesin [Car21].

The main idea for the construction of $H = H(G, r)$ was inspired by the aims of [Car22]. The idea is to apply the theory of tangles from graph minor theory, though not to G itself, but to the covering space G_r of G that reflects its short cycles, those of length $\leqslant r$, while unfolding in a tree-like way any longer cycles not generated by the short ones.[1] This tree-likeness of this graph G_r enables us to use tangles to structure G_r as customary in graph minor theory: by canonical tangle-distinguishing tree-decompositions. Projecting this structure back to G defines H and the decomposition of G modelled on H.

Our original proof of Theorem 1 also used group-theoretic tools in an essential way. In order to apply tangle-distinguishing tree-decompositions to our covering graph G_r, we needed that G_r resembles a Cayley graph of some finitely presented group that we could decompose using Stallings' theorem [Sta72], and then use [TW93] to analyse its ends. We formalised this resemblance by showing that G_r had a decomposition into bounded-sized finite parts modelled on a Cayley graph of its group of deck transformations over G. This group is finitely presented not only when G is finite (Theorem 1), but also when it is the Cayley graph of an arbitrary finitely generated group (Theorem 2). We have since been able to replace these group-theoretical tools by purely combinatorial arguments, which is how we present our proof here.

In the remainder of this introduction we indicate just enough background to be able to make the statement of Theorem 1 precise, and to give a brief sketch of its proof.

[1] The use of coverings to analyse graph connectivity from a local/global perspective goes back to Amit and Linial [AL02].

1.1.1 Tree-decompositions vs. graph-decompositions

Early in their seminal work on graph minors, Robertson and Seymour [RS] introduced *tree-decompositions* of graphs as a way to display their global structure when this is tree-like. However, not all graphs are globally tree-like. For those that are not, tree-decompositions are a poor fit to display their global structure.

Yet, conversely, few graphs are totally homogeneous at all levels of local or global focus: most have some clusters that are particularly highly connected. These clusters then form a global structure between them, even if this is not tree-like. Theorem 1 makes this global structure visible in the form of a graph H that need not be a tree. Indeed, H will change with the graph G being decomposed and with the locality parameter r we choose: the global structure of G is *found* by our theorem, not imposed as in the case of tree-decompositions.

Let us make this precise. Let H be a graph. A *decomposition* of a graph G *modelled on* H, or an *H-decomposition* of G, is a pair $(H, \mathcal{V})$ consisting of H and a family $\mathcal{V} = (V_h)_{h \in H}$ of sets of vertices of G that are associated with the nodes h of H in such a way that

- $\bigcup_{h \in H} G[V_h] = G$;
- for every vertex $v \in G$, the set W_v of nodes h of H whose V_h contains v is connected in H.

The decompositions of a graph that are modelled on trees are thus precisely its tree-decompositions [Die25].

Given a decomposition as above, we may choose spanning subgraphs $G_h \subseteq G[V_h]$ and $H_v \subseteq H[W_v]$ so that the G_h still cover all the edges of G and the H_v are still connected. We then call these G_h the *parts* of our decomposition, and the H_v its *co-parts*; see Definition 1.3.2 for details.

Our decompositions and choice of their parts will be *canonical* in that they commute with graph isomorphisms: if $(G_h)_{h \in H}$ and $(G'_{h'})_{h' \in H'}$ are the families of parts of decompositions of graphs G and G' that witnessTheorem 1, then any graph isomorphism $\sigma \colon G \to G'$ maps the parts G_h of G to the parts $G'_{h'}$ of G' in such a way that $h \mapsto h'$ is a graph isomorphism $H \to H'$. In particular, the graph H on which Theorem 1 models a given graph G is unique, up to isomorphism, for every choice of r.

1.1.2 Construction of $H = H(G, r)$

The construction of $H = H(G, r)$ is not hard to describe. It works as follows.

Given G and r, let $\pi_1^r(G, v_0)$ denote the subgroup of the fundamental group $\pi_1(G, v_0)$ of G, based at a vertex v_0, that is generated by the elements represented by a walk in G from v_0 to a cycle of length at most r, round it, and back along the access path. This $\pi_1^r(G, v_0)$ is a normal subgroup of the fundamental group of G, independent of the choice of basepoint; let $p_r \colon G_r \to G$ denote the normal covering of G with $\pi_1^r(G, v_0)$ as characteristic subgroup.

By our choice of r, the cover G_r reflects all the short cycles of G, as well as those of its longer cycles that are generated in $\pi_1(G)$ by the short ones. The other longer cycles of G are

usually unfolded to 2-way infinite paths, or *double rays*. (One can construct examples where G_r covers G with finitely many sheets, but those are rare.) Thus, G_r mirrors G in its 'r-local' aspects, but not in its 'r-global' aspects, where it is simply tree-like. This is why we call G_r the *r-local* cover of G.

As a consequence of this tree-likeness of the global aspects of G_r, tree-decompositions will be a better fit for G_r than they were for G, whose global aspects could include long cycles that would not fittingly be captured by tree-decompositions. We exploit this by applying to G_r, not to G, the tree-decompositions that represent the state of the art from the theory of graph minors: those that distinguish all the maximal blocks, tangles and ends of G efficiently, and are canonical in the sense described earlier [DHL19; EKT22]. We fix one such canonical tree-decomposition $\mathcal{T}_r = (T_r, (V_t)_{t\in T_r})$ that is particularly natural, and use it to define the r-global structure of G. More on this in Section 1.1.3.

Since our tree-decomposition $\mathcal{T}_r$ of G_r is canonical, the group $\mathscr{D}_r$ of deck-transformations of G_r over G acts on its model, the tree T_r. The canonical model H for our r-local decomposition $\mathcal{G}_r$ of G is then obtained as the orbit space $T_r/\mathscr{D}_r$, which is a graph. The parts G_h of $\mathcal{G}_r$ are the projections to G of the parts $G[V_t]$ of $\mathcal{T}_r$ whenever h is the $\mathscr{D}_r$-orbit of t.

1.1.3 Details of Theorem 1

In order to make Theorem 1 precise we have to define what it means for a graph H to 'display the r-global structure' of a given graph G. We shall first give a more formal definition of the 'r-global structure' of G, and then define what it means that H displays it. Let $r > 1$ be given.

Consider the r-local cover G_r of G from Section 1.1.2. This is unique up to isomorphisms (of coverings of G), and it is clearly a graph. The r-global structure of G is formally defined as a particular canonical *tree of tangles* of G_r along with the group $\mathscr{D}_r$ of deck transformations of G_r over G. Such a 'tree of tangles' is not a graph. It is a central notion in the theory of graph minors; let us indicate briefly what it means.

Our covering graph G_r may have highly cohesive regions, or 'clusters'. These will include lifts of the local clusters in G that our H-decomposition of G is meant to distinguish, in that they should live in distinct parts G_h. There are various notions of such clusters in graph theory. The most general of these are known as *blocks* [CDHS14; CG17] and as *tangles* [RS91]. When G_r is infinite, as it usually will be, it will also have *ends*. Ends can be formalised as infinite tangles, and all blocks except some irrelevant small ones also induce tangles. We shall therefore use the term 'tangle' from now on to cover all three of these: blocks, ends, and traditional tangles in the sense of Robertson and Seymour [RS]. Details will be given in Section 1.5.1.

All the tangles of our r-local cover G_r of G can be 'distinguished' by a nested set N of separations of G_r, in the following sense. Given any separation of G_r, of order k say, every tangle 'of order $> k$' in G_r will lie essentially on one of the two sides of this separation, and thereby *orient* it towards that side. A set N of separations is said to *distinguish* the tangles of G_r if for any two of them that orient some separation of G_r differently there is such a

separation even in N. If N is nested and even contains, for every pair of tangles, such a separation of minimum order (among all the separations of G_r that distinguish this pair of tangles), we say that N distinguishes the tangles of G_r *efficiently* and call it a *tree of tangles* for G_r.

While G_r can have many trees of tangles, one of these stands out. Let D be the set of all separations of G_r that distinguish some pair of tangles efficiently. Now given any pair of tangles that can be distinguished by some separation of G_r, associate with this pair the set of all those separations in D which distinguish it efficiently and cross as few other separations in D as possible. Let N be the union of all these sets. This set N clearly distinguishes every pair of tangles efficiently, and it is known to be nested [CK24].

Thus, N is a tree of tangles for G_r. It is canonical in the sense mentioned in Section 1.1.1, since it is defined purely in terms of invariants of G_r. In particular, it is invariant under the deck transformations of G_r over G (whose group we denote by $\mathscr{D}_r$). It is thus a well-defined set of separations of G_r that is unique, given G and r. In recognition of the fact that this tree of tangles of G_r distinguishes the lifts of all the local clusters in G and structures them in G_r, we call it, together with $\mathscr{D}_r$ (which reflects this global structure of G_r back to G), the *r-global structure* of G.

This r-global structure of G appears at the level of G as the quotient $N/\mathscr{D}_r$. This is not a graph. But we show that it can be displayed by a graph, the graph H from Theorem 1, as follows.

Every nested set N of separations of a finite graph can be converted into a tree-decomposition of that graph so that the separations of that graph which correspond to edges of the decomposition tree (the model of the tree-decomposition) are precisely those in N. Infinite graphs in general can have trees of tangles that cannot be represented by a tree-decomposition in this way, because they can have limits under the natural partial ordering of graph separations [Die18a]. For our particular G_r, however, we shall be able to prove that N, our canonical tree of tangles for G_r, does not in fact have problematic limits, and thus gives rise to a canonical tree-decomposition $\mathcal{T}_r = (T_r, (V_t)_{t\in T_r})$ of G_r.

The orbit space $H = T_r/\mathscr{D}_r$ of the tree T_r is a graph. We say that H, and the H-decomposition $\mathcal{G}_r = (H, (V_h)_{h\in H})$ of G such that $V_h = p_r(V_t)$ whenever h is the $\mathscr{D}_r$-orbit of t, *display* the r-global structure $(N, \mathscr{D}_r)$ of G. As parts G_h of $\mathcal{G}_r$ we choose the projections to G of the parts $G_r[V_t]$ of $\mathcal{T}_r$ with $t \in h$. These G_h need not be induced in G although the $G_r[V_t]$ are induced in G_r, since distinct lifts V_t of V_h can be adjacent in G_r. We remark that H is finite if G is finite.

1.1.4 Overview of the proofs of Theorems 1 and 2

We will prove the following common generalisation of these two theorems: every connected, locally finite, and quasi-transitive graph G has a graph-decomposition that displays its r-global structure. (A graph is *quasi-transitive* if its automorphism group has finitely many vertex-orbits.) We will see that, if G is quasi-transitive, then G_r is quasi-transitive too.

Our key task will be to prove that the quasi-transitive graph G_r is *tangle-accessible*; this means that all its tangles are pairwise distinguished by separations of some finitely bounded order. For such tangle-accessible G_r our tree of tangles N exists [CK24], and the separations in N also have finitely bounded order. This implies that N has no problematic limits. In particular, N gives rise to a canonical tree-decomposition of G_r, which will in turn define a canonical graph-decomposition of G that displays its r-global structure.

To show that the quasi-transitive graph G_r is tangle-accessible, we will first use a result of Hamann [Ham18] to see that G_r is *accessible*: that the ends of G_r can be pairwise distinguished by separations of finitely bounded order. By our above considerations there will then exist a canonical tree-decomposition of G_r, of finitely bounded adhesion, that distinguishes all the ends of G_r. We will find a specific such tree-decomposition, $\mathcal{T}_{\text{end}}$, whose infinite parts are 1-ended and themselves quasi-transitive, even just under automorphisms that extend to automorphisms of G_r.

This will allow us to refine $\mathcal{T}_{\text{end}}$ in each of its parts so that the arising canonical tree-decomposition $\mathcal{T}$ distinguishes all the tangles of G_r, not just its ends. This $\mathcal{T}$ will again have finitely bounded adhesion. It is not the canonical tree-decomposition of G_r we ultimately seek (the one whose projection to G defines our desired H-decomposition of G), but it witnesses that G_r is tangle-accessible and thus completes our proof.

1.1.5 A windfall from the proof

The concept of accessibility for locally finite graphs was introduced by Thomassen and Woess [TW93]. It follows their graph-theoretic characterisation of the finitely generated groups that are accessible, in the sense of Wall [Wal71], by Stallings's group splitting theorem [Sta72]. Hamann, Lehner, Miraftab, and Rühmann [HLMR22] extended this characterization to quasi-transitive graphs that need not be Cayley graphs.

In the course of our proof we show that not only the ends of an accessible quasi-transitive graph can be distinguished by separations of bounded order, but all its other tangles can be too:

Theorem 3. *Accessible locally finite, quasi-transitive graphs are in fact tangle-accessible.*

Applied to Cayley graphs of groups, Theorem 3 yields a structural description of all accessible groups that refines their structural decomposition by Stallings's theorem.

1.1.6 Applications

We think of the decomposition theory advanced in this chapter, and in particular of Theorem 1, not so much as a tool with which to attack existing problems in graph theory, but as a natural way to view graphs and to analyse their structure from first principles. We do believe that our decompositions have the potential to interact with other graph invariants, by splitting them into a local and a global aspect as indicated after the statement of Theorem 1. But how exactly this can happen may not be straightforward, and may require non-trivial structural

analysis. This will be worthwhile if, but also only if, our new invariant of $H(G, r)$ and the associated decomposition of G are considered as a natural lens through which graph structure may be viewed. We believe they are.

Still, let us sketch one particular potential application of graph decompositions of this type, just to indicate the idea. Theorem 1 includes, as the special case of $r = |G|$, the canonical strengthening of the so-called *tree of tangles* theorem of Robertson and Seymour [RS91]. This proves the existence of a tree-decomposition that separates out the most highly connected regions of a graph G in a tree-like way. If G is not globally tree-like, this decomposition may be trivial or otherwise fail to capture the global structure of G. If G is globally tree-like, which happens when the graph does not contain some other fixed graph as a minor but is much larger than it, the parts of such a tree-decomposition can be fleshed out in a way that Robertson and Seymour [RS03] have described in detail: they are each 'nearly' embeddable in one of a bounded number of surfaces. Both this bound and the notion of 'nearly' depend only on the excluded minor, not on the graph being decomposed.

For smaller r than $r = |G|$, our Theorem 1 no longer returns a canonical tree-decomposition of G as in the Robertson-Seymour result. Instead, it returns the projection to G of that tree-decomposition applied to its r-local cover. The parts of that projection, those of our canonical H-decomposition of G for $H = H(G, r)$, are not yet fleshed out by structural descriptions in terms of surface embeddings. But it should be possible to do that if it can be done for the tree-decomposition of the cover – which Esperet, Giocanti and Legrand-Duchesne [EGL24] have indeed done recently, albeit with a slightly different canonical tree-decomposition than ours. The result would be a structure theorem, for arbitrary finite graphs G and a given locality parameter r, that G has a canonical H-decomposition for $H = H(G, r)$ whose parts we expect to be nearly embeddable in surfaces in which G, as a whole, does not embed.

Pursuing applications such as this, which seek to apply global structural descriptions of tree-like graphs to the (tree-like) coverings of G in order to then translate them into self-contained theorems about G itself that describe its local structure, we in particular have to be able to describe the decomposition of G and its parts without reference to G_r and the covering map, as happens in its definition. We have indeed shown in a sequel to this chapter [CJKK25] that such a finite description always exists, under only mild assumptions about G and r. This can then form the basis for further structural analysis of the parts. Carmesin and Frenkel have used such a finite description of low-order graph-decompositions [Car22] to develop an efficient algorithm for computing them [CF23]. They have implemented their algorithm and tested it successfully on road networks.

1.1.7 Structure of this chapter

In Section 1.2 we collect together notation and basic facts with an emphasis on coverings. In Section 1.3 we introduce graph-decompositions, and show how we can obtain them from canonical tree-decompositions via coverings. Local coverings are introduced in Section 1.4, where we also investigate their properties as well as local coverings of some special classes of

graphs. In Section 1.5 we define formally what it means for a graph-decomposition to 'display the r-global structure' of a graph, and outline what remains to be done to prove Theorem 1. The proof is then completed in Sections 1.6–1.8.

1.2 Terminology

The purpose of this section is to introduce our terminology wherever this may not be obvious or differ from standard usage.

1.2.1 Graphs

Our graph-theoretic terminology follows [Die25], with the following adaptations.

A *graph* G consists of a set $V(G)$ of *vertices* or *nodes*, a set $E(G)$ of *edges* that is disjoint from $V(G)$, and an *incidence* map that assigns to every edge e an unordered pair $uv = \{u, v\}$ of vertices, its *endvertices*. Note that our graphs may have loops and parallel edges: we allow $u = v$, and distinct edges may have the same pair of endvertices.

The triples (e, u, v) and (e, v, u) are the two *orientations* of an edge e; we think of (e, u, v) as e *oriented from* u *to* v. When G has no parallel edges, e.g. when G is a tree, we abbreviate (e, u, v) to (u, v).

A graph is *locally finite* if each of its vertices is incident with only finitely many edges, and *finite* if its vertex and edge sets are both finite. An *isomorphism* between graphs G and H is a pair $\varphi = (\varphi_V, \varphi_E)$ of bijections $\varphi_V \colon V(G) \to V(H)$ and $\varphi_E \colon E(G) \to E(H)$ that commute with the incidence maps in G and H. We abbreviate both φ_V and φ_E to φ.

1.2.2 Groups

Our group-theoretic terminology largely follows [Löh17]. Given sets $R \subseteq S$ we write $F(S)$ for the free group that S generates, $\langle R \rangle^{\triangleleft}_{F(S)}$ for its normal subgroup generated by R, and $\langle\, S \mid R \,\rangle$ for the quotient group $F(S)/\langle R \rangle^{\triangleleft}_{F(S)}$ presented with generators S and relators R. Given a group Γ with a set S of generators, we write $\varphi_{\Gamma,S}$ for the canonical homomorphism $F(S) \to \Gamma$ that extends the identity on S.

1.2.3 Interplay of graphs and groups

When a group Γ acts on a graph G, we denote the action of an element $g \in \Gamma$ on a vertex or edge $x \in G$ by $x \mapsto g \cdot x$. We write G/Γ for the graph with vertex set $\{\Gamma \cdot v \colon v \in V(G)\}$ and edge set $\{\Gamma \cdot e \colon e \in E(G)\}$ in which the edge $\Gamma \cdot e$ has endvertices $\Gamma \cdot u$ and $\Gamma \cdot v$, where u and v are the endvertices of e in G. Note that G/Γ can have loops and parallel edges even if G does not. We call G/Γ the *orbit graph* of the action of Γ on G.

We say that Γ acts *transitively* on G if it does so with only one vertex orbit, and *quasi-transitively* if it does so with only finitely many vertex orbits. The graph G itself is *(quasi-) transitive* if its full automorphism group acts (quasi-) transitively on its vertices.

Let Γ be a group and $S \subseteq \Gamma$, usually a set of generators. The *Cayley graph* $\mathrm{Cay}(\Gamma, S)$ is the graph with vertex set Γ, edge set $\{(g,s) : g \in \Gamma,\ s \in S\}$ and incidence map $(g,s) \mapsto \{g, gs\}$. We allow S to contain inverse pairs of elements. If it contains both s and s^{-1}, say, and $gs = h$ in Γ, then $\mathrm{Cay}(\Gamma, S)$ has two edges incident with g and h: the edge (g,s) and the edge (h, s^{-1}). Every group acts on its Cayley graphs by left-multiplication, as $h \cdot g := hg$ and $h \cdot (g,s) := (hg, s)$.

1.2.4 Topological and covering spaces

In our topological notation we follow Hatcher [Hat02]. In particular, when X is a path-connected topological space and $\alpha\colon [0,1] \to X$ a path, we write $\alpha^-\colon t \mapsto \alpha(1-t)$ for the path that traverses α backwards. Given two paths α and β with $\alpha(1) = \beta(0)$ we write $\alpha\beta$ for their concatenation, the path which traverses α first and then β, both at double speed. Multiplication in the fundamental group is denoted by $[\alpha][\beta] := [\alpha\beta]$.

For a normal subgroup S of $\pi_1(X, x_0)$ and any $x \in X$ we write S_x for the subgroup $\{\, [\alpha\gamma\alpha^-] : [\gamma] \in S \,\}$ of $\pi_1(X,x)$, where α is any path from x to x_0 in X. Since S is normal, this definition is independent of the choice of α; in particular, S_x is again normal. When X is a graph G, we call S *canonical* if for every vertex $x \in G$ and for every automorphism φ of G we have $\varphi_*(S_x) = S_{\varphi(x)}$.

When we view a graph G as a 1-complex, as we freely will, we use the term 'path' for graph-theoretical paths in the sense of [Die25], and say *topological path* for continuous maps $[0,1] \to G$.

In addition to Hatcher's [Hat02, Ch.1.3] we use standard covering space terminology as in Jänich [Jän97, Ch.9]. We refer to covering projections $p\colon (C, \hat{x}_0) \to (X, x_0)$ as *coverings*, reserving the term *covering space*, or simply *cover*, for C. All covering spaces we consider are connected. The group of deck transformations of a covering p is denoted by $\mathscr{D}(p)$. As generally with automorphisms, we write multiplication in $\mathscr{D}(p)$ as $\varphi \circ \psi =: \varphi\psi$. A normal covering is *canonical* if its characteristic subgroup is canonical.

1.2.5 Topological paths versus graph-theoretic walks

In this section we provide a graph-theoretical description of the fundamental group of a graph, following Stallings [Sta83], on which our terminology is based.

A *walk* (of *length* k) in a graph G is a non-empty alternating sequence $v_0 e_0 v_1 \dots v_{k-1} e_{k-1} v_k$ of vertices and edges in G such that e_i has endvertices v_i and v_{i+1} for all $i < k$. A walk W is *closed* with *base vertex* v_0 if $v_0 = v_k$. We say that W is *based* at v_0, or simply a closed walk *at* v_0. A walk is *trivial* if it has length 0. A walk W' is a *subwalk* of W if $W' = v_i e_i v_{i+1} \dots v_{j-1} e_{j-1} v_j$ for some indices i, j with $0 \leqslant i \leqslant j \leqslant k$. The inverse sequence of a walk W is its *reverse*, or W *traversed backwards*; we denote this walk by W^-.

Let G be any connected graph. A walk in G is *reduced* if it contains no subwalk of the form $ueveu$ where u, v are vertices and e is an edge. Note that every walk W can be turned into a reduced walk W' based at the same vertex by iteratively replacing subwalks of the form $ueveu$ with the trivial walk u. We call W' the *reduction* of W and remark that this is well-defined.

We call two walks (*combinatorially*) *homotopic* if their reductions are equal. This defines an equivalence relation $\sim$ on the set of all walks in G.

Now let $x_0 \in G$ be any vertex. Let $\mathcal{W}(G, x_0)$ be the set of all closed walks in G that are based at x_0. Then $\pi_1'(G, x_0) := (\mathcal{W}(G, x_0)/\sim, \cdot\,)$ is a group, where $[W_1] \cdot [W_2] := [W_1 W_2]$ and $W_1 W_2$ is the concatenation of the two walks W_1 and W_2. Any walk in G defines (up to re-parametrisation) a topological path in the 1-complex G which traverses the vertices and edges of the walk in the same order and direction. This defines a group isomorphism between $\pi_1'(G, x_0)$ and the fundamental group $\pi_1(G, x_0)$ of the 1-complex G [Sta83]. For notational simplicity we will work with combinatorial walks instead of topological paths, and do not always distinguish between walks and the homotopy classes they represent. For example, we may say that a closed walk lies in a subgroup S of $\pi_1(G, x_0)$ when we mean that its combinatorial homotopy class maps to an element of S under the above group isomorphism.

1.3 Graph-decompositions

In this section we introduce graph-decompositions. We show how graph-decompositions of a graph G can be obtained canonically from canonical tree-decompositions of coverings graphs of G (Theorem 1.3.13).

1.3.1 Decompositions

Since graph-decompositions generalise tree-decompositions, we recall the notion of tree-decompositions before we introduce the new notion of graph-decompositions.

Definition 1.3.1 (**Tree-decomposition).** Let G be a graph, T a tree, and $\mathcal{V} = (V_t)_{t \in T}$ be a family of vertex sets $V_t \subseteq V(G)$ indexed by the nodes t of T. The pair $(T, \mathcal{V})$ is called a *tree-decomposition* of G if it satisfies the following two conditions:

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

(T2) for every vertex $v \in G$, the set of nodes t of T such that $v \in V_t$ is connected in T.

The sets V_t of vertices are the *bags* of this decomposition. We sometimes call the subgraphs $G[V_t]$ they induce the decomposition's *parts*, and T its *decomposition tree*.

In this chapter, all bags of tree-decompositions are required to be non-empty. If G is connected, this implies that for edges tt' of T the sets $V_t \cap V_{t'}$ are non-empty too [Die25, Lemma 12.3.1].

A precursor to our notion of graph-decompositions was suggested by Diestel and Kühn [DK05]. But their notion is not quite the same as ours, and we shall not rely on [DK05].

Definition 1.3.2 (**Graph-decomposition).** Let G and H be graphs, and let $\mathcal{V} = (V_h)_{h \in H}$ be a family of sets V_h of vertices of G indexed by the nodes h of H. The pair $(H, \mathcal{V})$ is called an *H-decomposition* of G, or more generally a *graph-decomposition*, if it satisfies the following two conditions:

(H1) $G = \bigcup_{h \in H} G[V_h]$;

(H2) for every vertex $v \in G$, the subgraph $H[W_v]$ of H induced by $W_v := \{ h \in H \mid v \in V_h \}$ is connected.

We call H the *decomposition graph* of this graph-decomposition, and say that it *models* G. The decomposition is *honest* if $V_h \neq \emptyset$ for all nodes h and $V_h \cap V_{h'} \neq \emptyset$ for all edges hh' of H.

The sets V_h in our decomposition are its *bags*; the sets W_v its *co-bags*. The maximum of its bag sizes $|V_h|$, if it exists, is the *width* of our decomposition, the maximum of its co-bag sizes $|W_v|$ its *spread*.[2] If all the bags are finite but their sizes are not necessarily bounded, we say that the decomposition has *finite width*, and similarly for spread.

Note the duality of

$$v \in V_h \Leftrightarrow h \in W_v.$$

This enables us to reconstruct the families $\mathcal{V} = (V_h)_{h \in H}$ and $\mathcal{W} = (W_v)_{v \in G}$ from each other when G and H are given. We call them *dual* to each other.

We often flesh out the bags and co-bags by choosing subgraphs G_h of G on V_h and H_v of H on W_v in such a way that (H1) and (H2) still hold, as

(H1′) $G = \bigcup_{h \in H} G_h$;

(H2′) for every vertex $v \in G$, the subgraph H_v of H is connected.

These G_h, then, are our chosen *parts* of the decomposition $(H, \mathcal{V})$; the H_v our chosen *co-parts*. We may then refer to the decomposition with these choices simply by citing the families $\mathcal{G} = (G_h)_{h \in H}$ and $\mathcal{H} = (H_v)_{v \in G}$ together with G and H. Note that $\mathcal{V}$ can be reconstructed from either of these families, and that it satisfies (H1) and (H2) if $\mathcal{G}$ and $\mathcal{H}$ satisfy (H1′) and (H2′).

Honest graph-decompositions into connected parts are closed under duality:

Lemma 1.3.3. *If $(H, \mathcal{V})$ is an honest graph-decomposition of G that has connected parts, then its dual $(G, \mathcal{W})$ is an honest graph-decomposition of H that has connected parts.*

Proof. The decomposition $(G, \mathcal{W})$ satisfies (H1), because $(H, \mathcal{V})$ is honest. It satisfies (H2), since $(H, \mathcal{V})$ has connected parts G_h, so the $G[V_h] \supseteq G_h$ from (H2) for $(G, \mathcal{W})$ are connected too. Thus, $(G, \mathcal{W})$ is a graph-decomposition. We can choose connected parts $H_v := H[W_v]$ for it by (H2) for $(H, \mathcal{V})$.

Let us show that $(G, \mathcal{W})$ is honest. Its bags W_v are non-empty by (H1) for $(H, \mathcal{V})$, which ensures that v lies in some V_h so that $W_v \ni h$. Similarly, every edge $e = vv'$ of G lies in some $G[V_h]$ by (H1), so that $h \in W_v \cap W_{v'} \neq \emptyset$. □

Within the realm of honest graph-decompositions into connected parts, dualising is clearly an involution.

[2] When H is a tree, this differs from the traditional width parameter for tree-decompositions by 1.

An *isomorphism* between graph-decompositions $(H, \mathcal{V})$ and $(H', \mathcal{V}')$ of graphs G and G', where $\mathcal{V} = (V_h)_{h\in H}$ and $\mathcal{V}' = (V'_{h'})_{h'\in H'}$, is a pair (φ, ψ) of isomorphisms $\varphi\colon G \to G'$ and $\psi\colon H \to H'$ such that $\varphi(V_h) = V'_{\psi(h)}$ for every node $h \in H$. If we have chosen families $\mathcal{G} = (G_h)_{h\in H}$ and $\mathcal{G}' = (G'_{h'})_{h'\in H'}$ of parts $G_h \subseteq G[V_h]$ and $G'_{h'} \subseteq G'[V'_{h'}]$ for these decompositions, we call our pair (φ, ψ) an *isomorphism* between $(H, \mathcal{G})$ and $(H', \mathcal{G}')$ if $\varphi(G_h) = G'_{\psi(h)}$ for every node $h \in H$.

1.3.2 Properties of graph-decompositions

A *separation* of a set X is an unordered pair $\{A, B\}$ of subsets A, B of X such that $A \cup B = X$. The sets A and B are the *sides* of this separation, their intersection $A \cap B$ is the associated *separator* of X, and $|A \cap B|$ is its *order*. The separation $\{A, B\}$ is *proper* if neither A nor B equals X. The *orientations* of $\{A, B\}$ are the two ordered pairs (A, B) and (B, A); these are *oriented separations* of X. We denote the two orientations of a separation s as $\vec{s}$ and $\overleftarrow{s}$. There are no default orientations: once we have called one of the two orientations $\vec{s}$, the other will be $\overleftarrow{s}$, and conversely. A *separation* of a graph G is a separation $\{A, B\}$ of its vertex set $V(G)$ such that G has no edge between $A \setminus B$ and $B \setminus A$.

Let $\mathcal{T} = (T, \mathcal{V})$ be a tree-decomposition of a graph G. Every orientation (t, t') of an edge $e \in T$ *induces* an oriented separation $\alpha_{\mathcal{T}}(t, t')$ of G, as follows. Let T_t and $T_{t'}$ be the components of $T - e$ containing t and t', respectively. By (T1), $\alpha_{\mathcal{T}}(t, t') := (A_t, A_{t'})$ is an oriented separation of G with sides $A_t = \bigcup_{s\in T_t} V_s$ and $A_{t'} = \bigcup_{s\in T_{t'}} V_s$. Similarly, $\alpha_{\mathcal{T}}(e) := \{A_t, A_{t'}\}$ is an unoriented separation of G. By (T2), the separator $A_t \cap A_{t'}$ of $\{A_t, A_{t'}\}$ equals $V_t \cap V_{t'}$ [Die25, Lemma 12.3.1]; this is the *adhesion set* of $(T, \mathcal{V})$ at the edge $e \in T$. A tree-decomposition has *finitely bounded adhesion* if there exists $K \in \mathbb{N}$ such that all adhesion sets of $(T, \mathcal{V})$ have order at most K.

Notice that, in the above example, the vertex sets of the components T_t and $T_{t'}$ of $T - e$ formed a separation of the set $V(T)$. It is not a separation of the graph T, since the edge e lies in neither componenent, but it induces a separation of the graph G. In general, given a graph-decomposition of G modelled on a graph H, separations of the set $V(H)$ induce separations of the graph G in the same way:

Lemma 1.3.4. *Let $(H, (V_h)_{h\in H})$ be a graph-decomposition of a graph G. Every separation $\{U, W\}$ of the set $V(H)$ induces a separation $\{A, B\}$ of the graph G with $A := \bigcup_{h\in U} V_h$ and $B := \bigcup_{h\in W} V_h$. The separator of $\{A, B\}$ is equal to $S := (\bigcup_{h\in U\cap W} V_h) \cup \bigcup_{hh'\in F}(V_h \cap V_{h'})$, where $F = E_H(U \setminus W, W \setminus U)$.*

Proof. By (H1), every edge e of G is contained in some part $G[V_h]$. Then $h \in U$ or $h \in W$, and $e \in G[A]$ or $e \in G[B]$ respectively. In particular, there is no edge in G between $A \setminus B$ and $B \setminus A$. Thus, $\{A, B\}$ is a separation of G.

It is immediate from the definition of A and B that $A \cap B \supseteq S$. To show $A \cap B \subseteq S$, let $v \in A \cap B$ be any vertex. By definition of A and B there exist $h \in U$ with $v \in V_h$ and $h' \in W$ with $v \in V_{h'}$. Thus, $h, h' \in W_v$. Since $H[W_v]$ is connected (H2), it contains an h–h' path from

U to W; choose h and h' so that this path has minimum length. Since $\{U, W\}$ is a separation of $V(H)$, this means that either $h = h' \in U \cap W$ or $hh' \in F$. In either case, $v \in V_h \cap V_{h'} \subseteq S$ as desired. □

Recall that an *end* of a graph G is an equivalence class of rays in G, where two rays in G are *equivalent* if for every finite set X of vertices of G the two rays have subrays in the same component of $G - X$. As a first reflection of the relationship between the separations of H and G observed in Lemma 1.3.4, let us show that the ends of a graph G often correspond naturally to those of its models H.

To define this relationship, we use a tangle-like description of ends known as directions. A *direction* in a graph G is a map f, with domain the set of all finite vertex sets of G, that assigns to every finite vertex set $X \subseteq V(G)$ a component $f(X)$ of $G - X$ so that $f(X) \supseteq f(X')$ whenever $X \subseteq X'$. Every end ω of G defines a direction f_ω in G by letting $f_\omega(X)$ be the component of $G - X$ that contains a subray of one (equivalently: every) ray in ω. The map $\omega \mapsto f_\omega$ is a bijection between the ends of G and its directions [DK03, Theorem 2.2]. Therefore, in order to define a bijection between the ends of H and G for an H-decomposition of G, it suffices to define a bijection between the directions of H and of G.

Let G be any graph, and let $(H, (V_h)_{h \in H})$ be a graph-decomposition of G of finite width. Every direction f in G defines a direction g_f in H, as follows. For every finite set $Y \subseteq V(H)$, the set $X_Y := \bigcup_{h \in Y} V_h$ of vertices in G is finite. Every W_v with $v \in f(X_Y)$ lies entirely outside Y, and hence by (H2) lies in one component of $H - Y$. This is the same component of $H - Y$ for all these v by (H1), since $h \in W_u \cap W_v$ and thus $W_u \cap W_v \neq \emptyset$ whenever an edge uv lies in $G[V_h]$ and $f(X_Y)$ is connected. Thus, all of $\bigcup_{v \in f(X_Y)} W_v$ lies in one component of $H - Y$; let $g_f(Y)$ be this component. It is straightforward to verify that g_f is a direction in H.

Lemma 1.3.5. *For every honest H-decomposition of G of finite width and spread, and with connected parts, the map $f \mapsto g_f$ between the directions of G and those of H is bijective.*

The bijection in Lemma 1.3.5 between the directions of G and H induces a bijection between their ends, by [DK03, Theorem 2.2] as mentioned earlier. That bijection is in fact a homeomorphism between the end spaces of G and H; see [Die25] for definitions.

Proof of Lemma 1.3.5. Let the given decomposition of G be denoted as $(H, \mathcal{V})$, where $\mathcal{V} = (V_h)_{h \in H}$ as usual, and denote its (connected) parts by G_h. We establish an inverse to the map $f \mapsto g_f$ by constructing a map $g \mapsto f_g$ from the directions of H to those of G such that $f_{g_f} = f$ for all f and $g_{f_g} = g$ for all g.

Let a direction g of H be given. To define f_g, let a finite set $X \subseteq V(G)$ be given. Let Y_X be the set of nodes $h \in H$ whose bag V_h contains some vertex of X. Note that Y_X is finite, since all the W_v are finite. The subgraph $\bigcup_{h \in g(Y_X)} G_h$ of $G - X$ is connected, because $g(Y_X) \subseteq H - Y$ is connected, the parts $G_h \subseteq G$ are connected, and $(H, \mathcal{G})$ is honest. Let $f_g(X)$ be the unique component of $G - X$ that contains $\bigcup_{h \in g(Y_X)} G_h$.

To check that f_g is a direction, let $X \subseteq X'$ be two finite sets of vertices of G. Then $Y_X \subseteq Y_{X'}$. Since g is a direction, we have $g(Y_X) \supseteq g(Y_{X'})$. Hence $f_g(X) \supseteq \bigcup_{h \in g(Y_X)} G_h \supseteq \bigcup_{h \in g(Y_{X'})} G_h$. Now since $X \subseteq X'$, every component of $G - X'$ is contained in a unique component of $G - X$. As $f_g(X')$ and $f_g(X)$ both contain $\bigcup_{h \in g(Y_{X'})} G_h \neq \emptyset$, the unique component of $G - X$ containing $f_g(X')$ is $f_g(X)$.

For a proof that $f = f_{g_f}$ for every direction f of G, we show that for every finite set $X \subseteq V(G)$ both $f(X)$ and $f_{g_f}(X)$ contain $f(X_{Y_X}) \neq \emptyset$, and must therefore coincide. As $X \subseteq X_{Y_X}$ we have $f(X) \supseteq f(X_{Y_X})$, since f is a direction. For a proof of $f_{g_f}(X) \supseteq f(X_{Y_X})$ let $v \in f(X_{Y_X})$ be given. By (H1) there is an $h \in H$ such that $v \in V_h$. Then $h \in W_v \subseteq g_f(Y_X)$ by definition of g_f and the choice of v, so $G_h \subseteq f_{g_f}(X)$ by definition of f_{g_f}. In particular, $v \in G_h \subseteq f_{g_f}(X)$ as desired.

To show $g = g_{f_g}$ for every direction g of H, we show that for every finite set Y of nodes of H both $g(Y)$ and $g_{f_g}(Y)$ contain $g(Y_{X_Y}) \neq \emptyset$, and must therefore coincide. Since $Y \subseteq Y_{X_Y}$, we have $g(Y) \supseteq g(Y_{X_Y})$ because g is a direction of H. For a proof of $g_{f_g}(Y) \supseteq g(Y_{X_Y})$, let $h \in g(Y_{X_Y})$ be given. Then $V_h \neq \emptyset$ since the parts of $(H, \mathcal{V})$ are connected; pick $v \in V_h$. Now $v \in V_h \subseteq f_g(X_Y)$ by definition of f_g and the choice of h, so $h \in W_v \subseteq g_{f_g}(Y)$ by the definition of g_{f_g}, as desired. □

Next, let us show that not only do the ends of H and G correspond bijectively, under the assumptions from Lemma 1.3.5, but the sizes of the separators needed to distinguish these ends correspond too.

A finite-order separation $\{A, B\}$ of G *distinguishes* two ends ω_1, ω_2 of G if for some (and hence every) ray in ω_1 some subray lies in $G[A]$ and for some (and hence every) ray in ω_2 some subray lies in $G[B]$. Following Thomassen and Woess [TW93], let us call a graph *accessible* if there exists an integer $K \in \mathbb{N}$ such that every two of its ends are distinguished by a separation of order at most K. Compare Section 1.1.5 for the origin of this notion.

Lemma 1.3.6. *Let G be a graph with an honest H-decomposition of finitely bounded width, finite spread, and with connected parts. If H is accessible, so is G.*

Proof. Let $(H, (V_h)_{h \in H})$ be the decomposition of G assumed to exist, let $n \in \mathbb{N}$ be an upper bound on the size of its parts, and let $k \in \mathbb{N}$ be such that every two ends of H are distinguished by a separation of order at most k. We show that every two ends of G are distinguished by a separation of order at most nk.

Let ω_1 and ω_2 be two distinct ends of G, and let $f_1 := f_{\omega_1}$ and $f_2 := f_{\omega_2}$ be the two directions of G they define. Consider the directions $g_1 := g_{f_1}$ and $g_2 := g_{f_2}$ of H as defined just before Lemma 1.3.5. By [DK03, Theorem 2.2], there exist ends η_1 and η_2 of H which induce these directions g_1 and g_2. Since H is accessible, it has a separation $\{U_1, U_2\}$ of order at most k such that, for $i = 1, 2$, the end η_i is represented by a ray in $H[U_i]$, and thus $g_i(U_1 \cap U_2) \subseteq H[U_i]$.

Let us apply Lemma 1.3.4 to $\{U_1, U_2\}$ to obtain a separation $\{A_1, A_2\}$ of G with $A_i = \bigcup_{h \in U_i} V_h$. Note that $A_1 \cap A_2 = \bigcup_{h \in U_1 \cap U_2} V_h$: as $\{U_1, U_2\}$ is a separation not just of $V(H)$ but

of the graph H, the set F in Lemma 1.3.4 is empty. Hence $\{A_1, A_2\}$ has order at most nk; it remains to show that it distinguishes ω_1 and ω_2.

By definition of f_i, each ω_i has a ray in $f_i(A_1 \cap A_2)$. So it suffices to show that $V(f_i(A_1 \cap A_2)) \subseteq A_i$, for $i = 1, 2$. For every $v \in G$ that is not in $A_1 \cap A_2 = \bigcup_{h \in U_1 \cap U_2} V_h$, the graph $H[W_v]$ avoids $U_1 \cap U_2$, and hence by (H2) lies in a component of $H - (U_1 \cap U_2) \neq \emptyset$. For $v \in f_i(A_1 \cap A_2)$ this component is $g_i(U_1 \cap U_2) \subseteq H[U_i]$, by definition of $g_{f_i} = g_i$. Pick $h \in W_v$ for any such v. Then $v \in V_h$ and $h \in U_i$. This implies $v \in A_i$, as desired, by definition of A_i. □

1.3.3 From tree-decompositions to graph-decompositions

In the remainder of this section we show how graph-decompositions of a graph can be obtained from tree-decompositions of its normal covers, and how all these can be chosen canonically. The main result in this section will be Theorem 1.3.13.

A tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of a connected graph G is *regular* if all the separations of G induced by the edges of T are proper: if none has the form $\{A, V(G)\}$. In particular, then, $\mathcal{T}$ cannot have empty adhesion sets and hence no empty bags, so regular tree-decompositions of connected graphs are honest. All tree-decompositions relevant to us will be regular.

We shall be interested particularly in tree-decompositions $\mathcal{T}$ on which the automorphisms φ of the graph being decomposed act naturally, in that they extend to automorphisms (φ, ψ) of $\mathcal{T}$. These automorphisms ψ are easily seen to be unique when $\mathcal{T}$ is regular:

Lemma 1.3.7. *If $\mathcal{T} = (T, \mathcal{V})$ is a regular tree-decomposition of a graph G, then for every automorphism φ of G there exists at most one automorphism ψ of T such that (φ, ψ) is an automorphism of $\mathcal{T}$.*

Proof. Consider automorphisms ψ, ψ' of T such that both (φ, ψ) and (φ, ψ') are automorphisms of $\mathcal{T}$. Then so is $(\mathrm{id}, \psi' \circ \psi^{-1})$. If $\psi' \circ \psi^{-1} = \mathrm{id}$, then $\psi' = \psi$. Thus, it is left to show that the only automorphism ψ of T for which (id, ψ) is an automorphism of $\mathcal{T}$ is $\psi = \mathrm{id}$.

Let us show first that it is enough to show that for every oriented separation (A, B) induced by $\mathcal{T}$, the graph $T_{(A,B)}$ which consists of all oriented edges of T that induce (A, B) is a finite directed path. Assuming this, let (id, ψ) be any automorphism of $\mathcal{T}$. Let (t, t') be an oriented edge of T, and let (A, B) be the separation induced by (t, t'). Since (id, ψ) is an automorphism of $\mathcal{T}$, the automorphism ψ of T maps every oriented edge of T to an oriented edge of T that induces the same separation. Thus, ψ restricts to an automorphism of $T_{(A,B)}$. As the only automorphism of a finite directed path is the identity map, the restriction of ψ to $T_{(A,B)}$ is the identity map. Hence, $(\psi(t), \psi(t')) = (t, t')$. Since the edge oriented (t, t') of T was chosen arbitrarily, ψ is the identity map on the whole decomposition tree T.

It remains to show, then, that $T_{(A,B)}$ is a finite directed path, for every separation (A, B) induced by $\mathcal{T}$. Let (A, B) be a separation induced by $\mathcal{T}$. Since $\mathcal{T}$ is regular, $T_{(A,B)}$ does not contain (t, t') and (t', t) for any edge tt' of T; otherwise, $(A, B) = (B, A)$, and thus $(A, B) = (V(G), V(G))$ - contradicting the fact that $\mathcal{T}$ is regular. Let (t_0, t_1) and (t_2, t_3) be two oriented edges of T which both induce (A, B), then either $(t_0, t_1) \leqslant (t_2, t_3)$ or $(t_2, t_3) \leqslant (t_0, t_1)$,

i.e., the unique $\{t_0, t_1\}$–$\{t_2, t_3\}$ path in T is t_1Tt_2 or t_3Tt_0, respectively. Indeed, if not it follows from the compatibility of the partial order on the oriented edges of $\mathcal{T}$ and its induced separations that $(A, B) \leqslant (B, A)$ or $(B, A) \leqslant (A, B)$; in particular, $A = V(G)$ or $B = V(G)$ – a contradiction to $\mathcal{T}$ being regular. In both cases this compatibility also shows that all oriented edges of T between these two edges induce (A, B) as well. This shows that the underlying graph of $T_{(A,B)}$ is connected and also that at each node t of $T_{(A,B)}$ there is at most one outgoing edge which induces (A, B) and at most one outgoing edge which induces (B, A). Using (T1), it is not hard to show that for every ray $t_0t_1t_2\dots$ in T we have $V(G) = \bigcup_{i\in\mathbb{N}} A_i$, where (A_i, B_i) is the separation induced by (t_i, t_{i+1}). Since $\mathcal{T}$ is regular, $A, B \neq V(G)$, and thus the underlying graph of $T_{(A,B)}$ does not contain a ray. Hence, $T_{(A,B)}$ is finite. So, $T_{(A,B)}$ is a finite directed path. □

We believe that Lemma 1.3.7 extends to H-decompositions for graphs H other than trees, with an appropriate adaptation of the regularity requirement. For example, the extended version is easy to prove when H is 2-connected and its bonds induce proper separations of G, since this implies that $V_h \neq V_{h'}$ for distinct nodes $h, h' \in H$.

Let Γ be a group of automorphisms of G. A graph-decomposition $(H, \mathcal{V})$ or $(H, \mathcal{G})$ of G is Γ-*canonical* if Γ acts on $(H, \mathcal{V})$ or $(H, \mathcal{G})$, respectively, via $\varphi \mapsto (\varphi, \psi)$ for suitable automorphisms $\psi =: \psi_\varphi$ of H. Once the action of Γ on $(H, \mathcal{V})$ or $(H, \mathcal{G})$ is fixed, or if it is unique by Lemma 1.3.7, we usually write φ instead of ψ_φ to reduce clutter. If $(H, \mathcal{V})$ or $(H, \mathcal{G})$ is $\mathrm{Aut}(G)$-canonical, we just call it a *canonical* graph-decomposition of G. When $(H, \mathcal{V})$ is a tree-decomposition, this definition formalises the notion of canonicity used informally for tree-decompositions in the literature [CDHH16a; CDHS14].

The graph-decompositions given by the following construction will be the key objects in Theorem 1.3.13.

Construction 1.3.8 (Graph-decompositions defined by tree-decompositions via coverings)**.** Let $p\colon C \to G$ be any normal covering of a connected graph G. Let $\mathscr{D} \subseteq \mathrm{Aut}(C)$ be the group of deck transformations, and let $\mathcal{T} = (T, \mathcal{V}_T)$ be any regular $\mathscr{D}$-canonical tree-decomposition of C. By Lemma 1.3.7, every $\varphi \in \mathscr{D}$ acts uniquely on $\mathcal{T}$, and in particular on T, so that $\varphi(C[V_t]) = C[V_{\varphi(t)}]$ for all $t \in T$. The map $t \mapsto p(C[V_t])$ is constant on the orbits in $V(T)$ under the action of $\mathscr{D}$ on T, since for every $\varphi \in \mathscr{D}$ we have

$$p(C[V_t]) = p(\varphi(C[\,V_t\,])) = p(C[V_{\varphi(t)}]). \tag{1.1}$$

Let H be the orbit graph $T/\mathscr{D}$ of the action of $\mathscr{D}$ on T. Let $\mathcal{V}_G := (V_h)_{h\in H}$, where $V_h = p(V_t)$ for any $t \in h$. These V_h are well-defined by (1.1). We shall prove in a moment that $(H, \mathcal{V}_G)$ is a graph-decomposition of G.

Let $\mathcal{G} := (G_h)_{h\in H}$, where $G_h = p(C[V_t])$ for any $t \in h$. These G_h are well-defined by (1.1); we choose them as the parts of $(H, \mathcal{V}_G)$. Note that, while the parts $C[V_t]$ of $(T, \mathcal{V}_T)$ were induced subgraphs of C, their images G_h under p need not be induced in G.

Let $\mathcal{H} = (H_v)_{v \in G}$ be the family of subgraphs H_v of H whose nodes and edges are the $\mathscr{D}$-orbits in T of the nodes and edges of $T_{\hat{v}} := T[\,t \in T \mid \hat{v} \in V_t\}]$, where $\hat{v}$ is any vertex of C in $p^{-1}(v)$. These H_v are independent of the choice of $\hat{v}$ since, as p is normal, $\mathscr{D}$ acts transitively on the fibres of p. Note that these H_v are spanning subgraphs of $H[\{\, h \mid v \in V_h\}]$; we may thus choose these H_v as the co-parts of $(H, \mathcal{V}_G)$. Once more, the H_v need not be induced subgraphs of H, although the co-parts $T_{\hat{v}}$ of $(T, \mathcal{V}_T)$ are induced in T.

We say that $\mathcal{T}$ *defines, via* p, the graph-decomposition $(H, \mathcal{V}_G)$ of G with parts G_h and co-parts H_v.

Construction 1.3.8 can also be applied to graph-decompositions, rather than tree-decompositions, of the normal cover C. But we shall not need that level of generality here.

Lemma 1.3.9. *Every $(H, \mathcal{V}_G)$ constructed as in Construction 1.3.8 is an honest graph-decomposition of G. Its parts G_h and co-parts H_v satisfy* (H1′) *and* (H2′).

Proof. With its induced parts $C[V_t]$ and co-parts $T_{\hat{v}}$, the given tree-decomposition $(T, \mathcal{V}_T)$ of C satisfies not only (H1) and (H2) but even (H1′) and (H2′). It is straightforward to check that $(H, \mathcal{V}_G)$ inherits (H1′) and (H2′) from $(T, \mathcal{V}_T)$ for our choice of parts G_h and co-parts H_v. As $(T, \mathcal{V}_T)$ is regular, and hence honest, $(H, \mathcal{V}_G)$ is honest too. □

The graph-decompositions $(H, \mathcal{G})$ from Construction 1.3.8 also satisfy an analogue of (H2) for edges:

(H3) For every edge $e \in G$, the graph $H[\{\, h \in H \mid e \in G_h \,\}]$ is connected.

Note that (H3) refers to a given choice of parts G_h. This is deliberate, and we are always free to choose these as $G[V_h]$. But in the context of Construction 1.3.8 that would be unnatural when G_h is not an induced subgraph of G, because G has 'external' edges on V_h that are not p-images of edges of some $C[V_t]$ with $t \in h$.

The decompositions $(H, \mathcal{G})$ from Construction 1.3.8 also satisfy the analogue

(H3′) For every edge $e \in G$, the subgraph H_e of H is connected

of (H2′) for edges if we define the H_e in the natural way: as the subgraphs of H whose nodes and edges are the $\mathscr{D}$-orbits in T of the subtree $T_{\hat{u}} \cap T_{\hat{v}}$ of T, where $\hat{e} = \hat{u}\hat{v}$ is any lift of $e = uv$ under p. The proofs of (H3) with the G_h from Construction 1.3.8, and of (H3′) with the H_e chosen as above, are analogous to the proofs for vertices: they reduce these properties to the analogous properties of the tree-decomposition $(T, \mathcal{V}_T)$ of C.

Condition (H3) is a familiar property of tree-decompositions. It hinges on the fact that any non-empty intersection of two subtrees of a (decomposition) tree is connected. Connected subgraphs of other graphs H can have disconnected intersections, though, which is why (H3) can fail in graph-decompositions not induced by Construction 1.3.8:

Example. Let $G := K^2$ with vertices u, v and let $H := C_4 = abcda$. Let $G_a := G_c := G$, let $G_b := (\{u\}, \emptyset)$ and let $G_d := (\{v\}, \emptyset)$. Then $(H, \mathcal{G})$ satisfies (H1′) and (H2). However, the set $\{\, h \in H \mid uv \in G_h\} = \{a, c\}$ is disconnected in H. □

The graph H obtained in Construction 1.3.8 can in general be infinite even when G is finite. However this will not happen if $\mathcal{T}$ has finite spread (which will always be the case):

Lemma 1.3.10. *The co-parts H_v of the graph-decomposition $(H, \mathcal{V}_G)$ constructed in Construction 1.3.8 are finite if the tree-decomposition $\mathcal{T}$ defining it has finite spread. If both the H_v and G are finite, then H is finite.*

Proof. The first assertion is immediate from the definition of H_v; note that H_v has only finitely many edges as well as nodes, since all its edges are orbits of edges of some finite tree $T_{\hat{v}} \subseteq T$.

As $\mathcal{T}$, being regular by assumption, is honest, we have $T \subseteq \bigcup_{\hat{v} \in V(C)} T_{\hat{v}}$. Since $H = T/\mathscr{D}$, this implies $H \subseteq \bigcup_{v \in V(G)} H_v$. Hence if both the H_v and G are finite, then so is H. □

For our proof of Theorem 1 we want the graph-decompositions $(H, \mathcal{G})$ of G obtained via Construction 1.3.8 to be canonical. If all the automorphisms of G are induced via p by automorphisms of C then, as we shall see in Lemma 1.3.12, the canonicity of the tree-decomposition $\mathcal{T}$ in Construction 1.3.8 will ensure this. In Lemma 1.3.11 we provide a sufficient condition on p to ensure this, that all the automorphisms of G 'lift' to C in this way.

In general, given a covering $p\colon C \to G$ of a connected graph G, a *lift* of an automorphism φ of G is an automorphism $\hat{\varphi}$ of C such that $p \circ \hat{\varphi} = \varphi \circ p$. Lifts of automorphisms are unique up to composition with deck transformations, since for any two lifts $\hat{\varphi}, \hat{\varphi}'$ of φ their composition $\hat{\varphi}' \circ \hat{\varphi}^{-1}$ is a deck transformation.

Lemma 1.3.11 below says that the automorphisms of G all lift to C if the covering $p\colon C \to G$ is canonical (see Section 1.2.4):

Lemma 1.3.11 ([Djo74], Theorem 1)**.** *Let G be a connected graph, and $p\colon C \to G$ a canonical normal covering. Then every automorphism φ of G lifts to an automorphism of C.*

Lemma 1.3.12. *The graph-decompositions $(H, \mathcal{G})$ from Construction 1.3.8 are canonical if $\mathcal{T}$ and p are canonical.*

Proof. To prove $(H, \mathcal{G})$ to be canonical, we have to find for every automorphism φ of G an automorphism (φ, ψ) of $(H, \mathcal{G})$ such that $\varphi \mapsto (\varphi, \psi)$ is a group homomorphism.

Let φ be an automorphism of G. As p is canonical, Lemma 1.3.11 tells us that φ lifts to an automorphism $\hat{\varphi}$ of C. as of course does $\mathscr{D}$. Since $\hat{\varphi}$ is a lift of φ, its action on T is well defined on the orbits of $\mathscr{D}$ in T, i.e., $\psi\colon \mathscr{D} \cdot t \mapsto \mathscr{D} \cdot \hat{\varphi}(t)$ is a well-defined automorphism of $H = T/\mathscr{D}$. The definition of ψ is also independent of the choice of $\hat{\varphi}$, given φ, since $\hat{\varphi}$ is unique up to composition with deck transformations.

It remains to check that our choice of ψ makes $\varphi \mapsto (\varphi, \psi)$ into a group homomorphism from the automorphisms of G to those of $(H, \mathcal{G})$. Let φ_0, φ_1 be automorphisms of G, with lifts $\hat{\varphi}_0$ and $\hat{\varphi}_1$. Then $\hat{\varphi}_0 \circ \hat{\varphi}_1$ is a lift of $\varphi_0 \circ \varphi_1$. So it follows directly from the definition of ψ that the automorphism induced on the orbit graph $H = T/\mathscr{D}$ by the action of $\hat{\varphi}_0 \circ \hat{\varphi}_1$ on T is $\psi_0 \circ \psi_1$, where ψ_0 and ψ_1 are the automorphisms induced by the action of $\hat{\varphi}_0$ and $\hat{\varphi}_1$, respectively. So, $\varphi \mapsto (\varphi, \psi)$ is a group homomorphism, and thus $(H, \mathcal{G})$ is canonical. □

Our main result of this section now follows directly from Lemmas 1.3.9, 1.3.10 and 1.3.12:

Theorem 1.3.13. *Let G be any connected graph, and let $p\colon C \to G$ be a normal covering of G. Then every regular $\mathscr{D}(p)$-canonical tree-decomposition $\mathcal{T}$ of C defines an honest graph-decomposition $(H, \mathcal{G})$ of G via Construction 1.3.8.*

If $\mathcal{T}$ and p are canonical, then $(H, \mathcal{G})$ is canonical. If $\mathcal{T}$ has finite spread, then the co-parts H_v of $(H, \mathcal{G})$ are finite. If the H_v and G are finite, then H is finite. □

1.3.4 Examples

In this section we illustrate Theorem 1.3.13 by a number of examples highlighting various aspects of Construction 1.3.8. All the coverings considered in this section are local coverings as defined in Section 1.4.

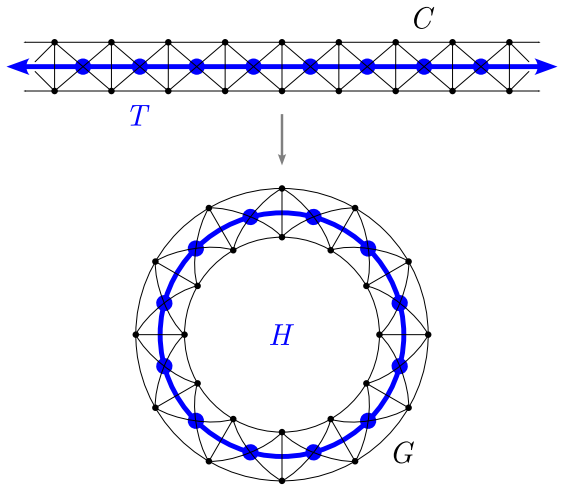

Figure 1.2: The global structure of G is displayed by a cycle. Its local parts are K^4s.

Example 1.3.14. Let G be the graph depicted in Figure 1.2. It is formed from twelve K^4s by glueing them along pairs of vertices in the shape of a cycle. The graph C at the top of Figure 1.2 is a normal cover of G. The deck transformation group $\mathscr{D}$ of this covering is isomorphic to the infinite cyclic group $\mathbb{Z}$; it acts on C by shifting the K^4s to the left or right by multiples of 12.

Let $\mathcal{T} = (T, \mathcal{V})$ be the canonical tree-decomposition of C whose decomposition tree T is a double ray and whose parts are the copies of K^4 in C. The decomposition graph $H = T/\mathscr{D}$ of the graph-decomposition $(H, \mathcal{G})$ defined by $\mathcal{T}$ (Construction 1.3.8) is a cycle of length 12; its family $\mathcal{G}$ of parts consists of the twelve K^4s which form G. □

Our next example justifies why we do not require the parts of a graph-decomposition of a graph G to be induced subgraphs of G. The example consists of a graph-decomposition of G that is defined by a tree-decomposition of a normal cover of G, just as we need it for Construction 1.3.8. But one of its parts is not an induced subgraph of G.

Example 1.3.15. Let G be the graph depicted in Figure 1.3 which is formed by a K^6 and a ladder with contracted first and last rungs such that the two respective contraction vertices

are identified with two distinct vertices of the K^6. The graph C at the top of Figure 1.3 is a normal cover of G. The deck transformation group $\mathscr{D}$ of this cover is the infinite cyclic group $\mathbb{Z}$, which acts on C by shifting.

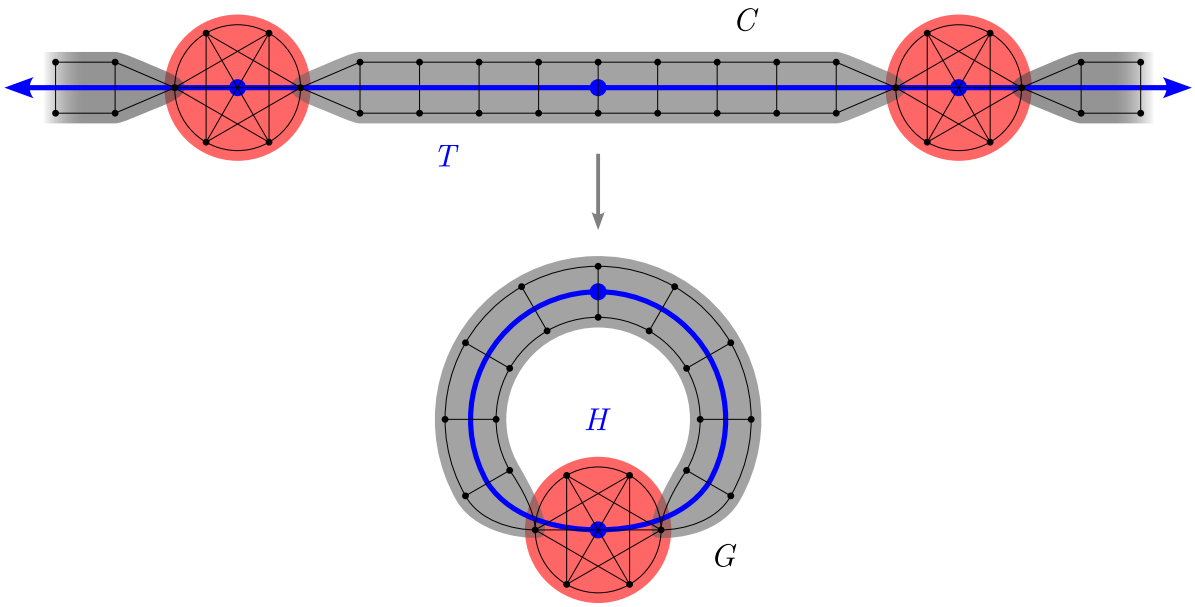

Figure 1.3: The grey part of H is not an induced subgraph of G

Figure 1.3 shows a canonical tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of C. Its decomposition tree T is a double ray; its parts alternate between red K^6s and copies of the modified ladder. All adhesion sets have size one. The graph-decomposition $(H, \mathcal{G})$ of G defined by $\mathcal{T}$ as in Construction 1.3.8 has a decomposition graph $H = T/\mathscr{D}$ on two vertices, which are joined by two parallel edges. It has two parts, the K^6 and the modified ladder. The latter is not an induced subgraph of G: its two endpoints are joined by an edge of the K^6, which is an edge of G but not an edge of the ladder part. □

Our final example shows that if $(H, \mathcal{G})$ is a graph-decomposition defined by $(T, \mathcal{V})$ via a normal covering $p\colon C \to G$, then the decomposition tree T of C need not be a covering space of the decomposition graph H of G. Indeed, one can show that T is a cover of H if and only if the deck transformation group $\mathscr{D}(p)$ acts properly discontinuously on T, i.e., if and only if every deck transformation of p either fixes all the parts of $(T, \mathcal{V})$ or none (cf. [Hat02, Proposition 1.40]).

Example 1.3.16. Let G be a graph obtained from two copies of some quadrangulation of the torus by glueing them together at a single vertex. Note that every quadrangulation of the torus is covered by the infinite $(\mathbb{Z} \times \mathbb{Z})$-grid. Thus, a normal cover C of G is obtained by identifying infinitely many $(\mathbb{Z} \times \mathbb{Z})$-grids in single vertices to form an $\aleph_0$-regular 'tree of grids'. The deck transformation group $\mathscr{D}$ of this cover is isomorphic to the free product $(\mathbb{Z} \times \mathbb{Z}) * (\mathbb{Z} \times \mathbb{Z})$. Let $(T, \mathcal{V})$ be the canonical tree-decomposition of C into its $(\mathbb{Z} \times \mathbb{Z})$-grids. The graph-decomposition $(H, \mathcal{G})$ defined by $(T, \mathcal{V})$ then has $H = K^2$ as its decomposition graph, one node for

each of two the quadrangulated tori of G. Since H has nodes of degree 1 but T does not, T cannot cover H. □

1.4 Local coverings

Our aim in this section is to construct graph-decompositions $(H, \mathcal{G})$ of a given graph G that display its global structure only, leaving the local aspects of the structure of G to the parts of $\mathcal{G}$. We shall use the framework of Construction 1.3.8 to achieve this, but will have to choose the covering $p\colon C \to G$ on which it is based specifically to suit our aim. Our choice of this covering will depend only on our intended level of 'locality', a parameter r we shall be free to choose in order to set a threshold between 'local' and 'global'.

After defining such 'local coverings' formally in Sections 1.4.1 and 1.4.2, we show in Section 1.4.3 that these coverings admit a natural hierarchy as their locality parameter r grows.

In Section 1.4.5 we show that these local covers G_r of finite graphs G always resemble Cayley graphs of finitely presented groups: every G_r has a graph-decomposition of finitely bounded width modelled on a Cayley graph of $\mathscr{D}(p)$, the group of deck transformations of p, which we show is finitely presented.

In Section 1.4.6 we translate the entire relationship between a graph and its local covers to groups: we show that when G is a locally finite Cayley graph of a finitely generated group Γ, then its local covers G_r are Cayley graphs of finitely presented groups that extend Γ.

1.4.1 Closed walks which stem from cycles

Let G be any connected graph, and let $x_0 \in G$ be any vertex. We say that a closed walk W in G based at x_0 *stems* from a closed walk Q in G if W can be written as $W = W_0 Q W_0^-$ where the *base walk* W_0 starts at x_0 and ends at the base vertex of Q. A closed walk *once around* a cycle O is a closed walk in O which traverses every edge of O exactly once. If W stems from a closed walk once around a cycle O in G, then W *stems* from O.

Let $\mathcal{O}$ be some fixed set of cycles in G. Given a vertex $x_0 \in G$, denote by $\pi_1^{\mathcal{O}}(G, x_0)$ the subgroup of $\pi_1(G, x_0)$ generated by the closed walks at x_0 that stem from a cycle in $\mathcal{O}$. For such a walk $W = W_0 Q W_0^-$ its conjugation $UWU^- = (UW_0)Q(UW_0)^-$ with any closed walk U based at x_0 stems from the same cycle as W, as only the base walk has changed. Thus, $\pi_1^{\mathcal{O}}(G, x_0)$ is normal in $\pi_1(G, x_0)$.

We say that a closed walk W in G is *generated* by $\mathcal{O}$ if the following equivalent

(i) there exists a vertex $x_0 \in G$ such that some walk in $\pi_1^{\mathcal{O}}(G, x_0)$ stems from W;

(ii) for every vertex $x_0 \in G$, every closed walk at x_0 in G which stems from W is in $\pi_1^{\mathcal{O}}(G, x_0)$.

In particular, the closed walks at any vertex x_0 that are generated by $\mathcal{O}$ are precisely those in $\pi_1^{\mathcal{O}}(G, x_0)$. Moreover, $S^{\mathcal{O}} := \pi_1^{\mathcal{O}}(G, x_0)$ is canonical as a subgroup of $\pi_1(G, x_0)$ if and only if $\mathcal{O}$ is invariant under the automorphisms φ of G: for every vertex $x \in G$ we have $\varphi_*(S_x^{\mathcal{O}}) = S_{\varphi(x)}^{\varphi(\mathcal{O})} = S_{\varphi(x)}^{\mathcal{O}}$.

While the implication (ii)→(i) is trivial, let us show (i)→(ii). To see (ii), let $x_0 \in G$ be any vertex and let $W_0 W W_0^-$ be any closed walk at x_0 that stems from W. By (i), there are a vertex $x_0' \in G$ and a closed walk $W' := W_0' W (W_0')^- \in \pi_1^{\mathcal{O}}(G, x_0')$ that stems from W.

As G is connected, we may assume that $x_0 = x_0'$. Indeed, let U be any walk in G from x_0 to x_0'. Then $UW'U^-$ clearly stems from W; let us show that it lies in $\pi_1^{\mathcal{O}}(G, x_0)$ and hence satisfies (i). Write $W' \in \pi_1^{\mathcal{O}}(G, x_0')$ in terms of generators, say $W' = W_1 \dots W_n$ where each W_i is a walk from x_0' to a cycle in $\mathcal{O}$, round it, and back. Then every UW_iU^- lies in $\pi_1^{\mathcal{O}}(G, x_0)$, and hence so does $UW_1U^- \dots UW_nU^- = UW'U^-$.

Thus, $x_0' = x_0$. Since $W_0 W W_0^-$ can be obtained from $W_0' W (W_0')^- \in \pi_1^{\mathcal{O}}(G, x_0)$ by conjugation with $W_0 (W_0')^-$, and since $\pi_1^{\mathcal{O}}(G, x_0)$ is normal, it follows that $W_0 W W_0^- \in \pi_1^{\mathcal{O}}(G, x_0)$.

We shall need the following well-known relationship between the fundamental group of a graph and the fundamental cycles of its spanning trees (see [Hat02] for a proof). Let T be a spanning tree of a graph G, rooted at a vertex x_0. Given an edge $e = uv$ not on T, there is a unique cycle in the graph $T + e$, the *fundamental cycle* of e with respect to T. The walk $x_0 T u v T x_0$ and its reverse $x_0 T v u T x_0$ stem from this cycle; they are the only two reduced closed walks at x_0 in $T + e$ that traverse e exactly once. For every such edge e we pick one of these two closed walks and call it the *fundamental closed walk* through e with respect to T.

Lemma 1.4.1. *Let G be a connected graph, and let T be a spanning tree of G rooted at a vertex x_0. Then the fundamental group $\pi_1(G, x_0)$ is freely generated by the fundamental closed walks with respect to T.* □

In particular, every closed walk in G is generated by the set of fundamental cycles with respect to any fixed spanning tree of G. If G is planar, its closed walks are also freely generated by walks around its inner faces [KT], which need not arise as the fundamental cycles with respect to any spanning tree. We shall not make use of this fact here.

1.4.2 Definition of local coverings

Let G be any connected graph, let $x_0 \in G$ be any vertex, and let $r \in \mathbb{N}$. The *r-local subgroup* of $\pi_1(G, x_0)$ is the group

$$\pi_1^r(G, x_0) := \pi_1^{\mathcal{O}_r}(G, x_0)$$

where $\mathcal{O}_r$ denotes the set of all cycles in G of length at most r. When the base point is given implicitly or does not matter, we refer to it simply as $\pi_1^r(G)$.

The *r-local covering* of G, denoted by $p_r \colon G_r \to G$, is the unique covering of G with characteristic subgroup $\pi_1^r(G)$. As $\pi_1^r(G)$ is normal and canonical (Section 1.4.1), so is p_r. We call G_r the *r-local covering graph* of G, or its *r-local cover* for short. In the context of r-local coverings, we may think of r as the *locality parameter*. We call a covering $p \colon C \to G$ *local* if p is equivalent to p_r for some $r \in \mathbb{N}$.

If the value of r is clear from the context, we refer to the cycles in $\mathcal{O}_r$ as the *short* cycles in G, and say that a closed walk in G is *generated by short cycles* if it is generated by $\mathcal{O}_r$. Note

that, by standard covering space theory, a walk in G_r is closed if and only if its projection to G is a closed walk in G generated by short cycles.

1.4.3 Local coverings are ball-preserving

All coverings, by definition, are local homeomorphisms. For r-local coverings of graphs we can extend these local homeomorphisms to larger neighbourhoods. How large exactly depends on r, and our next goal is to quantify this.

Let G be any connected graph, and $\varrho \in \mathbb{N}$. As usual, we write $d_G(x, y)$ for the distance in G between two vertices x and y. Given a vertex $v \in G$, let us write $B_G(v, \varrho/2)$ for the (combinatorial) $(\varrho/2)$-*ball* in G *around* v, the subgraph of G formed by the vertices at distance at most $\varrho/2$ from v and all the edges $xy \in G$ on these vertices that satisfy

$$d_G(v, x) + 1 + d_G(y, v) \leqslant \varrho.$$

Note that $B_G(v, \varrho/2)$ lacks all the edges xy where x and y have distance exactly $\varrho/2$ from v – which, of course, can occur only when ϱ is even. Instead, it consists of precisely those vertices and edges of G that lie on closed walks at v of length at most ϱ. In particular, $B_G(v, \varrho/2)$ contains all the cycles of length at most ϱ in G that meet v. Given a set X of vertices of G, we write $B_G(X, \varrho/2) := \bigcup_{x \in X} B_G(x, \varrho/2)$ for the $(\varrho/2)$-*ball* in G *around* X.

A covering $p\colon C \to G$ is said to *preserve* $(\varrho/2)$-*balls* if, for every lift $\hat{v}$ of a vertex $v \in G$, the covering p maps $B_C(\hat{v}, \varrho/2)$ isomorphically to $B_G(v, \varrho/2)$.[3] In particular, then, p projects all the cycles of length at most ϱ in C isomorphically to cycles of length at most ϱ in G. Note that every covering $C \to G$ preserves 1-balls.

The preservation of $(\varrho/2)$-balls can be reduced to local injectivity of the covering map:

Lemma 1.4.2. *The following two statements are equivalent for coverings $p\colon C \to G$ and $\varrho \in \mathbb{N}$:*

(i) *p preserves $(\varrho/2)$-balls;*

(ii) *for every vertex $v \in G$, the distance in C between two distinct lifts of v is greater than ϱ.*

Proof. For (i)→(ii), suppose for a contradiction that there exists a vertex $v \in G$ that has distinct lifts $\hat{v}_1$ and $\hat{v}_2$ at a distance of at most ϱ. Choose v, $\hat{v}_1$ and $\hat{v}_2$ so that $\hat{v}_1$ and $\hat{v}_2$ have minimum distance in C. Let P be a shortest $\hat{v}_1$–$\hat{v}_2$ path in C; its internal vertices then lie in distinct fibres of p. Hence $p(P) =: O$ is a cycle of length $d_C(\hat{v}_1, \hat{v}_2) \leqslant \varrho$, which thus lies in the $(\varrho/2)$-ball around v. By (i), the $(\varrho/2)$-ball around $\hat{v}_1$ contains a cycle Q such that $p(Q) = O$. Now Q and P, viewed as walks from $\hat{v}_1$, are distinct lifts of the same closed walk at v in G, which contradicts the uniqueness of path-lifting.

For (ii)→(i), consider a lift $\hat{v}$ of a vertex v of G. Recall that $B_G(v, \varrho/2)$ consists of precisely those vertices and edges of G that lie on closed walks at v of length at most ϱ, and similarly for $\hat{v}$ in C. Since closed walks of length at most ϱ at $\hat{v}$ project to closed walks of length at

[3] A notion that could be regarded as a precursor of the preservation of $(\varrho/2)$-balls has been studied by Georgakopoulos [Geo17].

most ϱ at v, our p maps $B_C(\hat{v}, \varrho/2)$ to $B_G(v, \varrho/2)$. It does so surjectively, since such walks at v lift to such walks at $\hat{v}$ by (ii). It does so injectively on the vertices of $B_C(\hat{v}, \varrho/2)$ by (ii). This implies that p is injective also on non-parallel edges of $B_C(\hat{v}, \varrho/2)$. But parallel edges cannot project to the same edge either, because p is 1–ball-preserving. □

Local coverings, by definition, preserve local structure as encoded in short cycles. Our next lemma says that they preserve local structure also in the more usual sense of copying small neighbourhoods:

Lemma 1.4.3. *For every $r \in \mathbb{N}$, the r-local covering $p_r \colon G_r \to G$ preserves $(r/2)$-balls.*

Proof. By Lemma 1.4.2 it is enough to show that no two lifts $\hat{v}_1, \hat{v}_2$ of a vertex $v \in G$ are connected in G_r by a path of length at most r. The projection of any such path P is a closed walk W at v of length at most r. By a straightforward direct argument, W is generated by closed walks at v stemming from short cycles, and thus lies in $\pi_1^r(G, v)$. Hence W lifts to a closed walk at $\hat{v}_1$ in G_r, by definition of p_r, which contradicts the fact that its unique such lift P is a path. □

Our r-local covering p_r is universal amongst all the $(r/2)$-ball-preserving coverings of G:

Lemma 1.4.4. *For every $(r/2)$-ball-preserving covering $p \colon C \to G$ there is a covering $q \colon G_r \to C$ such that $p_r = p \circ q$, and all such coverings q are equivalent.*

Proof. By the Galois correspondence of coverings we only have to show that the characteristic subgroup of p contains that of p_r. The latter is generated by the closed walks W in G that stem from a short cycle O, so it suffices to show that such W lift, via p, to closed walks in C.

Our W has the form $W_0 Q W_0^-$, where Q is a walk once around O. Pick a lift $\hat{W}$ of W to C. This includes, as subwalks, lifts $\hat{W}_0$, $\hat{Q}$ and $\hat{W}_0^-$ of W_0, Q and W_0^-. Since O is short and p preserves $(r/2)$-balls, $\hat{Q}$ is closed. Therefore $\hat{W}_0^-$ is the reverse of $\hat{W}_0$, and $\hat{W} = \hat{W}_0 \hat{Q} \hat{W}_0^-$ is a closed walk in C. □

Let us now see how the r-local coverings of G are related for various values of r. When $r \leqslant r'$, the $(r'/2)$-balls in G, which the r'-local covering $p_{r'} \colon G_{r'} \to G$ preserves by Lemma 1.4.3, include the $(r/2)$-balls in G. We may thus apply Lemma 1.4.4 with $p := p_{r'}$ to obtain a covering $q_{r,r'} \colon G_r \to G_{r'}$ such that $p_r = p_{r'} \circ q_{r,r'}$, as shown in Figure 1.4.

Moreover, $G_{r'}$ has its own r-local covering $p_{r,r'} \colon (G_{r'})_r \to G_{r'}$. We can compose this with $p_{r'}$ to obtain a covering $p_{r'} \circ p_{r,r'} \colon (G_{r'})_r \to G$. This composition of two universal local coverings is itself universal, as it is equivalent to the r-local covering p_r of G:

Lemma 1.4.5. *Given $r \leqslant r'$, the coverings $p_{r,r'}$ and $q_{r,r'}$ of $G_{r'}$ are equivalent, while p_r and $p_{r'} \circ p_{r,r'}$ are equivalent as coverings of G. In particular, $(G_{r'})_r$ and G_r are isomorphic graphs.*

Choosing $r = r'$ in Lemma 1.4.5 shows that taking r-local covers is idempotent: $(G_r)_r = G_r$.

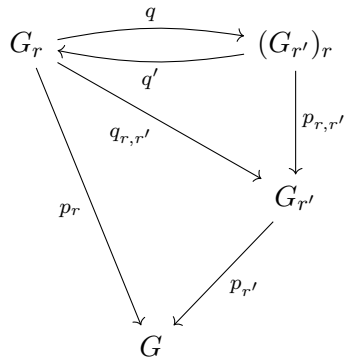

Figure 1.4: The interaction of local coverings for different values of r

Proof. Since both $p_{r'}$ and $p_{r,r'}$ preserve $(r/2)$-balls, so does their composition $p_{r'} \circ p_{r,r'}$. Choosing this as p in Lemma 1.4.4 yields a covering $q\colon G_r \to (G_{r'})_r$ such that $p_r = p_{r'} \circ p_{r,r'} \circ q$. Since $p_r = p_{r'} \circ q_{r,r'}$ and both p_r and $p_{r'}$ preserve $(r/2)$-balls, so does $q_{r,r'}$. So by Lemma 1.4.4 with $p := q_{r,r'}$, we obtain a covering $q'\colon (G_{r'})_r \to G_r$ with $p_{r,r'} = q_{r,r'} \circ q'$. Combining all the above, we have

$$p_r = p_{r'} \circ p_{r,r'} \circ q = p_{r'} \circ q_{r,r'} \circ q' \circ q = p_r \circ q' \circ q.$$

As a concatenation of two coverings, $q' \circ q$ is a covering (of G_r by itself), as of course is the identity id_{G_r}. Since $p_r \circ \mathrm{id}_{G_r} = p_r = p_r \circ (q' \circ q)$, as shown above, Lemma 1.4.4 applied with $p = p_r$ yields that id_{G_r} and $q' \circ q$ are equivalent coverings of G_r. In particular, $q' \circ q$ is a homeomorphism. As q is surjective, being a covering, this implies that q is a homeomorphism, which in turn implies that q' is a homeomorphism. Hence $p_{r,r'}$ and $q_{r,r'}$ are equivalent as coverings of $G_{r'}$ as witnessed by q', while p_r and $p_{r'} \circ p_{r,r'}$ are equivalent as coverings of G as witnessed by q. □

By and large, the connectivity of G_r increases as the locality parameter r grows and G_r changes from a tree, which is minimally connected, to the original graph G. How exactly the connectivity of G_r evolves along the hierarchy of local coverings given by Lemma 1.4.5 is studied in [CJKK25].

1.4.4 Cycle spaces

As is customary in graph theory, we call the first homology group of a graph its *cycle space*. Taken over the integers it is its *integral* cycle space; taken over $\mathbb{Z}/2\mathbb{Z}$ it is its *binary* cycle space.

To formalise this, let G be any connected graph and $r \in \mathbb{N}$. For every edge e of G pick an arbitrary fixed orientation $\vec{e}$; this makes G into a 1-complex. The *integral* or *binary cycle space* of G, then, is the R-module $\mathcal{Z}(G)$ for $R = \mathbb{Z}$ or $R = \mathbb{Z}/2\mathbb{Z}$ that consists of all $E(G) \to R$ functions with finite support whose sum of the values of the incoming edges at any given vertex equals the sum of the values of its outgoing edges.

Every closed walk $W = v_0 e_0 v_1 \dots v_{k-1} e_{k-1} v_k$ once around a cycle in G induces an element z_W of $\mathcal{Z}(G)$ by letting $z_W(e_i) := 1$ if $\vec{e}_i = (e_i, v_i, v_{i+1})$ and $z_W(e_i) := -1$ if $\vec{e}_i = (e_i, v_{i+1}, v_i)$

and $z_W(e_i) := 0$ elsewhere. Note that the reverse W^- of W induces $z_{W^-} = -z_W$. As is well known and easy to see, the (functions z_W of walks W once around) cycles generate all of $\mathcal{Z}(G)$. We write $\mathcal{Z}_r(G)$ for the submodule of $\mathcal{Z}(G)$ generated by the cycles of length at most r.

Lemma 1.4.6. *For the r-local cover G_r of G we have $\mathcal{Z}_r(G_r) = \mathcal{Z}(G_r)$.*

Proof. Let us show first that $\pi_1^r(G_r) = \pi_1(G_r)$. By definition of $q_{r,r}$ we have $p_r = p_r \circ q_{r,r}$. By Lemma 1.4.4 applied with $p = p_r$, it follows from $p_r \circ \mathrm{id}_{G_r} = p_r = p_r \circ q_{r,r}$ that id_{G_r} and $q_{r,r}$ are equivalent coverings of G_r. But $q_{r,r}$ is equivalent also to $p_{r,r}$, by Lemma 1.4.5. Hence $p_{r,r}$ and id_{G_r} are equivalent coverings of G_r, and thus have the same characteristic subgroups. These are $\pi_1^r(G_r)$ and $\pi_1(G_r)$, respectively.

Since the first homology group of a path-connected space is the abelianisation of its fundamental group, the fact that $\pi_1^r(G_r) = \pi_1(G_r)$ implies $\mathcal{Z}_r(G_r) = \mathcal{Z}(G_r)$. □

As pointed out for G_r in the last line of the above proof, we have $\mathcal{Z}(G) = \mathcal{Z}_r(G)$ if the fundamental group of G is generated by short cycles. When G is planar the converse holds too; this is shown in [KT].

1.4.5 Local coverings of finite graphs

While local coverings of finite graphs are often infinite, this is not necessarily the case. The following example due to Bowler [Bow21] exhibits finite graphs with n-sheeted coverings for any $n \in \mathbb{N}$.

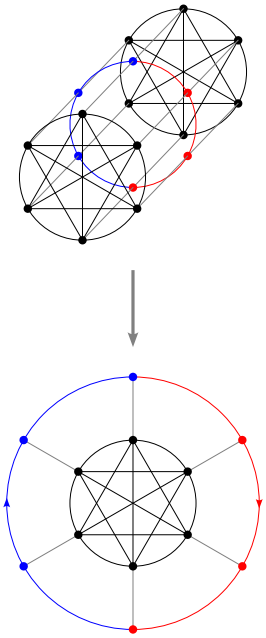

Figure 1.5: A graph whose 4-local covering is finite.

Example. Let G' be the graph depicted in Figure 1.5. We construct the desired graph G from G' in that we identify the two paths of length 3 drawn in red and blue which partition

the outer cycle of G', and we do this identification according to the indicated direction of the two paths. In other words, if we view G' embedded in the plane, symmetrically as drawn here, with the outer cycle equal to S^1, then we construct G from G' by taking the quotient of this outer cycle under the antipodal map $x \mapsto -x$ on S^1. Now taking two copies of G' and glueing them together by identifying their outer cycles yields a 2-sheeted covering of G, and one can check that this is indeed the 4-local covering of G by Lemma 1.4.4.

The above construction can easily be adapted to yield n-sheeted local coverings by partitioning the directed outer cycle of (a possibly larger version of) G' into n paths of the same length at least 3 which we then identify as above to obtain G. The respective 4-local coverings then arise by glueing n copies of this larger G' along their outer cycles. □

Our aim in this section is to prove that the local covers of finite graphs resemble Cayley graphs of finitely presented groups. We first show that every normal cover of a finite graph has a graph-decomposition, of finitely bounded width, modelled on a Cayley graph of its group of deck transformations. This group is always finitely generated. In the case of our local coverings we shall see that it is even finitely presented.

Our first result, the structure theorem for arbitrary normal covers of finite graphs, can be seen as a Švarc-Milnor lemma [Mil68; Šva55] that provides additional structural information:

Theorem 1.4.7. *Let $p\colon C \to G$ be a normal covering of a finite connected graph G. Then the group $\mathscr{D}$ of all deck transformations of p has a finite generating set S such that C has a graph-decomposition modelled on $D := \mathrm{Cay}(\mathscr{D}, S)$ with connected parts of finitely bounded size. If $\hat{x}_0 \in C$ is a lift of any vertex x of G then these parts C_φ, one for each $\varphi \in \mathscr{D}$, can be chosen as $C_\varphi = B_C(\varphi(\hat{x}_0), |G|)$.*

With these C_φ the graph-decomposition $(D, (C_\varphi)_{\varphi\in\mathscr{D}})$ of C is honest, and the map $f \mapsto g_f$ defined in Section 1.3.2 induces a bijection between the ends of C and those of $\mathscr{D}$. If $\mathscr{D}$ is an accessible group, then C is an accessible graph.

Proof. Choose $(C_\varphi)_{\varphi\in\mathscr{D}}$ as in the statement of the theorem. As these C_φ depend only on $\mathscr{D}$, not on the choice of generators for our Cayley graph $\mathrm{Cay}(\mathscr{D}, S)$, we can first verify (H1) and choose S later.

Since C is connected, every vertex is incident with some edge. To verify (H1) it is therefore enough to show that every edge of C is contained in some part C_φ. Every edge of C is a lift $\hat{e} = \hat{u}\hat{v}$ of an edge $e = uv$ of G. These have distance at most $|G| - 1$ from x in G. Pick a shortest u–x path P in G, lift it to a path $\hat{P}$ starting at $\hat{u}$, and let $\hat{x}_1$ be the last vertex of $\hat{P}$. As $\hat{P}$ has length at most $|G| - 1$, this shows that $\hat{e}$ lies in $B_C(\hat{x}_1, |G|)$. Since p is normal, there exists a deck transformation $\varphi \in \mathscr{D}$ such that $\varphi(\hat{x}_0) = \hat{x}_1$. Thus $\hat{e} \in C_\varphi$, completing the proof of (H1).

Let us now find a finite generating set S of $\mathscr{D}$ such that, for $D = \mathrm{Cay}(\mathscr{D}, S)$, the pair $(D, (C_\varphi)_{\varphi\in\mathscr{D}})$ will satisfy (H2) and thus be a graph-decomposition of C. We shall have to choose S so that, for every lift $\hat{v} \in C$ of a vertex v of G, the subgraph $D_{\hat{v}}$ of D induced by $\{\varphi \in \mathscr{D}\colon \hat{v} \in C_\varphi\}$ is connected.

To define S, fix of every $v \in G$ one lift $\hat{v}_0$. We shall pick enough generators for S to make every $D_{\hat{v}_0}$ connected, and then deduce the same for the other lifts $\hat{v}$ of each v. In fact, we shall add to S for every two nodes of $D_{\hat{v}_0}$ a generator that will make them adjacent in $\mathrm{Cay}(\mathscr{D}, S)$, and hence in $D_{\hat{v}_0}$.

The pairs of nodes of $D_{\hat{v}_0}$ are the pairs $(\varphi, \varphi') \in \mathscr{D}^2$ such that $\hat{v}_0 \in C_\varphi \cap C_{\varphi'}$. To give $D_{\hat{v}_0}$ an edge (φ, s) from φ to φ' we need to add $s = \varphi^{-1}\varphi'$ to S. Let S_v be the set of all these s, setting

$$S_v := \{\, \varphi^{-1}\varphi' \mid \hat{v}_0 \in C_\varphi \cap C_{\varphi'} \} \subseteq \mathscr{D},$$

and let $S := \bigcup_{v \in V(G)} S_v$. Then the subgraphs $D_{\hat{v}_0}$ of $D = \mathrm{Cay}(\mathscr{D}, S)$ are complete for all $v \in G$.

For a proof that each S_v, and hence S, is finite, it suffices to show that $\hat{v}_0 \in C_\varphi$ for only finitely many φ. For every such φ there is a path in C of length at most $|G|$ from $\hat{v}_0$ to $\varphi(x_0)$. But $B_C(\hat{v}_0, |G|)$ is finite, and the $\varphi(x_0)$ are distinct for different φ, by the uniqueness of path lifting. So there are only finitely many such φ, as required.

To prove (H2) for lifts $\hat{v}$ other than $\hat{v}_0$, recall from Section 1.2.3 that $\mathscr{D}$ acts on D by left-multiplication. Since our covering p is normal, there exists $\psi \in \mathscr{D}$ such that $\hat{v} = \psi(\hat{v}_0)$. Then $\psi \cdot D_{\hat{v}_0} = \{\, \psi\varphi \mid \varphi \in D_{\hat{v}_0} \}$, and any two of these nodes $\psi\varphi$ and $\psi\varphi'$ are joined by the edge $(\psi\varphi, s)$ for $s = \varphi^{-1}\varphi' \in S_v$. So it suffices to check that the node set of $\psi \cdot D_{\hat{v}_0}$ equals that of $D_{\psi(\hat{v}_0)} = D_{\hat{v}}$.

By definition, $D_{\psi(\hat{v}_0)}$ consists of those $\chi \in \mathscr{D}$ for which $\psi(\hat{v}_0) \in C_\chi$, i.e., for which $\psi(\hat{v}_0)$ has distance at most $|G|$ from $\chi(\hat{x}_0)$. With $\chi =: \psi\varphi$ the latter is equivalent to $\hat{v}_0$ having distance at most $|G|$ from $\varphi(\hat{x}_0)$, i.e. to $\varphi \in D_{\hat{v}_0}$. Hence

$$\chi \in D_{\psi(\hat{v}_0)} \Leftrightarrow \varphi \in D_{\hat{v}_0} \Leftrightarrow \psi\varphi \in \psi \cdot D_{\hat{v}_0} \Leftrightarrow \chi \in \psi \cdot D_{\hat{v}_0}$$

as desired. This completes the proof that $(D, (C_\varphi)_{\varphi \in \mathscr{D}})$ is a graph-decomposition.

For a proof that S generates $\mathscr{D}$, we have to show that $D = \mathrm{Cay}(\mathscr{D}, S)$ is connected. This follows from the connectedness of C by Lemma 1.3.4, since the parts C_φ of our graph-decomposition are non-empty.

Next, we check that our graph-decomposition $(D, (C_\varphi)_{\varphi \in \mathscr{D}})$ is honest. As the C_φ are non-empty, we only have to show that $C_\varphi \cap C_{\varphi'} \neq \emptyset$ for any adjacent nodes φ, φ' of D. As $D = \mathrm{Cay}(\mathscr{D}, S)$ there exists $s \in S$ such that $\varphi' = \varphi s$. This s lies in some S_v. This means, by definition of S_v, that there exist $\varphi_0, \varphi_0' \in \mathscr{D}$ such that $s = \varphi_0^{-1}\varphi_0'$ and $\hat{v}_0 \in C_{\varphi_0} \cap C_{\varphi_0'}$. Let $\psi := \varphi\varphi_0^{-1}$. Then $\varphi = \psi\varphi_0$ and $\varphi' = \varphi s = \psi\varphi_0 s = \psi\varphi_0'$. Hence $\hat{v}_0 \in C_{\varphi_0} \cap C_{\varphi_0'}$ implies that $\psi(\hat{v}_0) \in C_{\psi\varphi_0} \cap C_{\psi\varphi_0'} = C_\varphi \cap C_{\varphi'}$, since the action of $\mathscr{D}$ on itself by left-multiplication commutes with its natural action on C. In particular, $C_\varphi \cap C_{\varphi'} \neq \emptyset$, as required for our honesty proof.

We have shown that our graph-decomposition $(D, (C_\varphi)_{\varphi \in \mathscr{D}})$ is honest and has finite connected parts. In our proof that the sets S_v are finite we saw that each vertex of C lies in C_φ for only finitely many $\varphi \in \mathscr{D}$: that $(D, (C_\varphi)_{\varphi \in \mathscr{D}})$ has finite spread. By Lemma 1.3.5, therefore, the directions of C and of $D = \mathrm{Cay}(\mathscr{D}, S)$ are in bijective correspondence, and hence so are their ends [DK03, Theorem 2.2]. The ends of $\mathscr{D}$, by definition, are those of D.

Finally, since C covers the finite graph G it has finitely bounded degrees. The parts $C_\varphi = B_C(\varphi(\hat{x}_0), |G|)$ thus have finitely bounded size. If the group $\mathscr{D}$ is accessible then so is its Cayley graph D [TW93], and hence by Lemma 1.3.6 so is C. □

The groups of deck transformations of our local coverings of finite graphs are not only finitely generated but even finitely presented:

Proposition 1.4.8. *Let G be a finite graph, and let $r \in \mathbb{N}$. Then the group $\mathscr{D}_r$ of deck transformations of the r-local covering of G is finitely presented.*

Proof. Fix a vertex $x_0 \in G$. The fundamental group $\pi_1(G, x_0)$ of G is finitely presented by Lemma 1.4.1. To show that $\mathscr{D}_r = \pi_1(G, x_0)/\pi_1^r(G, x_0)$ is finitely presented, it remains to prove that $\pi_1^r(G, x_0)$ is the normal closure of a finite set of closed walks in G.

Recall that $\pi_1^r(G, x_0)$ is generated by the closed walks at x_0 stemming from short cycles in G. There are only finitely many such cycles, but infinitely many such walks. However, the walks stemming from the same short cycle O are all conjugates of each other in $\pi_1(G, x_0)$. Hence we may pick one of them for each O, finitely many in total, and their normal closure in $\pi_1(G, x_0)$ will be exactly $\pi_1^r(G, x_0)$. □

Corollary 1.4.9. *Local covers of finite graphs are accessible. They have graph-decompositions into connected parts of finitely bounded size modelled on Cayley graphs of finitely presented groups.*

Proof. Finitely presented groups are accessible [Dun85]. Apply Theorem 1.4.7 and Proposition 1.4.8. □

In contrast to local covers, arbitrary normal covers of finite graphs can be inaccessible:

Example. Let $\Gamma = \langle S \mid R \rangle$ be an inaccessible group finitely generated by S, and let G be any connected graph with $|V(G)| - 1 + |S|$ edges. Then $\pi_1(G)$ is isomorphic to the free group $F(S)$ (Lemma 1.4.1), so we may view the normal closure $N = \langle R \rangle_S^\triangleleft$ of R in S as a normal subgroup of $\pi_1(G)$.

Let $p\colon C \mapsto G$ be the normal covering of G with characteristic subgroup N. Then the group $\mathscr{D}$ of all deck transformations of p is isomorphic to $\pi_1(G)/N$ and hence to Γ. Thus, $\mathscr{D}$ is inaccessible, and hence so are all its locally finite Cayley graphs [TW93].

Now consider the graph-decomposition $(D, (C_\varphi)_{\varphi \in \mathscr{D}})$ of C modelled on a locally finite Cayley graph D of $\mathscr{D}$ as given by Theorem 1.4.7. This decomposition is honest, and its parts $C_\varphi = B_C(\varphi(\hat{x}_0), |G|)$ are connected. Hence by Lemma 1.3.3 there is a dual decomposition $(C, (D_v)_{v \in C})$ of D, again with connected parts D_v. These D_v have finitely bounded size by the proof of Theorem 1.4.7. Since D is inaccessible, Lemma 1.3.6 applied to its decomposition $(C, (D_v)_{v \in C})$ yields that C is inaccessible too. □

1.4.6 Local coverings of Cayley graphs

When our graph G is itself a Cayley graph, of some finitely generated (but not necessarily finite) group Γ, say, we can say even more about its local covers than we could in Section 1.4.5. These covers are now Cayley graphs too, of groups extending Γ:

Theorem 1.4.10. *Let G be a connected locally finite Cayley graph of some group Γ, and $r \in \mathbb{N}$. Then G_r is a connected locally finite Cayley graph of a finitely presented group Γ_r of which Γ is a quotient.*

For our proof of Theorem 1.4.10 we need a tool from combinatorial group theory.

Let G be a Cayley graph of a group Γ with respect to a finite generating set S. The fundamental group $\pi_1(G, x_0)$ of G is then isomorphic to the kernel of the canonical homomorphism $\varphi_{\Gamma,S}\colon F(S) \to \Gamma$, where $F(S)$ is the free group generated by S. Let us define such an isomorphism explicitly.

The *edge-labelling* $\ell\colon E(G) \to S$ of a Cayley graph $G = \mathrm{Cay}(\Gamma, S)$ is defined as $\ell(e) := s$ when $e = (g, s)$. We extend ℓ to a map from the walks in G to $F(S)$ by mapping a walk $W = g_0 e_0 g_1 \dots g_{k-1} e_{k-1} g_k$ to the group element $\ell(e_0)^{\epsilon_0} \cdots \ell(e_{k-1})^{\epsilon_{k-1}}$ of $F(S)$, where ϵ_i is 1 if $e_i = (g_i, s)$ for some $s \in S$ and -1 otherwise. Since two walks map to the same element of $F(S)$ under ℓ if their reductions coincide, ℓ defines a group homomorphism $\ell_*\colon \pi_1(G, x_0) \to F(S)$. This ℓ_* is our explicit isomorphism between $\pi_1(G, x_0)$ and $\ker(\varphi_{\Gamma,S}) \leqslant F(S)$:

Lemma 1.4.11 ([MKS04], Theorem 1.6)**.** *Let G be a Cayley graph of a group Γ with respect to a finite generating set S, and let ℓ be its edge-labelling. Fix any vertex x_0 of G. Then ℓ_* is an isomorphism from the fundamental group $\pi_1(G, x_0)$ to the kernel of the canonical homomorphism $\varphi_{\Gamma,S}\colon F(S) \to \Gamma$.*

Proof of Theorem 1.4.10. Let G be the Cayley graph of Γ with respect to the finite generating set S. Then Γ has the group presentation $\langle S \mid R \rangle$ with $R = \ker(\varphi_{\Gamma,S})$. We aim to show that G_r is the Cayley graph of a group Γ' with the same set S of generators as G, but with a different and finite set $R' \subseteq R$ of relators. We shall first construct Γ' and R', then check that $\mathrm{Cay}(\Gamma', S)$ is a covering graph of G, and finally show that this covering is equivalent to our r-local covering of G.

In $G = \mathrm{Cay}(\Gamma, S)$ let us fix the identity $1 \in \Gamma$ as our basepoint x_0. Since G is locally finite, there are only finitely many cycles of length at most r that contain x_0. Hence the set of closed walks at x_0 that only go once around such a cycle is finite too; let $R' \subseteq F(S)$ be its image under ℓ. By Lemma 1.4.11 we have $R' \subseteq \ell_*(\pi_1(G, x_0)) = \ker(\varphi_{\Gamma,S}) \subseteq F(S)$. As R' is finite, the group $\Gamma' := \langle S \mid R' \rangle$ is finitely presented.

Let us construct a covering $p'\colon G' \to G$ for $G' = \mathrm{Cay}(\Gamma', S)$. Since $R' \subseteq \ker(\varphi_{\Gamma,S})$, we have $\ker(\varphi_{\Gamma',S}) = \langle R' \rangle^{\lhd}_{F(S)} \subseteq \ker(\varphi_{\Gamma,S})$. Hence $\Gamma = F(S)/\ker(\varphi_{\Gamma,S})$ is a quotient of $\Gamma' = F(S)/\ker(\varphi_{\Gamma',S})$. The quotient map $p'\colon \Gamma' \to \Gamma$ between these groups, the vertex sets of their Cayley graphs G' and G, extends to a covering $p'\colon G' \to G$ by $(g', s) \mapsto (p'(g'), s)$ for edges, since both Cayley graphs use the same set S as edge labels. Explicitly, we have $\ell \circ p' = \ell'$ for the edge-labelling ℓ' of G'.

Let us show that p' is equivalent to our r-local covering $p_r\colon G_r \to G$. We do so by showing that their characteristic subgroups coincide. By Lemma 1.4.11, this will be the case if $\ell'_*(\pi_1(G', \hat{x}_0))$ and $\ell_*(\pi_1^r(G, x_0))$ are the same subgroup of $F(S)$, where $\hat{x}_0$ is any element in $p'^{-1}(x_0)$. Let R_r be the image under ℓ of the set of closed walks at x_0 which stem from a cycle in G of length at most r. Note that R_r is usually infinite, whereas $R' \subseteq R_r$ is finite. Again by Lemma 1.4.11 we have $\langle R' \rangle^{\lhd}_{F(S)} = \ell'_*(\pi_1(G', \hat{x}_0))$ and $\langle R_r \rangle^{\lhd}_{F(S)} = \ell_*(\pi_1^r(G, x_0))$. It thus remains to show that $\langle R_r \rangle^{\lhd}_{F(S)} = \langle R' \rangle^{\lhd}_{F(S)}$.

Since $R' \subseteq R_r$, we immediately have $\langle R' \rangle^{\lhd}_{F(S)} \subseteq \langle R_r \rangle^{\lhd}_{F(S)}$. For the converse inclusion let $t \in R_r$ be given. Then there exists a reduced closed walk W at x_0 which stems from a cycle O of length at most r such that $\ell(W) = t$. This W consists of a base walk W_0 and a closed walk Q once around O, so that $W = W_0 Q W_0^-$. Let $g \in V(G) = \Gamma$ be the base vertex of Q. Left-multiplication by g^{-1} defines an automorphism $\psi\colon h \mapsto g^{-1}h$ of G. Now $\psi(Q)$ is a closed walk, based at $\psi(g) = 1 = x_0$, once around the cycle $\psi(O)$ of length at most r. Since ψ is an automorphism of a Cayley graph, the labellings on the edges of O and $\psi(O)$ agree, so that $\ell(Q) = \ell(\psi(Q)) \in R'$. Thus, $t = \ell(W) = \ell(W_0)\ell(Q)\ell(W_0^-) = \ell(W_0)\ell(Q)\ell(W_0)^{-1}$ is contained in the normal subgroup $\langle R' \rangle^{\lhd}_{F(S)}$. This yields $R_r \subseteq \langle R' \rangle^{\lhd}_{F(S)}$ and completes the proof. □

Corollary 1.4.12. *Local covers of connected locally finite Cayley graphs are accessible.*

Proof. Finitely presented groups are accessible [Dun85], and hence so are their locally finite Cayley graphs [TW93]. Apply Theorem 1.4.10. □

1.5 Displaying global structure

In this section we make the statements of Theorems 1 and 2 precise and give an overview of their proof. This will include a proof of Theorem 3.

In Section 1.5.1 we recall how a tree of tangles can describe the tree-like aspects of the global structure of any graph. In Section 1.5.2 we apply this to the typically infinite graphs G_r that reflect the r-local structure of the finite graphs G given in Theorem 1, to describe their (tree-like) r-global structure. Projected back to G, this r-global structure of G_r ceases to be infinite and tree-like, but defines what we call the (finite) r-global structure of G.

Section 1.5.2 ends with a precise definition of the term of 'r-global structure' used in Theorems 1 and 2, and of what it means for another graph to 'display' that structure. We proceed from there to give an outline of the proof of Theorem 1. This includes a common generalisation, Theorem 1.5.5, of Theorems 1 and 2. The outline will serve as our road map when we prove Theorems 1 and 2 (by way of proving Theorem 1.5.5) in Sections 1.6–1.8.

1.5.1 Separations, tangles, and trees of tangles

In this section we recall basic facts and terminology about tangles; for a more detailed account see [Die25, Ch.12.5].

Let G be a graph. Given a set S of separations of G, we write $\vec{S}$ for the set of orientations of separations in S, two for every $s \in S$. An *orientation* O of S contains for every separation in S exactly one of its two orientations. Recall that $(A, B) \leqslant (C, D)$ for two oriented separations (A, B) and (C, D) of G if $A \subseteq C$ and $B \supseteq D$. If O does not contain an oriented separation (B, A) of G whenever $(A, B) \leqslant (C, D) \in O$, then O is *consistent*. A set N of separations of G is *nested* if every two of its elements, $\{A, B\}$ and $\{C, D\}$ say, have orientations $(A, B) \leqslant (C, D)$. Separations that are not nested are said to *cross*.

Let Γ be a group acting on a graph G. A set S of separations of G is Γ-*canonical* if for every $g \in \Gamma$ and every $\{A, B\} \in S$ we have $g \cdot \{A, B\} := \{g \cdot A, g \cdot B\} \in S$. If S is $\mathrm{Aut}(G)$-canonical, we call it *canonical*.

Let us now define tangles. Given any $k \in \mathbb{N} \cup \{\aleph_0\}$, let S_k denote the set of all separations of G of order $< k$. A k-*tangle* in G is an orientation τ of S_k such that for all $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$, not necessarily distinct, we have $G[A_1] \cup G[A_2] \cup G[A_3] \neq G$. A *tangle* is a k-tangle for some k. Note that every tangle is consistent.

Examples of tangles include blocks and ends. A k-*block* in G is a maximal set of at least k vertices that is included in one of the two sides of each separation in S_k. It *orients* that separation *towards* this side: as (A, B) if it lies in B. Every k-block X thus defines an orientation τ_X of S_k. If $|X| > 3(k-1)/2$, then τ_X is a tangle [CDHH14, Theorem 6.1]. By the maximality in the definition of a k-block, distinct k-blocks define different tangles.

Given an end ω of G, let f_ω be the corresponding direction of G (see Section 1.3.2). Orienting every separation $\{A, B\} \in S_{\aleph_0}$ towards the side that contains $f_\omega(A \cap B)$, we obtain a tangle τ_ω of $S_{\aleph_0}$ [Die17]. As in the case of blocks, distinct ends define different tangles.

A separation of G *distinguishes* two tangles in G if they orient it differently. The separation distinguishes these two tangles *efficiently* if no separation of smaller order distinguishes them. Two tangles are *distinguishable* if some separation distinguishes them; note that any two k-tangles for the same k are distinguishable. A set N of separations *distinguishes* a set of tangles (efficiently) if every two distinguishable tangles in this set are (efficiently) distinguished by some separation in N. A *tree of tangles* of G, in this chapter, is a nested set of separations of G which efficiently distinguishes the entire set of tangles in G.

A separation of G is *relevant* if it efficiently distinguishes some two tangles τ and τ' in G. One can show that such separations $\{A, B\}$ are *tight*: for $X := A \cap B$ the graph $G - X$ has components $C \subseteq G[A]$ and $D \subseteq G[B]$ such that every vertex in X sends edges to both C and D. Indeed, these are the unique such components for which $(V(G - C), V(C) \cup X) \in \tau$ and $(V(G - D), V(D) \cup X) \in \tau'$ or vice versa [Die25, Ex. 43(iii)].

Note that a finite-order separation of G distinguishes the tangles induced by two ends of G if and only if it distinguishes these ends in the usual sense that they contain rays on different sides of that separation. It distinguishes these two ends *efficiently* if no separation of lower order distinguishes them.

Let us remember all the above properties of trees of tangles for later use in Section 1.8:

Lemma 1.5.1. *Let N be a tree of tangles of G in which every separation is relevant. Then N*

efficiently distinguishes every two ends of G. Every separation in N has finite order and is tight. □

1.5.2 The r-global structure of G

In general, a locally finite graph G may have several canonical trees of tangles. However, there is one specific canonical tree of tangles $N(G)$ that is particularly natural, which we therefore fix as an invariant of those locally finite graphs in which it exists. In this section we describe the construction of $N(G)$, and then show in which sense it is particularly natural.

Recall [HLMR22; TW93] that a graph G is *accessible* if there exists a number $K \in \mathbb{N}$ such that every two ends of G are distinguished by a separation of order at most K. Extending this notion from ends to tangles, we call a graph G *tangle-accessible* if there exists a number $K \in \mathbb{N}$ such that every two distinguishable tangles in G are distinguished by a separation of order at most K.

Let G be a connected, locally finite, and tangle-accessible graph. Denote by D the set of all finite-order separations of G that distinguish some two tangles in G efficiently. Since G is tangle-accessible, all the separations in D have order at most some fixed $K \in \mathbb{N}$, and hence they each cross only finitely many other separations in D [EKT22, proof of Proposition 6.2].

Definition 1.5.2. For every pair τ, τ' of distinguishable tangles in G let $N_{\tau,\tau'}$ be the set of all separations that efficiently distinguish τ from τ' and, among these, cross as few separations in D as possible. Let $N(G)$ denote the union of all these sets $N_{\tau,\tau'}$.

Note that the separations in $N_{\tau,\tau'}$ in Definition 1.5.2 are relevant. The set $N(G)$ is clearly canonical, it distinguishes all the tangles in G, and its elements have finitely bounded order. It is also still nested:

Lemma 1.5.3 ([CK24], Theorem 4.2)**.** *For every connected, locally finite, and tangle-accessible graph G, the set $N(G)$ of separations defined in Definition 1.5.2 is a canonical tree of tangles. All the separations in $N(G)$ are relevant; in particular, they have finitely bounded order.*

We sometimes refer to $N(G)$ informally as *the* canonical tree of tangles of G.

We shall prove in Section 1.7 that the r-local covers G_r of the graphs G we consider in Theorems 1 and 2 satisfy the premise of Lemma 1.5.3. We shall see that they have unique canonical tree-decompositions $\mathcal{T} = (T, \mathcal{V})$ such that $\alpha_{\mathcal{T}} \colon E(T) \to N(G_r)$ is a bijection. We say that $N(G_r)$ *induces* this tree-decomposition $\mathcal{T}$ of G_r.

To see in which sense $N(G)$ is particularly natural, consider any tree of tangles N in G. By definition, N has to contain for every pair τ, τ' of tangles a separation $s(\tau, \tau')$ that distinguishes them efficiently. If we want N to be canonical, however, we cannot simply pick for N just any such separation $s(\tau, \tau')$ for each pair τ, τ' of tangles: our choice has to be definable in terms of invariants of G. This may force us to pick not just one such separation $s(\tau, \tau')$ for each pair τ, τ', but a set $N(\tau, \tau')$ of more than one separation. As $N(\tau, \tau')$ is to be included in N, it will be a nested subset of the set $D(\tau, \tau')$ of separations that distinguish τ from τ' efficiently.

Our set $N(G)$ now chooses these $N(\tau, \tau')$ with reference to $D(\tau, \tau')$ alone: they consist of *all* the elements of $D(\tau, \tau')$ that are the most promising candidates for inclusion in $N(\tau, \tau')$, in that they cross as few other elements of $D(\tau, \tau')$ as possible.

As described in Section 1.4, the r-local cover G_r of a graph G agrees with G on its r-local structure, the structure encoded in the r-local subgroup $\pi_1^r(G)$ of $\pi_1(G)$. The remaining structure of G is unfolded in G_r in a tree-like way, which encodes it in conjunction with the group $\mathscr{D}_r = \pi_1(G)/\pi_1^r(G)$ of deck transformations of G_r over G. This tree-like global structure of G_r is captured by its tree of tangles $N(G_r)$. We refer to $N(G_r)$ and $\mathscr{D}_r$ together as the r-global structure of G:

Definition 1.5.4. The *r-global structure* of a graph G whose r-local cover G_r is tangle-accessible is the pair $(N(G_r), \mathscr{D}_r)$, where $\mathscr{D}_r$ is the group of deck transformations of $p_r \colon G_r \to G$ and $N(G_r)$ is our canonical tree of tangles of G_r from Definition 1.5.2.

The *r-global structure* of a group Γ presented with a finite set S of generators is the r-global structure of its Cayley graph $\mathrm{Cay}(\Gamma, S)$.

The unique canonical graph-decomposition of G that *displays* its r-global structure is that defined via p_r by the canonical tree-decomposition which $N(G_r)$ induces on G_r, as in Construction 1.3.8. □

Definition 1.5.4 rests on the assumption that G_r is tangle-accessible. We shall prove that this is the case whenever G is finite, as in Theorem 1, or a Cayley graph as in Theorem 2.

Note that when a graph-decomposition $(H, \mathcal{G})$ displays the r-global structure of a graph G, then H comes with an edge labelling that reveals additional internal structure of H, which in turn displays additional connectivity aspects of G not reflected in the graph structure of H itself.

Indeed, recall that every edge e of the tree T from the tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ that defines $(H, \mathcal{G})$ as in Construction 1.3.8 maps to a separation $\alpha(e)$ of G_r. This separation has some order k_e. When one analyses trees of tangles in graph minor theory, one often contracts in T all those edges e whose label k_e is at least some chosen threshold k. This results in a tree T_k that is a minor of T, and whose edges e map to separations $\alpha(e)$ of order $< k$. These still distinguish all the k-tangles of the graph decomposed, in our case G_r. These trees T_k form a chain of minors of T as k grows: $T_1 \preccurlyeq T_2 \preccurlyeq \ldots \preccurlyeq T$.

This structure is now inherited by $H = T/\mathscr{D}$. Its edges arise from ($\mathscr{D}$-orbits of) edges e of T, which carry labels k_e that specify the orders of separations $\alpha(e)$ of G_r. We may interpret these as local separations of G: they tell us how the parts G_h of $\mathcal{G}$ are separated locally from each other in G. (In our original example from the Introduction, these local separators are the 2-sets in which two adjacent K^5s overlap.) And as with those tree minors T_k of T, contracting in H the edges with labels at least some chosen threshold k gives us a minor H_k of H. We thus obtain a hierarchy $H_1 \preccurlyeq H_2 \ldots$ of minors of H that defines an internal structure of H, which in turn displays more details of the connectivity structure of G than H does alone. This internal structure is explored further in [CJKK25].

It is sometimes desirable that the parts G_h of the graph-decomposition that displays the r-global structure of G should inherit more properties of the parts $G_r[V_t]$ of $\mathcal{T}$ whose projections they are than the projection map p_r can transmit. Such additional information is encoded in what we call a *blueprint* of G: a family whose members are not themselves subgraphs of G but which map homomorphically to the parts G_h. We explore the notion of blueprints and the properties of the parts of $\mathcal{T}$ that they preserve in [BJKKP]. This yields, among other things, a new proof of Carmesin's local 2-separator theorem [Car22] via local coverings.

While the r-global structure of a finite graph G is displayed by a graph-decomposition $(H, \mathcal{G})$ of G where H is another finite graph, this decomposition of G is obtained indirectly via its usually infinite r-local covering. One can ask, then, whether it is possible to characterise $(H, \mathcal{G})$ directly in terms of G. As a proof of concept for this type of problem, we obtain in [AJKK25] such a direct characterisation of the r-global structure of the locally finite graphs that are 'r-locally chordal', those whose parts G_h as above are complete.

1.5.3 Road map for the proof of main result

The remainder of this chapter is devoted to the proof of Theorem 1, for which we now give an outline. We shall prove Theorem 1 by way of proving a stronger result, Theorem 1.5.5, which will include Theorem 2 at no extra cost. A proof of Theorem 3 will also be included on the way.

As we will see in Section 1.7.1, every nested set of separations of a graph, all of finitely bounded order, induces a tree-decomposition of that graph. As seen in Construction 1.3.8 and Theorem 1.3.13, canonical tree-decompositions of local covers define canonical graph-decompositions of the graph they cover. So once we know that our canonical set $N(G_r)$ from Definition 1.5.2 is defined, for which we only need that G_r is tangle-accessible (Lemma 1.5.3), it will induce the tree-decomposition of G_r that defines the canonical graph-decomposition of G which displays its r-global structure.

Let us summarise these observations as the point of departure for the proofs of all our main results:

Reduction 1. *It suffices to show that G_r is tangle-accessible.*

In order to show that G_r is tangle-accessible we study another canonical tree of tangles of G_r, one whose existence was proved by Elbracht, Kneip and Teegen [EKT22] (Lemma 1.8.4). In contrast to $N(G_r)$ this tree of tangles $N'(G_r)$ is defined for all locally finite graphs, without any accessibility requirements. Correspondingly, the separations in $N'(G_r)$ need not, prima facie, have finitely bounded order. However if they do, then G_r is tangle-accessible. Hence:

Reduction 2. *It suffices to show that the separations in $N'(G_r)$ have finitely bounded order.*

Instead of trying to bound the order of the separations in $N'(G_r)$ directly, we will prove that $N'(G_r)$ exhibits a certain structural property: the property of being *exhaustive*. This means that for every infinite sequence $(A_0, B_0) < (A_1, B_1) < \dots$ in $\overrightarrow{N'}(G_r)$ we have $\bigcap_{i\in\mathbb{N}} B_i = \emptyset$. Sets of tight separations of finitely bounded order in locally finite graphs are easily seen to be

exhaustive. The converse fails in general. However we will show that if $N'(G_r)$ is exhaustive, then $\mathscr{D}_r$ acts on $N'(G_r)$ with only finitely many orbits. In particular, then, the separations in $N'(G_r)$ will have finitely bounded order. Thus:

Reduction 3. *It suffices to show that $N'(G_r)$ is exhaustive.*

We will prove that $N'(G_r)$ is exhaustive in two steps, as follows.

In the first step we show that G_r is accessible in the traditional sense: that the separations in $N'(G_r)$ that efficiently distinguish two ends of G_r have finitely bounded order. We denote this set as $N'_{\text{end}} \subseteq N'(G_r)$. The fact that the separations in N'_{end} have finitely bounded order, and hence are exhaustive because they are clearly tight, will be the key ingredient of the remainder of our proof.

In the second step we show that $N'(G_r)$ is exhaustive. For this let $(A_0, B_0) < (A_1, B_1) < \dots$ be any infinite sequence in $\overrightarrow{N'}(G_r)$; we have to show $\bigcap_{i\in\mathbb{N}} B_i = \emptyset$. Since N'_{end} is exhaustive, it suffices to find in $\overrightarrow{N'}_{\text{end}}$ an infinite sequence $(C_0, D_0) < (C_1, D_1) < \dots$ which our earlier sequence dominates: one such that for every $i \in \mathbb{N}$ there exists $j \in \mathbb{N}$ with $(C_i, D_i) \leqslant (A_j, B_j)$. If there is no such sequence we find a part V'_t of the tree-decomposition of G_r induced by N'_{end} such that $A_j \cap B_j \subseteq V'_t$ for all large enough j. We then enlarge V'_t to a certain set V_t such that, for every j_0 large enough, the subsequence $(A_j, B_j)_{j\geqslant j_0}$ induces an infinite sequence of tight separations of $G[V_t]$. Moreover, $G[V_t]$ will have the following properties.

The graph $G[V_t]$ will be connected; it will have exactly one end; and the stabiliser of t in $\mathscr{D}_r$, the group of the deck transformations of G_r, will act quasi-transitively on $G[V_t]$. These properties combined will imply that $G[V_t]$ has no infinite sequence of separations such as that induced by our $(A_j, B_j)_{j\geqslant j_0}$, a contradiction. Hence $N'(G_r)$ will be exhaustive, and our proof will be complete.

Following this road map we will be able to prove Theorems 1 and 2 together, by proving their following common strengthening:

Theorem 1.5.5. *Let G be a connected, locally finite, and quasi-transitive graph, and let $r > 0$ be an integer. Then G has a unique canonical graph-decomposition of finite spread which displays its r-global structure.*

Note that finite graphs, such as the G in Theorem 1, are trivially quasi-transitive, and so are locally finite Cayley graphs such as the $\mathrm{Cay}(\Gamma, S)$ in Theorem 2.

Our proof of Theorem 1.5.5 will implement the strategy outlined above by reversing the order of its steps, as follows. We start in Section 1.6 with a proof that if G is locally finite and quasi-transitive, then G_r is accessible. We deduce from this in Section 1.7 that $N'(G_r)$ is exhaustive (cf. Reduction 3). In Section 1.8 we infer that the separations in $N'(G_r)$ have finitely bounded order (cf. Reduction 2). This implies that G_r is tangle-accessible (cf. Reduction 1), which will allow us to prove Theorem 1.5.5.

1.6 Local covers are accessible

Our aim in this section is to establish Theorem 1.6.2 below. Recall the definition of accessibility in graphs:

Definition 1.6.1 ([TW93])**.** A graph G is *accessible* if there exists an integer K such that every two ends of G are distinguished by a separation of order at most K.

Theorem 1.6.2. *Let G be a connected, locally finite, quasi-transitive graph, and let $r > 0$ be an integer. Then the r-local cover G_r of G is accessible.*

We already proved Theorem 1.6.2 for the special cases of Cayley graphs (Corollary 1.4.12), and of arbitrary finite graphs by way of graph-decompositions (Corollary 1.4.9). These are the two cases that we would need for direct proofs of Theorems 1 and 2, proceeding exactly as we shall below. However, Theorem 1.6.2 is interesting in its own right, is considerably more general, and enables us to also prove Theorem 1.5.5.

For a comprehensive proof of Theorem 1.6.2, let us observe first that G_r is connected, locally finite, and quasi-transitive as soon as G is:

Lemma 1.6.3. *Let G be a connected, locally finite and quasi-transitive graph, and let $r > 0$ be an integer. Then the r-local cover G_r of G is connected, locally finite, and quasi-transitive.*

Proof. The r-local cover G_r is connected and locally finite by definition. To prove that G_r is quasi-transitive, let us show that lifts $\hat{u}$ of u and $\hat{v}$ of v lie in the same orbit of $\mathrm{Aut}(G_r)$ as soon as $v = \varphi(u)$ for some $\varphi \in \mathrm{Aut}(G)$. By Lemma 1.3.11, and since local coverings are canonical, φ lifts to an automorphism $\hat{\varphi}$ of G_r. Then $\hat{v}$ differs from $\hat{\varphi}(\hat{u})$ only by a deck transformation, so $\hat{u}$ and $\hat{v}$ lie in the same orbit of $\mathrm{Aut}(G_r)$. □

Our key ingredient for the proof of Theorem 1.6.2 is the following result of Hamann:

Lemma 1.6.4 ([Ham18], Corollary 3.3)**.** *Every locally finite quasi-transitive graph whose binary cycle space is generated by cycles of bounded length is accessible.*

Proof of Theorem 1.6.2. By Lemma 1.4.6, the binary cycle space of G_r is generated by its cycles of length at most r. The theorem now follows from Lemma 1.6.4, which we may apply to G_r by Lemma 1.6.3. □

1.7 Tools and lemmas for the proof of Theorem 1

Throughout this section, G will be a connected locally finite infinite graph. Our main goal is to prove the following key lemma. Note that if G satisfies its premise, it must be accessible.

Lemma 1.7.1. *Let N be a nested set of tight finite-order separations of G, and suppose that N is Γ-canonical for some group Γ that acts quasi-transitively on G. Suppose further that there exists an integer K such that some Γ-canonical set $N_{\mathrm{end}} \subseteq N$ of separations of order $< K$ distinguishes all the ends of G. Then N is exhaustive and induces a Γ-canonical tree-decomposition of G.*

1.7.1 Tools: tree-decompositions and nested sets of separations

Let $\mathcal{T} = (T, \mathcal{V})$ be a tree-decomposition of G, and let Γ be a group acting on G. Recall from Section 1.3.2 that every edge t_1t_2 of T induces a separation $\{A_{t_1}, A_{t_2}\}$ of G, and that each of its two orientations (t_1, t_2) induces the oriented separation $\alpha_{\mathcal{T}}(t_1, t_2) = (A_{t_1}, A_{t_2})$ of G. If we orient distinct edges t_1t_2 and t_2t_3 of T each from t_i to t_{i+1}, then $\alpha_{\mathcal{T}}(t_1, t_2) \leqslant \alpha_{\mathcal{T}}(t_2, t_3)$ in the partial ordering $\leqslant$ of oriented separations (see Section 1.5.1).

By definition of the sets A_{t_1} and A_{t_2}, the separations of G that $\mathcal{T}$ induces in this way are pairwise nested. If $\mathcal{T}$ is regular, they are proper separations. If $\mathcal{T}$ is Γ-canonical for some group Γ as defined in Section 1.3.3, then the nested set N of separations of G which $\mathcal{T}$ induces is Γ-canonical as defined in Section 1.5.1.

Conversely, let us try to construct from a given nested set N of proper separations of G a regular tree-decomposition of G that induces N. We begin by constructing, from N alone, the nodes of a potential decomposition tree. A *splitting star* of N is any set $\sigma \subseteq \overrightarrow{N}$ that consists of all the maximal separations of some consistent orientation O of N such that for every $\vec{r} \in O$ there exists $\vec{s} \in \sigma$ with $\vec{r} \leqslant \vec{s}$. Every separation in $\overrightarrow{N}$ lies in at most one splitting star of N. If N is induced by a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G, as is our aim, then every $\vec{s} \in \overrightarrow{N}$ does lie in a splitting star: it will have the form $\vec{s} = \alpha_{\mathcal{T}}(t', t)$ for some oriented edge (t', t) of T, and $\vec{s}$ will lie in splitting star of N formed by the separations $\alpha_{\mathcal{T}}(t'', t)$ where t'' varies over the neighbours of t in T.[4]

Taking our cue from the above observation, let us define a graph T from N as follows. Take the splitting stars of N as the nodes of T, and make two nodes adjacent if they contain, as splitting stars of N, distinct orientations of some element of N. For each splitting star $t \in V(T)$ let $V_t := \bigcap_{(A,B)\in t} B$, and put $\mathcal{V} = (V_t)_{t\in T}$. Let us say that N *induces* this pair $(T, \mathcal{V})$.

As indicated, $(T, \mathcal{V})$ need not be a tree-decomposition of G when N is infinite. Indeed, while T is always acyclic it need not be connected, and the family $\mathcal{V}$ might violate (T1) [EKT22, Example 4.9]. We shall mention some sufficient conditions for $(T, \mathcal{V})$ to be a tree-decomposition later.

If $\mathcal{T} := (T, \mathcal{V})$ is a tree-decomposition, it will induce the set N of separations we started with, in that $N = \{\, \alpha_{\mathcal{T}}(e) : e \in E(T) \,\}$ (as was our aim), with $\alpha_{\mathcal{T}}$ an order isomorphism between the oriented edges of T and $\overrightarrow{N}$ [EKT22, Proof of Lemma 2.7]. If N is Γ-canonical for some group Γ, then $\mathcal{T}$ will be a Γ-canonical tree-decomposition. Indeed, if N is Γ-canonical, then every $g \in \Gamma$ maps splitting stars of N to splitting stars of N. Since $g \cdot (A, B) = (g \cdot A\,,\, g \cdot B)$ for $g \in \Gamma$ and $(A, B) \in \overrightarrow{N}$, this action of Γ on T agrees with that in our definition of Γ-canonicity of $(T, \mathcal{V})$ in Section 1.3.3.

A strictly increasing infinite sequence $(A_i, B_i)_{i\in\mathbb{N}}$ of oriented separations of G is *exhaustive* if $\bigcap_{i\in\mathbb{N}} B_i = \emptyset$. A set S of separations is *exhaustive* if every such sequence in $\vec{S}$ is exhaustive. Any nested set of separations induced by a tree-decomposition of G of finite spread is clearly exhaustive. We shall need the following converse:

[4] In general, when N is infinite, the elements of $\overrightarrow{N}$ need not lie in a splitting star of N; see [Die18a; GK21].

Lemma 1.7.2 ([EKT22], Lemma 2.7)**.** *Let N be a nested set of proper separations of G. If N is exhaustive, then the pair $(T, \mathcal{V})$ induced by N is a regular tree-decomposition of G. If a group Γ acts on G so that N is Γ-canonical, then this tree-decomposition $(T, \mathcal{V})$ is Γ-canonical.*[5]

The task of constructing tree-decompositions thus reduces to that of constructing exhaustive nested sets of separations. Let us collect a few lemmas that will help with this task. When $(A_i, B_i)_{i\in\mathbb{N}}$ and $(C_j, D_j)_{j\in\mathbb{N}}$ are sequences of separations such that for every i there exists j such that $(A_i, B_i) \leqslant (C_j, D_j)$, we say that $(C_j, D_j)_{j\in\mathbb{N}}$ *dominates* $(A_i, B_i)_{i\in\mathbb{N}}$.

Lemma 1.7.3. *If $(A_i, B_i)_{i\in\mathbb{N}}$ and $(C_j, D_j)_{j\in\mathbb{N}}$ are sequences of separations such that $(C_j, D_j)_{j\in\mathbb{N}}$ dominates $(A_i, B_i)_{i\in\mathbb{N}}$, then $\bigcap_{i\in\mathbb{N}} B_i \supseteq \bigcap_{j\in\mathbb{N}} D_j$. In particular, if $(A_i, B_i)_{i\in\mathbb{N}}$ is exhaustive then so is $(C_j, D_j)_{j\in\mathbb{N}}$.* □

All the separations from which we shall construct tree-decompositions will be tight, and hence proper. Lemma 1.7.5 below, which is implicit in [EKT22], shows that a nested set of tight separations is exhaustive if these have finitely bounded order. For its proof we need the following well-known fact about tight separations:

Lemma 1.7.4 ([TW93], Proposition 4.2)**.** *Every vertex of G lies in only finitely many separators of tight separations of order less than K, for each $K \in \mathbb{N}$.*

Lemma 1.7.5 ([EKT22])**.** *Let K be an integer, and let $(A_i, B_i)_{i\in\mathbb{N}}$ be a strictly increasing sequence of tight separations of G, all of order $< K$. Then $(A_i, B_i)_{i\in\mathbb{N}}$ is exhaustive.*

Proof. By Lemma 1.7.4, every vertex of G lies in $A_i \cap B_i$ for only finitely many $i \in \mathbb{N}$. These sets $A_i \cap B_i$ are finite, so applying this to every vertex $v \in A_i \cap B_i$ shows that $(A_i \cap B_i) \cap (A_j \cap B_j) = \emptyset$ for all large enough j. We can thus find an infinite subsequence of $(A_i, B_i)_{i\in\mathbb{N}}$ whose separators are disjoint. But this implies that $\bigcap_i B_i = \emptyset$: otherwise G contains an $A_0 - \bigcap_i B_i$ path that meets all these infinitely many disjoint separators, a contradiction. Hence $(\vec{s}_i)_{i\in\mathbb{N}}$ is exhaustive. □

1.7.2 Proof of Lemma 1.7.1 for 1-ended graphs

Let us now prove Lemma 1.7.1 for 1-ended graphs. In fact, we show the following stronger result:

Lemma 1.7.6. *Assume that G has only one end. Let N be a nested set of tight separations of G, all of finite order. Suppose that N is Γ-canonical for some group Γ acting quasi-transitively on G. Then $\vec{N}$ contains no strictly increasing infinite sequence.*

We shall need the following fact about 1-ended, connected, locally finite graphs:

Lemma 1.7.7. *If G has only one end, then every finite-order separation of G has a finite and an infinite side.*

[5] The last sentence, which we already verified, is not explicitly stated in [EKT22].

Proof. Let $\{A, B\}$ be any finite-order separation of G. At least one of A, B must be infinite. Since $A \cap B$ is finite, and G is locally finite and connected, $G - A$ and $G - B$ each have only finitely many components. If they are both infinite, they each contain an infinite component. By König's infinity lemma [Die25, Proposition 8.2.1], each of these components contains a ray. These rays, then, are separated by the finite set $A \cap B$, which contradicts our assumption that G has only one end. □

Our next lemma shows that, in a strictly increasing infinite sequence of separations, every separation is oriented towards its infinite side:

Lemma 1.7.8. *In any strictly increasing sequence $(A_i, B_i)_{i\in\mathbb{N}}$ of separations of G, all B_i are infinite.*

Proof. Suppose some B_i is finite; we may assume that $i = 0$. Since $B_0 \supseteq B_i$ for all i, all the B_i are finite. Since $(A_i, B_i)_{i\in\mathbb{N}}$ is a strictly increasing sequence, we have for each i either $A_i \subsetneq A_{i+1}$ or $B_i \supsetneq B_{i+1}$. As B_0 is finite, we have $B_i \supsetneq B_{i+1}$ for only finitely many i, so $A_i \subsetneq A_{i+1}$ for infinitely many i. For all these i the non-empty sets $A_{i+1} \setminus A_i$ are disjoint and contained in the finite set B_0, a contradiction. □

Recall that the set N in Lemma 1.7.6, which we are seeking to analyse, is Γ-canonical. The following lemma shows how separations with a finite side interact with their images under certain automorphisms:

Lemma 1.7.9. *Let $\{A, B\}$ be a separation of G with A finite. Consider any automorphism φ of G and a vertex $v \in A \cap B$ with a neighbour $w \in B \setminus A$ such that $\varphi(v) \in A \setminus B$. Then $\{A, B\}$ and $\{\varphi(A), \varphi(B)\}$ cross.*

Proof. Let us show first that neither $A \subseteq \varphi(A)$ nor $\varphi(A) \subseteq A$. To see this, it is enough to show $\varphi(A) \neq A$, since A is finite. If $\varphi(A) = A$, then $\varphi(v) \in A \setminus B$ is a neighbour of $\varphi(w) \in \varphi(B \setminus A) = B \setminus A$, which contradicts the fact that $\{A, B\}$ is a separation.

As B is infinite, we also have $\varphi(B) \not\subseteq A$. So if $\{A, B\}$ and $\{\varphi(A), \varphi(B)\}$ were nested, we would have $(A, B) \leqslant (\varphi(B), \varphi(A))$, and thus $B \supseteq \varphi(A)$. But $\varphi(v) \in \varphi(A) \setminus B$ by assumption, a contradiction. □

We remark that Lemma 1.7.9 has no analogue for finite graphs, or for separations with two infinite sides. Indeed, for finite graphs consider a cycle of odd length at least 5 and an arbitrary tight separation of order 2 together with an automorphism which fixes precisely one of the vertices in the separator. For separations with two infinite sides, consider a double ray and an arbitrary separation of order 1 together with any automorphism which shifts the double ray.

With these three lemmas at hand, we are now ready to prove Lemma 1.7.6.

Proof of Lemma 1.7.6. Suppose, for a contradiction, that $(A_i, B_i)_{i\in\mathbb{N}}$ is a strictly increasing sequence in $\vec{N}$. Then $A_i \setminus B_i \subseteq A_j \setminus B_j$ for all $i < j$, and by Lemmas 1.7.7 and 1.7.8 all the A_i are finite.

We first show that there exists $n \in \mathbb{N}$ such that $A_i \cap B_i \subseteq A_{i+1} \cap B_{i+1}$ for all $i \geqslant n$. If not, there are infinitely many i such that $A_i \cap B_i$ contains a vertex $v_i \in A_{i+1} \setminus B_{i+1}$. By our initial observation, these v_i are distinct for different i. Since Γ acts quasi-transitively on G, there exist $i < j$ among these with $v_i = g \cdot v_j$ for some $g \in \Gamma$. Then $\{A_j, B_j\}$, together with $v_j \in A_j \cap B_j$ and the automorphism φ of G defined by g, satisfies the premise of Lemma 1.7.9, since $g \cdot v_j = v_i \in A_j \setminus B_j$ and $\{A_j, B_j\}$ is tight. Since N is Γ-canonical, applying Lemma 1.7.9 yields a contradiction to N being nested.

We next show that there exists $m \geqslant n$ such that $A_i \cap B_i = A_{i+1} \cap B_{i+1}$ for all $i \geqslant m$. If not, then $A_i \cap B_i \subsetneq A_{i+1} \cap B_{i+1}$ for infinitely many $i \geqslant n$; pick a subsequence $(A_{i_j}, B_{i_j})_{j \in \mathbb{N}}$ with $i_0 = n$ witnessing this. Since the separations (A_{i_j}, B_{i_j}) are tight, every $G[A_{i_j} \setminus B_{i_j}]$ has a component K_{i_j} with neighbourhood $A_{i_j} \cap B_{i_j}$. Then $K_{i_j} \cap K_{i_{j+1}} = \emptyset$ for all j, since K_{i_j} is also a component of $G[A_{i_{j+1}} \setminus B_{i_{j+1}}]$ but distinct from $K_{i_{j+1}}$, since the two have different neighbourhoods. As G is connected, there exists $v \in A_n \cap B_n$. This v lies in $A_{i_j} \cap B_{i_j}$ for every j, and hence has a neighbour in K_{i_j}. As these neighbours are distinct, this contradicts our assumption that $v \in G$ has finite degree.

The (A_i, B_i) with $i \geqslant m$ form a strictly increasing sequence of separations with the same separator X. This is impossible since, as G is connected and locally finite, the graph $G - X$ has only finitely many components. □

1.7.3 Proof of Lemma 1.7.1

In the remainder of this section we prove Lemma 1.7.1. Let N, N_{end}, Γ and $K > 0$ be as stated in the lemma. Our goal is to show that N is exhaustive: that every strictly increasing sequence $(A_i, B_i)_{i \in \mathbb{N}}$ of separations in $\vec{N}$ satisfies $\bigcap_i B_i = \emptyset$.

We already know from Lemma 1.7.5 that N_{end} is exhaustive, since by assumption the separations in N_{end} have bounded order. If we can show that our sequence $(A_i, B_i)_{i \in \mathbb{N}}$ dominates a sequence in $\vec{N}_{\text{end}}$, it will also be exhaustive (Lemma 1.7.3). Essentially, this amounts to showing that our sequence is interlaced with a sequence in $\vec{N}_{\text{end}}$: that at most finitely many (A_i, B_i) have the same $\leqslant$-relationships to all the elements of $\vec{N}_{\text{end}}$. Our plan is to show that if there were infinitely many such (A_i, B_i) then these would, essentially, disect the same 1-ended subgraph of G, which would contradict Lemma 1.7.6.

In graph terms this means that, essentially, for the tree-decomposition $\mathcal{T}' = (T, \mathcal{V}')$ of G that N_{end} induces (see below), where $\mathcal{V}' = (V'_t)_{t \in T}$ say, there is no $t \in T$ such that infinitely many (A_i, B_i) differ only in how they cut up V'_t, because this would contradict Lemma 1.7.6. This will not work directly, since the graphs $G[V'_t]$ need not satisfy the conditions on G laid down in the premise of Lemma 1.7.6. However we shall be able to extend the V'_t to sets V_t that do satisfy these conditions, and then apply Lemma 1.7.6 to the graphs $G[V_t]$ to complete our proof.

To get started, recall from Lemmas 1.7.2 and 1.7.5 that N_{end} does induce the required tree-decomposition:

Lemma 1.7.10. *N_{end} is exhaustive, and the tree-decomposition $\mathcal{T}'$ of G it induces is Γ-canonical.* □

As N_{end} distinguishes all the ends of G, each part of $\mathcal{T}'$ meets rays from at most one end of G infinitely.

Since N_{end} is Γ-canonical, the group Γ acts on the decomposition tree T of $\mathcal{T}'$ (see Section 1.7.1). Since G is quasi-transitive under the action of Γ, there are only finitely many Γ-orbits of tight separations of G of order less than K [TW93, Proposition 4.2]. The action of Γ on T, therefore, is quasi-transitive too:

Lemma 1.7.11. *Γ acts on $E(T)$ with finitely many orbits.* □

Let st be any edge of T. Since G is locally finite and $V'_{st} := V'_s \cap V'_t$ is finite, the graph $G - V'_{st}$ has only finitely many components. Hence there exists an integer $d_{st} \geqslant 1$ such that the d_{st}-ball $B_G(V'_{st}, d_{st})$ around V'_{st} (see Section 1.4.3) meets every component C of $G - V'_{st}$ in a set that includes all the shortest paths in C between neighbours of V'_{st} in C. (There are only finitely many such paths, since G is locally finite.) By Lemma 1.7.11, Γ has finitely many orbits on $E(T)$. Let us assume that the integers d_{st} were all chosen as small as possible; then they coincide for any two edges st in the same orbit. Thus, the maximum d over all d_{st} for $st \in E(T)$ is well-defined; let $V_t := V(B_G(V'_t, d))$ and $\mathcal{V} := (V_t)_{t \in T}$.

Lemma 1.7.12. *$\mathcal{T} = (T, \mathcal{V})$ is a Γ-canonical tree-decomposition of G of finitely bounded adhesion such that the following statements hold for every $t \in T$:*

(i) *$G[V_t]$ is connected;*

(ii) *the stabiliser Γ_t of t under the action of Γ on T acts quasi-transitively on $G[V_t]$;*

(iii) *$G[V_t]$ is either finite or 1-ended.*

The proof of Lemma 1.7.12 closely follows [HLMR22]; details are given in Section 1.9 below.

Recall that our plan is to show that our sequence $(A_i, B_i)_{i \in \mathbb{N}}$ in $\overrightarrow{N}$ dominates a suitable sequence in $\overrightarrow{N}_{\mathrm{end}}$. To achieve this, it will suffice to show that no part V'_t of $\mathcal{T}'$ contains all the separators from a final segment $(A_i, B_i)_{i \geqslant i_0}$ of our sequence. We will prove this by applying Lemma 1.7.6 to the corresponding part $G[V_t]$ of $\mathcal{T}$.

To apply Lemma 1.7.6 in this way, we need to show that a separation of G whose separator is contained in V'_t induces a tight separation of $G[V_t]$. It does, because the graph $G[V_t]$ behaves like a torso of $G[V'_t]$ (see [Die25, Ch.12.3] if desired):

Lemma 1.7.13. *Let t be a node of T and $X \subseteq V'_t$. Then the map $C \mapsto C \cap G[V_t]$ is a bijection between the components of $G - X$ and those of $G[V_t] - X$. Moreover, $N_{G[V_t]}(C \cap G[V_t]) = N_G(C)$ for all such C.*

Proof. Let C be any component of $G - X$. Since G is connected, C contains a neighbour of X. All the neighbours of X in C lie in V_t, since $N_G(X) \subseteq B_G(V'_t, 1) \subseteq B_G(V'_t, d) = V_t$.

Hence C meets some component D of $G[V_t] - X$. We shall prove that D is the only component of $G[V_t] - X$ that C meets; then clearly $D = C \cap G[V_t]$, which establishes the desired bijection.

Since $G[V_t]$ is connected (Lemma 1.7.12), every component of $G[V_t] - X$ contains a neighbour of X. It thus suffices to show that all the neighbours u, v of X in C lie in the same component of $G[V_t] - X$ (our D). We prove this by showing that every shortest u–v path P in $G - X$ lies in $G[V_t]$. Note that $P \subseteq C$.

Let x be a neighbour of u in X, and y a neighbour of v in X, possibly $x = y$. Then $xuPvy$ is a concatenation of walks in G that either lie in $G[V_t']$ or have their first and last vertex in V_t' and all their inner vertices (at least one) in the same component of $G - V_t'$. Let $Q = x' \dots y'$ be any subwalk of $xuPvy$ of the latter type if one exists (which will be a path or the entire walk $xuPvy$, in which case that is a cycle), and let C' be the component of $G - V_t'$ that contains $Q - x' - y'$. It will suffice for our proof of $P \subseteq G[V_t]$ to show that $Q \subseteq G[V_t]$.

As $X \subseteq V_t'$ and C' meets C (in $Q - x' - y' \neq \emptyset$, which lies in $P \subseteq C$), we have $C' \subseteq C$. Since $\mathcal{T}'$ is a tree-decomposition, $V(C')$ is contained in $\bigcup_{t' \in T'} V_{t'}'$ for some component T' of $T - t$. Let $s \in T'$ be the neighbour of t in that component. Then all the neighbours of C' in V_t' lie in $V_{st}' \subseteq V_t'$, and so the component of $G - V_{st}'$ that contains C' (which exists because $V_{st}' \subseteq V_t'$) in fact equals C'.

Thus, $Q - x' - y'$ is a path in a component of $G - V_{st}'$ between neighbours of V_{st}' (possibly identical). Since it is a subpath of P, it is also a shortest such path. Hence $Q \subseteq G[V_t]$ by the definition of V_t. □

To complete our proof of Lemma 1.7.1 we need one more definition. Let $\vec{r}, \vec{s}$ be orientations of proper separations of G, such as those in $\overrightarrow{N}$. If $r = s$ we say that $\vec{r}$ *points towards* $\vec{s}$ if $\vec{r} = \vec{s}$. If $r \neq s$ we say that $\vec{r}$ *points towards* $\vec{s}$ (and s) if either $\vec{r} \leqslant \vec{s}$ or $\vec{r} \leqslant \overleftarrow{s}$. Note that if $r \neq s$ are nested, they have unique orientations $\vec{r}$ and $\overleftarrow{s}$ that point towards each other, with $\vec{r} \leqslant \vec{s}$ and $\overleftarrow{r} \geqslant \overleftarrow{s}$. Also, since r and s are proper, we cannot have both $\vec{r} \leqslant \vec{s}$ and $\vec{r} \leqslant \overleftarrow{s}$.

Completion of proof of Lemma 1.7.1. Recall that we have a sequence $(A_i, B_i)_{i \in \mathbb{N}}$ of separations in $\overrightarrow{N}$, which we want to show to be exhaustive. Let us abbreviate (A_i, B_i) as $\vec{s}_i$. The tools at our disposal are that $N_{\text{end}} \subseteq N$ is exhaustive (Lemma 1.7.10), our tree-decomposition of G into finite or 1-ended parts provided by Lemma 1.7.12, and Lemma 1.7.6 by which 1-ended graphs contain no infinite sequences such as $(\vec{s}_i)_{i \in \mathbb{N}}$ at all.

We need three more steps. We first establish a partition of $\overrightarrow{N}$ defined by the splitting stars σ of N_{end}. Recall that these correspond to the nodes t of T (see Section 1.7.1); the corresponding partition class P_σ of $\overrightarrow{N}$ will consist, very roughly, of those separations in $\overrightarrow{N}$ that lie in σ or separate V_t. Next we show that for our proof that $(\vec{s}_i)_{i \in \mathbb{N}}$ is exhaustive it suffices to show that it contains separations from infinitely many of these partition classes P_σ. Finally, we check that $(\vec{s}_i)_{i \in \mathbb{N}}$ does meet infinitely many P_σ.

Since N_{end} is exhaustive, its splitting stars form a partition of $\overrightarrow{N}_{\text{end}}$: every element of $\overrightarrow{N}_{\text{end}}$ lies in exactly one splitting star [GK21, Lemma 2.4]. Let us extend this partition of $\overrightarrow{N}_{\text{end}}$ to

one of $\vec{N}$, as follows. For every splitting star σ of N_{end} let

$$P_\sigma := \{\vec{s} \in \vec{N} : \text{all the elements of } \sigma \text{ point towards } \vec{s}\}.$$

Note that $P_\sigma \cap \vec{N}_{\text{end}} = \sigma$.

Let us show that these P_σ form a partition of $\vec{N}$. As the separations in N are nested, every $\vec{s} \in \vec{N}$ defines a consistent orientation of N_{end}:

$$O_{\vec{s}} := \{\vec{r} \in \vec{N}_{\text{end}} : \vec{r} \text{ points towards } \vec{s}\};$$

let $\sigma_{\vec{s}}$ denote the set of its maximal elements. Since N_{end} is exhaustive and s is proper, $O_{\vec{s}}$ contains no infinite strictly increasing sequence. Hence every $\vec{r} \in O_{\vec{s}}$ lies below some maximal element of $O_{\vec{s}}$; thus, $\sigma_{\vec{s}}$ is a splitting star of N_{end}. So $P_{\sigma_{\vec{s}}}$ is defined, and it clearly contains $\vec{s}$. But $\vec{s}$ cannot lie in any other P_σ. Indeed, consider a splitting star $\sigma \neq \sigma_{\vec{s}}$ of N_{end}. The orientation O of N_{end} of which σ is the set of maximal elements then orients some separation $r \in N_{\text{end}}$ differently from the way $O_{\vec{s}}$ does. Hence if $\vec{s} \in P_\sigma \cap P_{\sigma_{\vec{s}}}$ then both orientations of r lie below an element of σ or $\sigma_{\vec{s}}$, and hence both point towards $\vec{s}$. This is impossible, since s is proper.

Next we show that $(\vec{s}_i)_{i\in\mathbb{N}}$ is exhaustive if it meets infinitely many P_σ. By Lemma 1.7.3 we may assume that the $\vec{s}_i$ lie in distinct P_σ. For each $i \in \mathbb{N}$ let σ_i be such that $\vec{s}_i \in P_{\sigma_i}$. Then $O_i := O_{\vec{s}_i}$ is the consistent orientation of N_{end} of which σ_i is the set of maximal elements. while the other, $\overleftarrow{r}_i$, points towards $\vec{s}_i$. Since $\vec{s}_i < \vec{s}_{i+1}$, this is possible only if $\vec{s}_i < \vec{r}_i \leqslant \vec{s}_{i+1}$. (The first inequality is strict, because otherwise $\overleftarrow{r}_i$ would point towards its own inverse $\vec{r}_i = \vec{s}_i$, which is impossible.) Hence the sequence $(\vec{r}_i)_{i\in\mathbb{N}}$ is strictly increasing too, and is dominated by $(\vec{s}_i)_{i\in\mathbb{N}}$. Since $(\vec{r}_i)_{i\in\mathbb{N}}$ consists of separations from $\vec{N}_{\text{end}}$ and is thus exhaustive by Lemma 1.7.10, our sequence $(\vec{s}_i)_{i\in\mathbb{N}}$ is exhaustive by Lemma 1.7.3.

It remains to check that $(\vec{s}_i)_{i\in\mathbb{N}}$ does meet infinitely many P_σ. To prove this we show that each P_σ contains only finitely many $\vec{s}_i$. Recall that $\mathcal{T}' = (T, \mathcal{V}')$ was the tree-decomposition of G induced by N_{end}, and that our group Γ of automorphisms of G acts on T. As shown in Section 1.7.1, the splitting stars σ of N_{end} correspond to the nodes t of T. Given such a pair σ and t, let $N_t \subseteq N$ be the set of all separations with an orientation in P_σ. Note that the stabiliser Γ_t of t under the action of Γ on T consists precisely of those elements of Γ that fix σ as a set.

Let us show that N_t is Γ_t-canonical: that for every $g \in \Gamma_t$ and every $s \in N_t$ we have $g \cdot s \in N_t$. By definition of N_t and P_σ the latter means that every $\overrightarrow{r'} \in \sigma$ points towards $g \cdot s$: that $s' := g \cdot s$ has an orientation $\overrightarrow{s'} \geqslant \overrightarrow{r'}$. Let $\vec{r} := g^{-1} \cdot \overrightarrow{r'}$. As $g \in \Gamma_t$ we have $\vec{r} \in \sigma$; thus, $\vec{r}$ also points towards s, say $\vec{r} \leqslant \vec{s}$. But then $\overrightarrow{r'} = g \cdot \vec{r} \leqslant g \cdot \vec{s} =: \overrightarrow{s'}$ as required.

Let us turn N_t, a set of separations of G, into a Γ_t-canonical set M_t of tight separations of $G[V_t]$. By the definition of $\mathcal{T}'$, the part V_t' is equal to $\bigcap_{(C,D)\in\sigma} D$. As for every $\{A, B\} \in N_t$ and every $(C, D) \in \sigma$ either A or B is a subset of D, we thus have $A \cap B \subseteq V_t'$. Every oriented separation $(A, B) \in \vec{N}_t$ induces an oriented separation $(A_t, B_t) := (A \cap V_t, B \cap V_t)$ of $G[V_t]$. By Lemma 1.7.13 one easily shows that these are distinct for different (A, B), that they are tight because the (A, B) are tight, and that for $(A, B) < (C, D)$ in $\vec{N}_t$ we have $(A_t, B_t) < (C_t, D_t)$. Let

$$M_t := \{\{A_t, B_t\} \colon \{A, B\} \in N_t\}.$$

M_t is a Γ_t-canonical set of tight separations of $G[V_t]$. If $G[V_t]$ is finite then so is M_t, and in particular $\vec{M}_t$ contains no strictly increasing sequence. Otherwise $G[V_t]$ is a connected locally finite 1-ended graph that is quasi-transitive under the action of Γ_t, by Lemma 1.7.12. Hence, by Lemma 1.7.6, $\vec{M}_t$ contains no strictly increasing infinite sequence. Since any strictly increasing sequence in $\vec{N}_t$ induces a strictly increasing sequence in $\vec{M}_t$, by Lemma 1.7.13 as noted earlier, there is thus no strictly increasing infinite sequence in $\vec{N}_t$ either. This completes our proof that $P_\sigma \subseteq \vec{N}_t$ contains only finitely many $\vec{s}_i$, and hence that our sequence $(\vec{s}_i)_{i\in\mathbb{N}}$ is exhaustive. By Lemma 1.7.2, N induces a Γ-canonical tree-decomposition of G, as required in the statement of Lemma 1.7.1. □

1.8 Proof of the main result of this chapter

In this section we combine all our lemmas to prove Theorem 1.5.5 – the more general version of Theorem 1 that also includes Theorem 2 – and derive some corollaries of the proof. The unique graph-decomposition of the given graph G whose existence is asserted in Theorem 1.5.5, the decomposition that displays its r-global structure (if it exists), was specified in Definition 1.5.4. It is defined (via the projection $p_r\colon G_r \to G$) by the tree-decomposition $\mathcal{T}$ of G_r that is induced by its canonical tree of tangles $N(G_r)$ from Definition 1.5.2 if this exists.

We shall derive the existence of $N(G_r)$ from Lemma 1.5.3, which requires G_r to be tangle-accessible. Lemma 1.8.1 helps ensure this when G_r is accessible, which we already know from Theorem 1.6.2.

Let us say that a set S of separations of a graph G is *point-finite* if for every vertex $v \in V(G)$ the set $S_v := \{\{A,B\} \in S\colon v \in A\cap B\}$ is finite. It is easy to see that a regular tree-decomposition has finite spread if (and only if) its induced nested set of separations is point-finite. Indeed, if a vertex of G lies in infinitely many parts then, by (T2) and [Die25, Prop. 8.2.1], the decomposition tree has an infinite star or a ray such that v lies in every part V_t corresponding to a node t in that star or ray. Then v also lies in all the separators that correspond to edges of that star or ray.

Lemma 1.8.1. *Let G be a connected locally finite graph, and let N be a nested set of tight finite-order separations of G. Let Γ be a group acting quasi-transitively on G so that N is Γ-canonical. Suppose that there exists a number $K \in \mathbb{N}$ such that some Γ-canonical set $N_{\mathrm{end}} \subseteq N$ of separations of order $< K$ distinguishes all the ends of G.*

Then N is exhaustive and point-finite, and the order of the separations in N is finitely bounded. The tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G which N induces is regular, Γ-canonical, of finitely bounded adhesion, and has finite spread. If G is 1-ended, then T is rayless.

For a proof of Lemma 1.8.1 we need two observations.

Lemma 1.8.2. *Let G be a connected locally finite graph, and let N be a nested set of tight separations of G. If N is exhaustive, then N is point-finite.*

Proof. Let $v \in V(G)$, and suppose the set N_v of separations in N whose separator contains v is infinite. Since N is exhaustive, so is N_v. By Lemma 1.7.2, N_v induces a regular tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G. If T contains a ray, then the forward-oriented edges of this ray define

an infinite strictly increasing sequence of separations in N_v via $\alpha_{\mathcal{T}}$. All these separations have v in their separator, which contradicts the fact that N_v is exhaustive.

So T is rayless, and hence has a vertex t of infinite degree [Die25, Prop. 8.2.1]. Applying $\alpha_{\mathcal{T}}$ to the edges incident with t and oriented towards t yields an infinite set $\{(A_i, B_i) \colon i \in \mathbb{N}\} \subseteq \vec{N}_v$ such that $(A_i, B_i) < (B_j, A_j)$ for all $i \neq j$. All the sets $A_i \setminus B_i$ are disjoint, and v has a neighbour in every $A_i \setminus B_i$ since $v \in A_i \cap B_i$ and $\{A_i, B_i\}$ is tight. This contradicts the local finiteness of G. □

In quasi-transitive graphs, canonical and point-finite sets of separations are again quasi-transitive under the action of the same group:

Lemma 1.8.3. *Let G be a connected graph, not necessarily locally finite, and let S be a point-finite set of finite-order separations of G. Let Γ be a group acting quasi-transitively on G so that S is Γ-canonical. Then Γ acts on S with finitely many orbits. In particular, the order of the separations in S is finitely bounded.*

Proof. As G is quasi-transitive under the action of Γ, there exists a finite set V' of vertices of G such that $\Gamma \cdot V' = V(G)$. Then $S \setminus \{\{\emptyset, V(G)\}\} = \Gamma \cdot S'$ where $S' = \bigcup_{v' \in V'} S_{v'}$. Since S is point-finite, all these $S_{v'}$ are finite, and hence so is S'. Thus, Γ acts on S with finitely many orbits, while $\{\emptyset, V(G)\}$, if it lies in S, forms an additional orbit in S. Since S' is finite, the order of the separations in S is finitely bounded. □

Proof of Lemma 1.8.1. By Lemma 1.7.1, N is exhaustive. By Lemma 1.8.2 it is point-finite, and by Lemma 1.8.3 its separations have finitely bounded order. By Lemma 1.7.2, N induces a regular Γ-canonical tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G. Its adhesion sets are the separators of the separations in N, so it has finitely bounded adhesion. As remarked just before Lemma 1.8.1, the point-finiteness of N implies that $\mathcal{T}$ has finite spread, too.

If G is 1-ended then, by Lemma 1.7.6, $\vec{N}$ contains no infinite strictly increasing sequence. Any ray in T would define such a sequence via $\alpha_{\mathcal{T}}$, so T is rayless. □

Lemma 1.8.1 essentially strengthens accessibility to tangle-accessibility in our context. We need this for our proof of Theorems 1 and 2, as outlined at the start of this section. More precisely, what we shall need is Theorem 3, which we will thus prove first. Our input N for Lemma 1.8.1 in the proof of Theorem 3 will come from the following non-trivial result:

Lemma 1.8.4 ([EKT22], Theorem 6.6)**.** *Every connected, locally finite graph has a canonical tree of tangles consisting of relevant separations.*[6]

We can now prove Theorem 3, which we restate:

Theorem 3. *Accessible locally finite, quasi-transitive graphs are even tangle-accessible.*

[6]The relevance of the separations is not explicitly stated in [EKT22, Theorem 6.6] but immediate from its proof.

Proof. Let G be a graph as stated. Since it is quasi-transitive, it has only finitely many isomorphism types of components, so we may assume G to be connected. Let $N := N'(G)$ be the tree of tangles given by Lemma 1.8.4.

By Lemma 1.5.1, N is a set of tight finite-order separations of G. Since ends define tangles, and these are distinguished by exactly the separations that distinguish those ends (Section 1.5.1), N distinguishes the ends of G efficiently.

Since G is accessible, the order of the separations of G needed to distinguish its ends is finitely bounded. As N is canonical and distinguishes the ends efficiently, it thus has a canonical subset N_{end} of separations of finitely bounded order that achieves this too – for example, its subset of all efficient end-distinguishers. By Lemma 1.8.1, then, the separations in N have finitely bounded order too. Since N distinguishes the set of all tangles of G, it thus witnesses that G is tangle-accessible. □

Note that the tree of tangles $N'(G)$ given by Lemma 1.8.4, which we just used in the proof of Theorem 3, exists for arbitrary connected locally finite graphs, not just quasi-transitive ones. In general it may have limit points, in which case it will not define a tree-decomposition. But even when G is tangle-accessible and $N'(G)$ does define a tree-decomposition, it need not coincide with our canonical tree of tangles $N(G)$ from Definition 1.5.2.

Turning now to our canonical tree of tangles $N(G)$ from Definition 1.5.2, recall that in order to ensure its existence we needed that G is tangle-accessible (Lemma 1.5.3). By Theorem 3, this now follows already from classical accessibility, which we shall obtain from Theorem 1.6.2.

Our next lemma records that, in this case, $N(G)$ induces the desired tree-decomposition:

Lemma 1.8.5. *Let G be a connected, locally finite, quasi-transitive, accessible graph. Then $N(G)$ is defined and induces a canonical, regular tree-decomposition of G of finitely bounded adhesion and with finite spread. If G is 1-ended, the decomposition tree is rayless.*

Proof. By Theorem 3 and Lemma 1.5.3, $N(G)$ exists, is canonical, and consists of relevant (and therefore tight) separations of finitely bounded order. Since $N(G)$ distinguishes all the ends of G, it induces the desired tree-decomposition by Lemma 1.8.1. □

Summing everything up, we can now prove Theorem 1.5.5:

Proof of Theorem 1.5.5. Let G and r be as stated in the theorem. The r-local cover G_r of G is connected, locally finite, and quasi-transitive by Lemma 1.6.3, and accessible by Theorem 1.6.2. By Lemma 1.8.5, $N(G_r)$ is defined and induces a canonical, regular tree-decomposition $\mathcal{T}$ of G_r of finite spread. As $p_r\colon G_r \to G$ is also canonical (see Section 1.4.2), $\mathcal{T}$ defines a canonical graph-decomposition of G of finite spread (Theorem 1.3.13). By Definition 1.5.4, this is the unique graph-decomposition of G that displays its r-global structure. □

Proof of Theorem 1. Theorem 1 is the special case of Theorem 1.5.5 where G is finite. □

Proof of Theorem 2. Since $\mathrm{Cay}(\Gamma, S)$ is locally finite and transitive, Theorem 1.5.5 implies Theorem 2. □

1.9 Tree-Decompositions in quasi-transitive graphs

In this section we give a detailed proof of Lemma 1.7.12. As our construction is related to that in [HLMR22, Proposition 4.1], most of the arguments in the proof are essentially adaptations of the respective proofs in [HLMR22, Section 4]. For better readability of this appendix, we split up Lemma 1.7.12 into four separate lemmas, Lemmas 1.9.1 to 1.9.4, which we prove in turn.

Let us briefly recall the setup we are working in: G is a connected, locally finite graph, and Γ is a group acting quasi-transitively on G. Given the Γ-canonical tree-decomposition $\mathcal{T}' = (T, \mathcal{V}')$ of G of finitely bounded adhesion from Lemma 1.7.10, we found in Section 1.7.3 an integer $d > 0$ such that for every edge $tt' \in T$ and every component C of $G - V'_{tt'}$, where $V'_{tt'} := V'_t \cap V'_{t'}$, the graph $C \cap B_G(V'_{tt'}, d)$ contains all the shortest paths in C between neighbours of $V'_{tt'}$. Clearly, all these graphs $C \cap B_G(V'_{tt'}, d)$ are connected. We then defined $\mathcal{T} = (T, \mathcal{V})$ by letting $V_t := V(B_G(V'_t, d))$ for each node $t \in T$.

Let us begin by showing Lemma 1.7.12 (i), that $(T, \mathcal{V})$ is a tree-decomposition into connected parts:

Lemma 1.9.1. *$\mathcal{T} = (T, \mathcal{V})$ is a Γ-canonical tree-decomposition of finitely bounded adhesion whose parts are connected.*

Proof. We first check that $\mathcal{T}$ is indeed a tree-decomposition. The Γ-canonicity of $\mathcal{T}$ then follows immediately from the construction and the Γ-canonicity of $\mathcal{T}'$. (T1) is clear. For (T2), let t_1, t_2, t_3 be nodes of T with $t_3 \in t_1 T t_2$, and v any vertex in $V_{t_1} \cap V_{t_2}$. We consider two cases.

Assume first that $v \in V'_{t_1} \cap V_{t_2}$. Let w be a vertex of V'_{t_2} such that $d_G(v, w) \leqslant d$, and fix a shortest v–w path P in G. Then P meets V'_{t_3}, by (T2) for $\mathcal{T}' = (T, \mathcal{V}')$. Hence, $v \in V(B_G(V'_{t_3}, d)) = V_{t_3}$. Note that the case of $v \in V_{t_1} \cap V'_{t_2}$ is analogous. This completes the first case.

For the second case, assume that $v \in (V_{t_1} \setminus V'_{t_1}) \cap (V_{t_2} \setminus V'_{t_2})$. By (T1) for $\mathcal{T}' = (T, \mathcal{V}')$, there exists a node t_4 of T such that $v \in V'_{t_4}$. Since the t_1–t_2 path in T is unique, the node t_3 lies on the t_1–t_4 path in T or on the t_2–t_4 path in T. So we are done by applying the first case to the pair t_1, t_4 or the pair t_2, t_4. This completes the proof of (T2), so $\mathcal{T}$ is a tree-decomposition.

For any edge tt' of T, the corresponding adhesion set in $\mathcal{T}$ has the form $V(B_G(V'_{tt'}, d))$. Thus, the adhesion sets of $\mathcal{T}$ are of finitely bounded size, since G is locally finite and $\mathcal{T}'$ has finitely bounded adhesion. Moreover, our choice of d guarantees that all the adhesion sets of $\mathcal{T}$ are connected. It remains to deduce from this that $G[V_t]$ is connected for every node t of T.

Recall that G is connected. So consider any two vertices $u, v \in V_t$ and a u–v path P in G with as few vertices outside V_t as possible. Now for any maximal subpath P' of P whose internal vertices are not in V_t, both endvertices are in $G[V_t]$ and more precisely in the same adhesion set $V_{tt'}$ for some $t' \in N_T(t)$. But since $V_{tt'}$ is connected, we may replace P' by a path in $G[V_{tt'}]$, hence obtaining a u–v walk in G with strictly fewer vertices outside V_t than P, and thus such a u–w path in G. So the minimal choice of P yields that P must be completely contained in $G[V_t]$, which completes the proof. □

Next we prove Lemma 1.7.12 (ii): that a subgroup of Γ acts quasi-transitively on each part $G[V_t]$ of $\mathcal{T}$.

Lemma 1.9.2. *For every node $t \in T$, the stabiliser Γ_t of t under the action of Γ on T acts quasi-transitively on $G[V_t]$.*

Proof. Since $\mathcal{T}$ is Γ-canonical, we have $\varphi(V_t) = V_t$ for every $\varphi \in \Gamma_t$; thus, the action of Γ on G induces an action of Γ_t on $G[V_t]$. Moreover, $\varphi(V_{st}) = V_{\varphi(s)\varphi(t)}$ for every $\varphi \in \Gamma$. Hence, Γ_t acts on the set $W_t := \{v \in V_t \mid \forall s \in N_T(t) \colon v \notin V_{st}\}$ of 'inner vertices' of V_t and on the set $V_t \setminus W_t$. It is enough to show that Γ_t does so quasi-transitively.

To show that Γ_t acts quasi-transitively on W_t, let $u, v \in W_t$ be such that $\varphi(u) = v$ for some $\varphi \in \Gamma$. Since $\mathcal{T}$ is a Γ-canonical tree-decomposition, $\varphi(W_t) = W_{\varphi(t)}$. By (T2), t is the unique node of T containing v as inner vertex, and hence $t = \varphi(t)$. Thus, $\varphi \in \Gamma_t$. So, Γ_t has finitely many orbits on W_t, since Γ acts quasi-transitively on G.

Second, let us consider the action of Γ_t on $V_t \setminus W_t$: the union of all adhesion sets associated with edges of T at t. Since each adhesion set V_{st} is finite, the action of Γ_t has only finitely many orbits on this adhesion set: at most $|V_{st}|$. Thus, it is enough to show that Γ_t acts quasi-transitively on the set $\{st \mid s \in N_T(t)\}$ of edges of T. By Lemma 1.7.11, there are only finitely many Γ-orbits of $E(T)$. We now claim that each Γ-orbit O of $E(T)$ has at most two Γ_t-orbits. For this, let s_0, s_1, s_2 be neighbours of t in T such that $ts_0, ts_1, ts_2 \in O$. Then there exist $\varphi_1, \varphi_2 \in \Gamma$ such that $\varphi_1(ts_0) = ts_1$ and $\varphi_2(ts_0) = ts_2$. So if φ_1 and φ_2 are not in Γ_t, then $\varphi_2\varphi_1^{-1} \in \Gamma_t$. This proves the claim and thus completes the proof. □

Towards a proof of Lemma 1.7.12 (iii), we first show that $\mathcal{T}$ distinguishes all ends of G, just as $\mathcal{T}'$ does by definition:

Lemma 1.9.3. *$\mathcal{T}$ distinguishes all ends of G.*

Proof. The tree-decomposition $\mathcal{T}'$ distinguishes all the ends of G, because it is induced by N_{end} (Lemma 1.7.10). Let tt' be an edge of T, and let $\alpha_{\mathcal{T}}(t, t') = (A, B)$ be the oriented separation of G associated with tt' in $\mathcal{T}$. By construction of $\mathcal{T}$, we have $(A, B) = (B_G(A', d), B_G(B', d))$, where (A', B') is the oriented separation $\alpha_{\mathcal{T}'}(t, t')$ associated with tt' in $\mathcal{T}'$. Thus, the separator $A \cap B$ of (A, B) is $B_G(A' \cap B', d)$. Hence it is finite and contains $A' \cap B'$. So (A, B) distinguishes every two ends of G that are distinguished by (A', B'). Altogether, the tree-decomposition $\mathcal{T}$ distinguishes all the ends of G because $\mathcal{T}'$ does. □

Our last lemma deduces Lemma 1.7.12 (iii) from Lemma 1.9.3: that the parts of $\mathcal{T}$ have at most one end:

Lemma 1.9.4. *Let t be a node of T. Then $G[V_t]$ is either finite or 1-ended.*

Proof. If V_t is infinite, then $G[V_t]$ is an infinite and locally finite graph, which is connected by Lemma 1.9.1. By König's infinity lemma [Die25, Proposition 8.2.1], this graph $G[V_t]$ has at least one end ω. Since $\mathcal{T}$ has finitely bounded adhesion, the end ω_t of G with $\omega \subseteq \omega_t$ has the

property that every ray $R \in \omega_t$ meets V_t infinitely often. We claim that ω_t is the unique end of G with this property.

Suppose there exists another such end ω_t' of G. Since $\mathcal{T}$ distinguishes all ends of G, by Lemma 1.9.3, T has an edge whose associated separation $\{A, B\}$ of G distinguishes ω_t from ω_t'. Then the part V_t lies on one side of this separation, say in A. Any ray R in ω_t or ω_t' meets the finite set $A \cap B$ only finitely often, and hence has a subray of R on one side of $\{A, B\}$. This side must be A, since R meets V_t infinitely often by our assumption about ω_t and ω_t'. But this contradicts the fact that $\{A, B\}$ distinguishes ω_t from ω_t'.

So ω_t is the unique end of G with the property that every ray $R \in \omega_t$ meets V_t infinitely often. The statement now follows directly from [HLMR22, Proposition 4.8 (iv)]. □

CHAPTER 2

Canonical graph decompositions via local separations

This chapter is based on [CJKK25], which is joint work with Johannes Carmesin, Raphael W. Jacobs, and Jan Kurkofka.

2.1 Introduction

2.1.1 From tree-decompositions to graph-decompositions

Tree-decompositions are an established concept to display the connectivity-structure of graphs when this resembles that of a tree. Roughly speaking, a tree-decomposition (Figure 2.1) is a blueprint which shows how a graph G (black) can be built from smaller graphs (grey), called *parts*, by gluing them together in the shape of a tree (blue), called the *decomposition tree*. Each edge of the decomposition tree determines a *separator* of G (circled in red), and more specifically induces a *separation* which describes G as a union of two halves that overlap precisely in that separator.

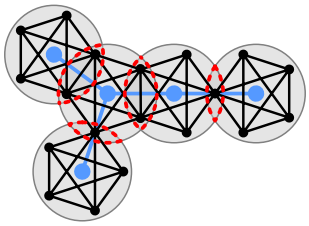

Figure 2.1: A tree-decomposition of a graph (black) into K_5's (grey).

Tree-decompositions are a key concept in Robertson and Seymour's proof of their Graph Minor Theorem [RS04]. In computer science, many algorithmic problems that are NP-hard in general can be decided in linear time for classes of graphs admitting tree-decompositions into parts of bounded size, by Courcelle's Theorem [Cou90]. In group theory, and in Bass-Serre theory more specifically, tree-decompositions are used to detect product-structure in Cayley graphs such as amalgamated free products or HNN-extensions, for example in modern proofs of Stallings' Theorem [HLMR22; Krö10].

One of the first applications of tree-decompositions sits at the interface of cluster analysis and graph theory. Even though clusters are fuzzy in nature, there exist precise notions that describe clusters in graphs. A well-known example is the abstract notion of a tangle, introduced by Robertson and Seymour alongside tree-decompositions [RS91]. Their *tree-of-tangles theorem* states that every finite graph G has a tree-decomposition such that every two tangles in G are minimally separated by some separation of the tree-decomposition. Such tree-decompositions are referred to as *trees of tangles*. An example is depicted in Figure 2.1, where each K_5-subgraph

corresponds to a tangle. The following recent work studies or constructs trees of tangles [Alb23; Alb24a; Alb24b; Car16; CDHH14; CDHH16a; CDHH16b; CDHS14; CG17; CK24; DEW19; DHL19; EKT21; EKT22; Erd17; Gro16b; JK24] or even more specific decompositions [BGHKMR23; CK23a; EGL24]. Algorithmic applications include [EFKKRTL23; GS15a; Kor24], see [Die24] for an overview.

Intuitively, trees of tangles capture the connectivity structure of graphs, as their separations represent the 'bottlenecks' of the graph along which they decompose it into 'highly connected parts', which in turn fit together as described by the decomposition tree. However, there are many graphs for which their trees of tangles fail to detect much of their interesting connectivity structure. Indeed, when we perceive a graph to be formed from regions that are 'locally highly connected' by arranging these regions in a 'graph-like' way rather than a 'tree-like' one, then trees of tangles fail to recover this arrangement.

(A) Large sparse networks form a wide class of examples, including networks from applications in Data Science such as infrastructure networks or social networks. For example, the orange regions in Figure 2.2a are locally highly connected, and they are arranged in a 'grid-like' way.

(B) Cayley graphs of finite groups yield examples. The locally highly connected regions of the Cayley graph in Figure 2.2b are copies of K_4 arranged in a 'cycle-like' way.

(C) Just as a chordal graph is obtained from cliques by gluing them together in the shape of a tree, a *locally chordal* graph G consists of cliques arranged in a 'high-girth-like' way [AJKK25], see Figure 2.2c.

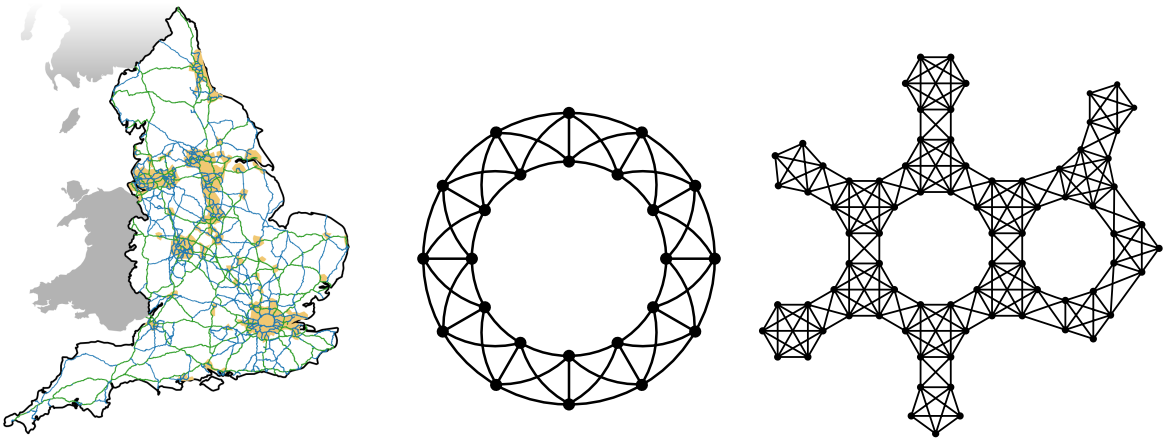

(a) The main road network of the UK

(b) A Cayley graph of $\mathbb{Z}_{12} \times \mathbb{Z}_2$

(c) A graph that resembles the molecule caffeine

Figure 2.2: Graphs whose connectivity-structure eludes trees of tangles

A recent development in graph minor theory, initiated by Diestel and Kühn [DK05] and followed up in [ADEFJKW23; AJKW24; DJKK22] and differently in [Car22], is the study

of *graph-decompositions*, which generalise tree-decompositions in that the decomposition tree may take the form of an arbitrary graph; see Figure 2.3. Diestel, Jacobs, Knappe and Kurkofka [DJKK22] have developed a 'graph-decomposition version' of the tree-of-tangles theorem.

We briefly recall the construction; see Section 2.2 for the precise definitions. Suppose we are given a connected graph G. Let $r > 0$ be an integer parameter that will quantify the meaning of 'local' in 'locally highly connected'. We construct a graph-decomposition $\mathscr{D}_r(G)$ of G in three steps.

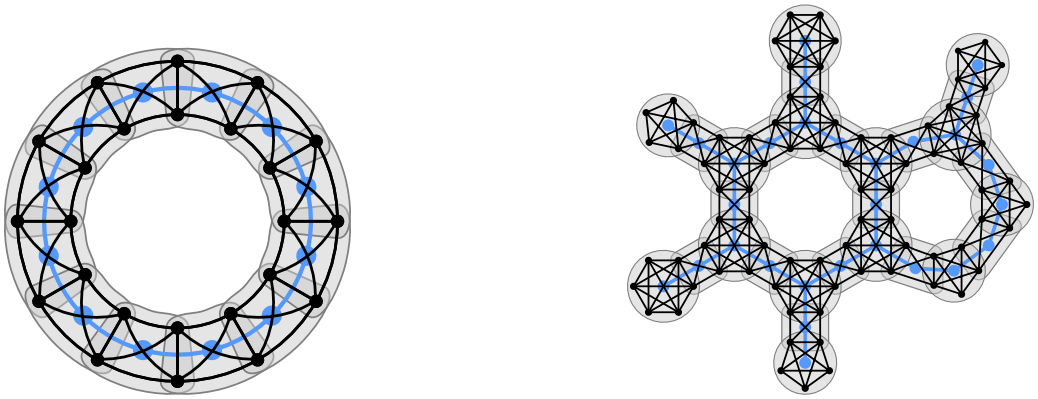

Figure 2.3: Graph-decompositions of black graphs into smaller parts (grey). The decomposition graphs are blue.

First, we take the *r-local covering* $p_r\colon G_r \to G$ of G introduced in [DJKK22], which is the unique covering of G that preserves all $(r/2)$-balls of G but unfolds all other structure in G as much as possible. The graph G_r is usually infinite, but it is highly symmetric; see Figure 2.4. Intuitively, the r-locally highly connected regions in G correspond to highly connected regions in G_r, and in G_r they are arranged in a 'tree-like' way due to the nature of coverings – so this arrangement is captured with a tree of tangles of G_r. In the second step, we apply a tree-of-tangles theorem for infinite graphs [CK24] to the covering G_r to obtain a tree of tangles of G_r that is even canonical (i.e. invariant under graph isomorphisms such as the symmetries of G_r). In the third and final step, the tree of tangles of G_r is 'folded back' via p_r to yield the graph-decomposition $\mathscr{D}_r(G)$ of G; see Construction 2.2.6. Following the above intuition, $\mathscr{D}_r(G)$ captures the arrangement of the r-locally highly connected regions in G.

The graph-decomposition $\mathscr{D}_r(G)$ is canonical, and unique given G and r. In the case of $r = 0$, the covering G_r is the universal one, which is a tree: then $\mathscr{D}_r(G)$ decomposes G into its edges and vertices, and the decomposition graph is the 1-subdivision of G. In the case of $r \geqslant |G|$, we have $G_r = G$, and so $\mathscr{D}_r(G)$ defaults to a tree of tangles of G.

The above construction of $\mathscr{D}_r(G)$ involves topology and infinite graphs. To show that graph-decompositions computable, and thus accessible to algorithms, we need to solve the following problem:

Problem 1. *Is there a finite combinatorial description of graph-decompositions like $\mathscr{D}_r(G)$?*

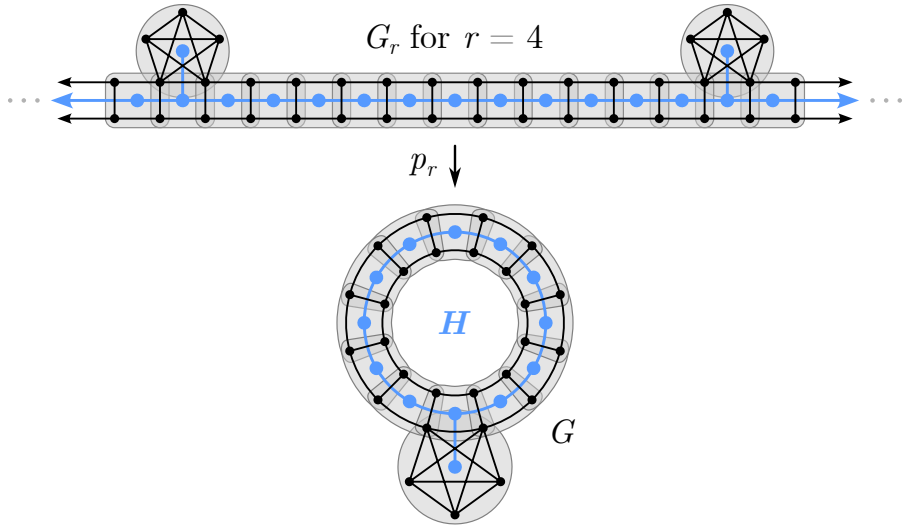

Figure 2.4: An H-decomposition of G (black) into 4-cycles and a K_5 (grey) over a cycle H (blue), obtained from a tree-decomposition of the r-local cover G_r for $r = 4$.

2.1.2 From separations to local separations

In a graph modelling a road network, small separators between large clusters form bottlenecks that often get congested by traffic. However, the same happens when such bottlenecks are only local: if all the *shortest* routes between two clusters lead through the bottleneck. Indeed, drivers will choose those short routes as long as the bottleneck through which they take them is not congested. Only then will it make sense to use a longer detour, even if such roads around bottlenecks exist.

For example, the two red vertices in Figure 2.5a block the two shortest routes between Birmingham (B) and Nottingham (N). Since the available detours are longer, they will not prevent congestion on the two shortest routes: only once these are congested will the alternative longer routes become faster. Thus local bottlenecks are real: they are places in the network that get congested first.

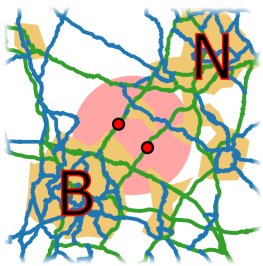

(a) A portion of the main road network of the UK.

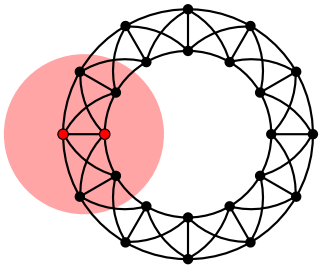

(b) The red r-local 2-separator marks a factor of a direct product.

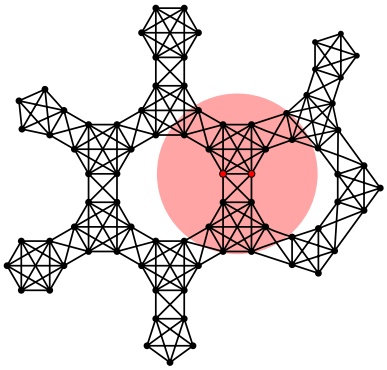

(c)

Figure 2.5: Local separators (red) that separate the reddish balls around them.

In this chapter we develop a formal way to study such local bottlenecks. Just as the edges of the decomposition tree of a tree-decomposition correspond to global separations of the underlying graph, the edges of the models H of the decompositions $\mathscr{D}_r(G)$ will correspond to its local separations.

Carmesin [Car22] initiated the study of r-local 1- and 2-separators, generalising the tree-like block-cutvertex decomposition and Tutte's 2-separator theorem [Tut84] to the r-local setting. Those early theorems provide canonical sets of pairwise non-crossing separations of order one and two, respectively, where the *order* of a separation is the size of its separator. The r-local ($\leqslant 2$)-separations from [Car22] have since found applications to sparse networks and finite groups:

ad (A): Carmesin and Frenkel [CF23] have developed an efficient algorithm for computing the canonical non-crossing r-local ($\leqslant 2$)-separations and tested it on road networks (Figure 2.5a). Their algorithm can use Parallel Computing, which is made possible by canonicity.

ad (B): They were used in [CKKT24] to prove a low-order Stallings-type theorem for finite nilpotent groups (Figure 2.5b).

Carmesin [Car22, §10] asked if his r-local ($\leqslant 2$)-separations have analogues of orders $k > 2$:

Problem 2. *Is there a decomposition theory for graphs based on r-local k-separations?*

Our aim in this chapter is to present one unified solution to both Problem 1 and Problem 2: we show that the graph-decomposition $\mathscr{D}_r(G)$ have a finite combinatorial description, by identifying the projections of ($\leqslant k$)-separations of the cover G_r as *r-local ($\leqslant k$)-separations* of G and characterising these in finite terms based on G alone.

2.1.3 Graph-decompositions and local separations

In this chapter, we develop a finite combinatorial theory of local separations, and we use this theory to come up with a finite combinatorial description of graph-decompositions like $\mathscr{D}_r(G)$. Our work is fully compatible with both the existing perspectives of r-local coverings [DJKK22] and of r-local ($\leqslant 2$)-separations [Car22].[1] This answers questions Problem 1 and Problem 2, up to the following caveat. While the concepts that we develop as the basis for our theory of r-local k-separations are defined for all k, the rest of the theory is only guaranteed to work for $k < K(G, r)$, where $K(G, r)$ is a graph invariant which we introduce next.

Given a graph G and an integer $r > 0$, we define the *displacement* $\Delta_r(G)$; see Section 2.5 for the definition. We guarantee that our theory works for r-local separations of order $< K(G, r)$ where $K(G, r) := \frac{\Delta_r(G)}{r} + 1$ (setting $K(G, r) := 2$ if $r = 0$). The value of $K(G, r)$ is quite large in many graphs, and we provide a construction scheme for examples supporting this claim, see

[1] We show this for r-local coverings in this chapter, whereas for r-local ($\leqslant 2$)-separations this will be proved in [BJKKP].

Example 2.5.13. Even in the worst case, the theory works for r-local ($\leqslant 2$)-separations, see Lemma 2.5.1.

The definition of the displacement $\Delta_r(G)$ involves the r-local covering, which is usually infinite, and so it is not clear whether there exist efficient algorithms that determine $\Delta_r(G)$. However, we only use the displacement to guarantee that our theory works for r-local k-separations if $k < K(G, r)$; examples suggest that the theory often works beyond this guarantee. Still, if determining $\Delta_r(G)$ becomes an issue, one can alternatively work with finite combinatorial lower bounds for $\Delta_r(G)$, such as the minimum length of an induced cycle of length $> r$ (Lemma 2.5.3).

2.1.4 From trees of tangles to graphs of local bottlenecks

The work [EKT21; EKT22] initiated a shift of perspective from tangles to the minimum-order separations that separate them. Recent research [CK24] has developed this perspective further by axiomatising these minimum-order separations through the new concept of *entanglements*, and then used entanglements to come up with a new canonical construction of trees of tangles.[2] As this construction works without referencing tangles, we refer to it as a *tree of entanglements*. The characteristic property of a tree of entanglements of G is that it *represents* all entanglements in G, that is, for every entanglement in G the decomposition tree has an edge that corresponds to some separation in that entanglement.

Hence, to find a finite combinatorial description of graph-decompositions like $\mathscr{D}_r(G)$, which are obtained from trees of tangles of the r-local covering G_r by folding them to graph-decompositions of G, one might try to come up with a notion of 'r-local entanglements' based on r-local separations, and then try to find a canonical construction of 'graphs of r-local entanglements'. Roughly speaking, a 'graph of r-local entanglements' of G would be a graph-decomposition of G that represents the r-local entanglements in G.

Entanglements, however, present a challenge from an algorithmic standpoint. To determine whether a given set of k-separations constitutes an entanglement, one must compare these k-separations with all separations of G across all orders, including those of order much larger than k. This requirement complicates iterative algorithmic processes, which typically seek to build a tree of entanglements layer by layer: by starting with examining all entanglements formed by ($\leqslant k$)-separations for $k = 1$, before progressively considering higher values of k. This requirement also causes a theoretical problem, as we can only guarantee that our finite combinatorial theory works as intended for r-local separations of order $< K(G, r)$.

To overcome the challenge associated with entanglements, we develop a new axiomatisation of the minimum-order separations between any two given tangles, which we call *bottlenecks*. All separations within a bottleneck share the same order, which is the *order* of the bottleneck. The main difficulty in conceptualising bottlenecks was to strike a delicate balance: On the one hand, bottlenecks need to be specific enough to make it possible for us to find an r-local version of them. In particular, whether a given set of k-separations forms a bottleneck of

[2] For ends instead of tangles, a similar construction but without entanglements had also been found by [Rüh17].

order k must only depend on their relationships with the $(\leqslant k)$-separations of G. On the other hand, bottlenecks need to be general enough so that, for any given two tangles, the set of minimum-order separations distinguishing them forms a bottleneck.

As a main result, we show that every finite connected graph G has for every r and $k < K(G,r)$ a *graph of r-local bottlenecks* $\mathcal{H}_r^{\leqslant k}(G)$. That is, $\mathcal{H}_r^{\leqslant k}(G)$ is a graph-decomposition of G that represents all the r-local bottlenecks β in G of order $\leqslant k$: for every β there is an edge of the decomposition graph H of $\mathcal{H}_r^{\leqslant k}(G)$ that corresponds to an r-local separation in β. The parameter $k < K(G,r)$ allows us to set the maximum order for all r-local separations and all r-local bottlenecks considered in the construction of $\mathcal{H}_r^{\leqslant k}(G)$. For fixed G and r, the decompositions $\mathcal{H}_r^{\leqslant 0}(G)$, $\mathcal{H}_r^{\leqslant 1}(G), \dots$ then refine their predecessors, and hence allow for iterative algorithmic constructions. The constructions may even continue for $k \geqslant K(G,r)$ if additional checks for the results are in place.

ad (C): The graph-decomposition $\mathcal{H}_r^{\leqslant k}(G)$ has been applied for values of k beyond the guarantee $K(G,r)$ to structurally characterise the r-locally chordal graphs G [AJKK25].

The graph-decompositions $\mathcal{H}_r^{\leqslant k}(G)$ are unique given G, r, k, and they are canonical.

If $r \geqslant |G|$, and more generally if G itself is an r-local covering of some graph, then $\mathcal{H}_r^{\leqslant k}(G)$ defaults to a tree of tangles of G. For all $r \in \mathbb{N}$, the graph-decomposition $\mathcal{H}_r^{\leqslant k}(G)$ coincides with the graph-decomposition obtained from the tree of tangles $\mathcal{H}_r^{\leqslant k}(G_r)$ of the r-local covering G_r by folding via p_r. In this sense, $\mathcal{H}_r^{\leqslant k}(G)$ exhibits the same qualities as the graph-decomposition $\mathscr{D}_r(G)$ that is defined via r-local coverings, up to the parameter k that bounds the order of the r-local separations considered in the construction of $\mathcal{H}_r^{\leqslant k}(G)$.

Theorem 4. *Let G be a finite connected graph and $r, k \in \mathbb{N}$ with $k < K(G,r)$.*

(i) *$\mathcal{H}_r^{\leqslant k}(G)$ is a canonical graph-decomposition of G that represents all the r-local bottlenecks in G of order $\leqslant k$.*

(ii) *$\mathcal{H}_r^{\leqslant k}(G_r)$ is a tree-decomposition induced by a tree-of-tangles of resolution k of the r-local covering G_r.*

(iii) *$\mathcal{H}_r^{\leqslant k}(G)$ is canonically isomorphic to the graph-decomposition of G that is obtained from the tree-decomposition $\mathcal{H}_r^{\leqslant k}(G_r)$ of the r-local covering G_r by folding via p_r.*

(iv) *$\mathcal{H}_r^{\leqslant k}(G)$ refines $\mathcal{H}_r^{\leqslant \ell}(G)$ for all $\ell \in \mathbb{N}$ with $\ell \leqslant k$.*

See Section 2.3 and Section 2.9 for definitions.

Theorem 5. *There is an algorithm that determines $\mathcal{H}_r^{\leqslant k}(G)$ given G, r, k as in Theorem 4.*

Theorems 4 and 5 together constitute one main result of this chapter. However we hope that the finite combinatorial theory of r-local k-separations will be of independent interest.

2.1.5 Organisation of the chapter

Section 2.2 recalls terminology about graph-decompositions. In Section 2.3, we introduce bottlenecks. A canonical version of the Tree-of-Tangles Theorem from [EKT22] can be generalised to bottlenecks. In Section 2.4, we introduce the notion of local separations and study their interplay with local coverings. In Section 2.5, we introduce the notion of displacement and study how large displacement can be harnessed. In Section 2.6, we expand the crossing relation from the usual separations to local separations and study its interplay with local coverings. In Section 2.7, we introduce a local version of bottlenecks, study its interplay with local coverings, and prove a local-separations version of the Tree-of-Tangles Theorem. In Section 2.8, we show how graph-decompositions can be built from sets of pairwise non-crossing local separations. In Section 2.9, we prove Theorems 4 and 5.

2.2 Review of graph-decompositions and coverings

Our notations and definitions follow [Die25]. Throughout this chapter, let G be a (simple and not necessarily finite) graph and $r \in \mathbb{N}$ an integer. When r is clear from context, which it usually will be, we refer to the cycles in G of length $\leqslant r$ as *short* cycles. Let X, Y, A be sets of vertices in G. An X–Y path P in G is called a path *through* A if all internal vertices of P are contained in A and P has at least one internal vertex. Recall that a *near-partition* of a set M is a family $(M_i : i \in I)$ such that the M_i are disjoint and $M = \bigcup_{i\in I} M_i$. If $|I| = 2$, we speak of a *near-bipartition*.

We remark that most of the following is a recap of definitions and concepts from [DJKK22] relevant to this chapter; in [DJKK22] you will find proofs or reference to proofs.

2.2.1 Local coverings

A *walk* (of *length* k) in a graph G is a non-empty alternating sequence $W = v_0e_0v_1 \dots v_{k-1}e_{k-1}v_k$ of vertices and edges in G such that each e_i joins v_i and v_{i+1}. If $v_0 = v_k$ we call W *closed* and say that W is a closed walk *at* v_0. The subsequences of the form $v_ie_i \dots v_j$ are called *subwalks* of W. The walk that traverses W backwards is denoted by W^-. We denote the set $\{v_0, \dots, v_k\}$ of vertices on the walk W by $V(W)$. A walk W is an X–Y *walk* for vertex-sets $X, Y \subseteq V(G)$ if W has precisely its first vertex in X and precisely its last vertex in Y. An X*-walk* is a walk that has precisely its first and last vertex in X.

We say that a walk W is *reduced* if it contains no subwalk of the form $ueveu$. The *reduction* of W is the unique walk obtained by iteratively replacing its subwalks of the form $ueveu$ with the trivial walk u. Two walks are *homotopic* if their reductions are equal. Let $\sim$ denote the equivalence relation this defines on the set of all walks in G.

For a vertex x_0 of G, we denote by $\mathcal{W}(G, x_0)$ the set of all closed walks in G at x_0. Then $\pi_1(G, x_0) := (\mathcal{W}(G, x_0)/\!\sim, \cdot)$ is the (*combinatorial*) *fundamental group* of G at x_0, where $[W_1] \cdot [W_2] := [W_1W_2]$ and W_1W_2 is the concatenation of the two walks W_1 and W_2. We often do not distinguish between W and $[W]$ when it is clear from context what we mean,

A closed walk W in G at x_0 *stems* from a closed walk Q if W can be written as $W = W_0QW_0^-$. A closed walk *once around* a cycle O is a closed walk in O that traverses every edge of O exactly once. We say that W *stems* from a cycle O if W stems from a closed walk Q once around O. The *r-local subgroup* $\pi_1^r(G, x_0)$ of $\pi_1(G, x_0)$ is the subgroup generated by the closed walks at x_0 that stem from short cycles in G (those of length $\leqslant r$). The r-local subgroup is normal.

A *covering* of a graph G is an epimorphism $p\colon C \to G$ for some graph C which exhibits the following property: for every vertex $v \in C$, the map p bijectively sends the edges incident to v to the edges incident to $p(v)$ in G. The vertices in $p^{-1}(v)$ are the *lifts* of v. Lifts of edges are defined similarly. A *lift* of a set X of vertices of G to G_r is a vertex-set $\hat{X} \subseteq V(G_r)$ that contains for each $x \in X$ exactly one lift $\hat{x}$ of x. Equivalently, p_r restricts to a bijection $\hat{X} \to X$. A *lift* of a walk W in G means a walk in C which p projects to W. There exists a covering $p_r\colon G_r \to G$ such that p_r induces an isomorphism from the fundamental group $\pi_1(G_r, \hat{x}_0)$ of G_r to the r-local subgroup $\pi_1^r(G, x_0)$, where $\hat{x}_0$ is some lift of x_0. The covering $p_r\colon G_r \to G$ is unique up to isomorphism given G and r, and it is independent of the choice of x_0. We call $p_r\colon G_r \to G$ the *r-local covering* of G. Recall that the 0-local, 1-local and 2-local covering of a simple graph G are all its universal covering. The group of *deck transformations* of G_r, those automorphisms of G_r whose action commutes with p_r, is denoted by $\Gamma_r(G)$.

Lemma 2.2.1. [DJKK22, Lemma 4.6] *The binary cycle space of G_r is generated by the short cycles.*

We will often use this lemma implicitly when applying lemmas to G_r.

2.2.2 Separations

Let G be any graph. For $X \subseteq V(G)$, a component K of $G - X$ is *tight* at X if $N_G(K) = X$. A *separation* of G is a set $\{A, B\}$ such that $A \cup B = V(G)$ and G contains no edge between $A \setminus B$ and $B \setminus A$. We refer to A and B as the *sides* of $\{A, B\}$, and we call $A \cap B$ the *separator* of $\{A, B\}$. The size $|A \cap B|$ of the separator is the *order* of $\{A, B\}$. Also we call a separation with separator X a *k-separation* if $k = |X|$. A separation $\{A, B\}$ is *proper* if $A \setminus B$ and $B \setminus A$ are non-empty. If there are tight components C_A, C_B of $G - (A \cap B)$ whose vertex sets are contained in A and B, respectively, then the separation $\{A, B\}$ is *tight*.

Every separation $\{A, B\}$ has two *orientations* (A, B) and (B, A), and we refer to (A, B) and (B, A) as *oriented separations*. We denote the *inverse* (B, A) of an oriented separation (A, B) by $(A, B)^*$. Given a set S of (unoriented) separtions, we denote the set of all orientations of the separations in S by $\mathrm{ori}(S)$. The *order* of (A, B) and (B, A) is the order of $\{A, B\}$. A partial ordering $\geqslant$ is defined on the ordered separations of G by letting[3]

$$(A, B) \geqslant (C, D) \quad :\Longleftrightarrow \quad A \subseteq C \quad \text{and} \quad B \supseteq D$$

[3]We remark that we use the partial order as defined in the 6th but not the 5th edition of Diestel's book [Die25]. That is, we write $\geqslant$ where readers familiar with previous literature would expect $\leqslant$.

Two separations $\{A, B\}$ and $\{C, D\}$ of G are *nested* if they have $\geqslant$-comparable orientations. Then also every orientation of $\{A, B\}$ is *nested* with every orientation of $\{C, D\}$. If $\{A, B\}$ and $\{C, D\}$ are not nested, then they are said to *cross*, and also every orientation of $\{A, B\}$ *crosses* every orientation of $\{C, D\}$. A set of separations of G is *nested* if its elements are pairwise nested.

The following definitions are supported by Figure 2.6. Let $\{A_1, A_2\}$ and $\{C_1, C_2\}$ be two separations of a graph G with separators X and Y, respectively. The *X-link for C_i* is the vertex-set $X \setminus C_{3-i}$ (which is included in C_i). Similarly, the *Y-link for A_i* is the vertex-set $Y \setminus A_{3-i} \subseteq A_i$. The two X-links are said to be *opposite* of each other, and similarly the Y-links are *opposite* of each other. The *centre* is the intersection $X \cap Y$. The separator X is near-partitioned into its two X-links and the centre.

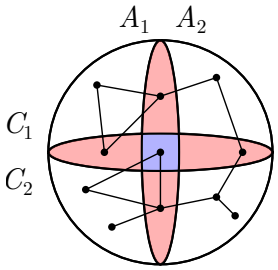

Figure 2.6: Links are red while the centre is blue

Lemma 2.2.2. *Let G be a connected graph, and let $\{A_1, A_2\}$ and $\{C_1, C_2\}$ be separations of G with separators X and Y, respectively. The following assertions are equivalent:*

(i) *$\{A_1, A_2\}$ and $\{C_1, C_2\}$ cross;*

(ii) *for every pair of indices $i, j \in [2]$, at least one of the following three sets is non-empty: the Y-link for A_i, the X-link for C_j, and $(A_i \setminus A_{3-i}) \cap (C_j \setminus C_{3-j})$.*

(iii) *for every pair of indices $i, j \in [2]$, at least one of the following three sets is non-empty: the Y-link for A_i, the X-link for C_j, and $E(X \cap Y, (A_i \setminus A_{3-i}) \cap (C_j \setminus C_{3-j}))$.*

Moreover, (ii) *and* (iii) *are violated for a pair of indices $i, j \in [2]$ if and only if $(A_i, A_{3-i}) \geqslant (C_{3-j}, C_j)$.*

Proof. (i)↔(ii) is [CDHS14, Statement (6)]. (iii)→(ii) is trivial. It remains (ii)→(iii).

For this, let $i, j \in [2]$ and assume that the Y-link for A_i and the X-link for C_j are empty. By (ii), there is a vertex v in $(A_i \setminus A_{3-i}) \cap (C_j \setminus C_{3-j})$. There is a vertex u of G outside of $(A_i \setminus A_{3-i}) \cap (C_j \setminus C_{3-j})$, since otherwise (ii) is violated for $i' := 3 - i$ and $j' := 3 - j$. Since G is connected, G contains a u–v path P. As $\{A_i \cap C_j, A_{3-i} \cup C_{3-j}\}$ is a separation, the path P must intersect its separator, which is equal to $X \cap Y$ by the assumption that the Y-link for A_i and the X-link for C_j are empty. Hence, the edge on P before its first vertex in $X \cap Y$ lies in $E(X \cap Y, (A_i \setminus A_{3-i}) \cap (C_j \setminus C_{3-j}))$.

The moreover-part follows immediately from the definition of $\geqslant$ and (ii). □

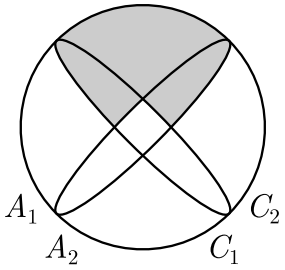

Figure 2.7: $(A_1, A_2) \geqslant (C_1, C_2)$ if and only if the grey area is empty. The grey area corresponds to $i = 1$ and $j = 2$ in Lemma 2.2.2 (ii)

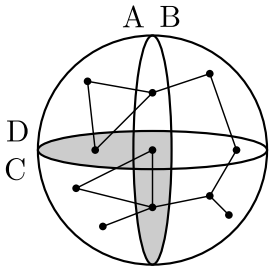

Figure 2.8: $\{A \cap C, B \cup D\}$ is one of the four corners of $\{A, B\}$ and $\{C, D\}$

For a depiction of the setting for the next definitions, see Figure 2.8. If $\{A, B\}$ and $\{C, D\}$ cross, then their four *corners* are the separations $\{A\cap C, B\cup D\}$, $\{A\cap D, B\cup C\}$, $\{B\cap D, A\cup C\}$ and $\{B \cap C, A \cup D\}$. The corners $\{A \cap C, B \cup D\}$ and $\{B \cap D, A \cup C\}$ are *opposite*, and so are the corners $\{A \cap D, B \cup C\}$ and $\{B \cap C, A \cup D\}$. Any two corners that are not opposite are *adjacent*. The two adjacent corners $\{A \cap C, B \cup D\}$ and $\{A \cap D, B \cup C\}$ are said to *lie on the same side* of $\{A, B\}$, namely *on the A-side*. Similarly, the other two adjacent corners $\{B \cap D, A \cup C\}$ and $\{B \cap C, A \cup D\}$ are said to *lie on the same side* of $\{A, B\}$, namely *on the B-side*.

The separations of G satisfy the following property (F), see for example [CK24, Lemma 3.2].

(F) For every two crossing separations s and s', every two opposite corners c and d of s and s', and every separation t, the following three assertions hold:

  (F1) If t crosses at least one of c and d, then t crosses at least one of s and s'.

  (F2) If t crosses both c and d, then t crosses both s and s'.

  (F3) Neither s nor s' crosses c or d.

The order-function $\{A, B\} \mapsto |A \cap B|$ is *subdmodular*[4] on the separations of G:

(†) If two separations s and s' cross and if c and d are opposite corners, then $|c|+|d| \leqslant |s|+|s'|$.

[4]This is compatible with how the term 'submodular' is used in the literature, as the set $\mathrm{ori}(N)$ equipped with $\geqslant$ forms a lattice in which the infimum and supremum of two oriented separations are given by orientations of opposite corners.

A set σ of oriented separations of G is a *star* if $(A,B) \geqslant (D,C)$ for every distinct two $(A,B),(C,D) \in \sigma$. Let N be a nested set of separations of G. A *splitting star* of N is a star σ such that $\{A,B\} \in N$ and $(B,A) \notin \sigma$ for every $(A,B) \in \sigma$ and such that every $\{C,D\} \in N$ has an orientation (C,D) with $(C,D) \geqslant (A,B)$ for some $(A,B) \in \sigma$. Let $\mathrm{int}(\sigma)$ be the *interior*[5] $G[\bigcap_{(A,B)\in\sigma} B]$ of a splitting star σ of N.

Note that every $(A,B) \in \mathrm{ori}(N)$ is contained in at most one splitting star. Moreover, if all separations in N are proper and $\mathrm{ori}(N)$ has no $(\omega+1)$-chain, Kneip and Gollin [GK21, Corollary 3.3] showed that the splitting stars of N even form a partition of $\mathrm{ori}(N)$.

2.2.3 Decompositions

Let G be a graph, and H a multigraph. Let $\mathcal{G} = (G_h : h \in V(H))$ be a family of subgraphs $G_h \subseteq G$ indexed by the nodes of H. The pair $(H,\mathcal{G})$ is an *H-decomposition* or *graph-decomposition* of G if

(H1) $G = \bigcup_{h\in V(H)} G_h$, and

(H2) for every $v \in V(G)$, the subgraph $H_v \subseteq H$ induced on $\{h \in V(H) : v \in V(G)\}$ is connected.

We refer to the G_h as its *parts* and to H as its *decomposition graph*. Sometimes we also say that H *models* $(H,\mathcal{G})$. We often introduce graph-decompositions as pair $(H,\mathcal{G})$ and tacitly assume that $\mathcal{G} = (G_h : h \in V(H))$. If H is a tree and the G_h are induced subgraphs of G, then $(H,\mathcal{G})$ is a *tree-decomposition* of G. We often introduce tree-decompositions as pair $(T,\mathcal{V})$ and tacitly assume that $\mathcal{V} = (V_t : t \in V(T))$, to better distinguish them from the more general graph-decompositions.[6] For a category-theoretic perspective on graph-decompositions and tree-decompositions see [BFT24; BKM22].

Construction 2.2.3 (Tree-decompositions from nested sets separations)**.** Let G be a graph with a nested set N of proper separations. We define the *candidate* $\mathcal{T}(N) = (T,\mathcal{V})$ for a tree-decomposition, as follows. The nodes of T are the splitting stars of N and the edges of T are the separations in N, where $\{A,B\} \in N$ joins the two unique nodes t_1, t_2 of T with $(A,B) \in t_1$ and $(B,A) \in t_2$. The part V_t of a node t of T is $\mathrm{int}(t)$.

We remark that the candidate $\mathcal{T}(N)$ is not necessarily a tree-decomposition of a given connected, locally finite graph G [EKT22, Example 4.9].

Lemma 2.2.4. [DJKK22, Lemma 7.5] *Assume that G is connected and locally finite. Let N be a nested set of tight separations of G which all have order $\leqslant k$ for some $k \in \mathbb{N}$. Then $\mathcal{T}(N)$ is a tree-decomposition.*

[5] Many (if not almost all) authors define the interior of a splitting star as just the vertex-set $\bigcap_{(A,B)\in\sigma} B$. In this chapter, especially in Section 2.8, it will be extremely convenient to define it as the induced subgraph instead.

[6] Many (if not almost all) authors define tree-decompositions so that their parts are sets of vertices rather than induced subgraphs. The two definitions are equivalent. We have opted to work with induced subgraphs, because this will be extremely convenient in this chapter, especially in Section 2.8.

For a tree T, we denote by $\mathrm{ori}(E(T))$ the set of all orientations (u,v), (v,u) of the edges $e = uv$ of T. The *tree-order* on $\mathrm{ori}(E(T))$ is given by letting $e_1 \geqslant e_2$ for $e_i = (u_i, v_i) \in \mathrm{ori}(E(T))$ and $i \in [2]$ if $e_1 = e_2$ or the (unique) v_1–u_2 path in T avoids u_1 and v_2.

Let $(T, \mathcal{V})$ be a tree-decomposition of a graph G. Each oriented edge $e = (t_1, t_2) \in \mathrm{ori}(E(T))$ corresponds to a separation (A_1^e, A_2^e) of G where $A_i^e := \bigcup_{t \in V(T_i)} V(V_t)$ and T_i is the component of $T - t_1t_2$ which contains t_i, see [Die25, Lemma 12.3.1].

Lemma 2.2.5. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N) =: (T, \mathcal{V})$ is a tree-decomposition of G. Then $e \mapsto (A_1^e, A_2^e)$ is an order-isomorphism from $(\mathrm{ori}(E(T)), \geqslant)$ to $(\mathrm{ori}(N), \geqslant)$.*

Proof. This is well known; see for example [EKT22, Proof of Lemma 2.7]. □

An *action* of a subgroup Γ of $\mathrm{Aut}(G)$ on a graph-decomposition $(H, \mathcal{G})$ of G is an action of Γ on the graph H that commutes with $h \mapsto G_h$, i.e. that satisfies $\varphi(G_h) = G_{\varphi(h)}$ for all $h \in V(H)$. If such an action exists, then $(H, \mathcal{G})$ is Γ-*canonical*. In the case $\Gamma = \mathrm{Aut}(G)$, we just say that $(H, \mathcal{G})$ is *canonical*.

If $\mathcal{T}(N)$ is a tree-decomposition and N is Γ-invariant, then $\mathcal{T}(N)$ is Γ-canonical by construction. Moreover, every automorphism φ of G commutes with at most one automorphism ψ of the decomposition tree T in that $\varphi(V_t) = V_{\psi(t)}$ [DJKK22, Lemma 3.7]. So if $(T, \mathcal{V})$ is Γ-canonical, then its Γ-canonicity is witnessed by a unique action of Γ on T which commutes with $t \mapsto V_t$.

Construction 2.2.6 (Graph-decompositions via local coverings [DJKK22, Construction 3.8])**.**

Let G be a connected graph and let $r \in \mathbb{N}$. Let the group $\Gamma_r := \Gamma_r(G)$ of deck transformations of the r-local covering of G act on a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G_r. We define the graph-decomposition $\mathcal{H} = (H, \mathcal{G})$ with family $\mathcal{G}$ of parts G_h, as follows. The decomposition graph H is the orbit-graph of T under the action of Γ_r on $(T, \mathcal{V})$, that is the vertex and edge set of H consist of the Γ_r-orbits of the vertex and edge set of T where the endvertices of the orbit of $e = st$ are the orbits of s and t. The part G_h of a node h of H is defined as $G_h := p_r(V_t)$ for $t \in h$. We say that $(H, \mathcal{G})$ is *defined* by $(T, \mathcal{V})$ *via* p_r, or that $(H, \mathcal{G})$ is *obtained* from $(T, \mathcal{V})$ by *folding via* p_r.

Pairs $(H, \mathcal{G})$ from Construction 2.2.6 are graph-decompositions by [DJKK22, Lemma 3.9]. We remark that even though G is simple, H does not have to be simple.

2.3 Bottlenecks

In this section we introduce *bottlenecks*, which will later become one of our main tools to transfer the Tree-of-Tangles Theorem to local separations (see Section 2.7). Our main result of this section, Theorem 2.3.7, asserts that every locally finite and connected graph admits a

nested set of separations meeting all bottlenecks. This implies the existence of trees-of-tangles, see Theorem 2.3.12 and Theorem 2.3.13.

Bottlenecks build on a recent trend to construct trees-of-tangles by working with the minimum-order separations between the tangles instead of with the tangles themselves. This shift of perspective was initiated by [EKT21; EKT22] and led to a first axiomatisation of such sets of separations through *entanglements* [CK24]. With bottlenecks, we here give a second, alternative axiomatisation, which in contrast to entanglements allows for an iterative construction of a tree-of-tangles, see Construction 2.3.6. This in turn makes possible the formulation of an r-local version that is compatible with the guarantee $K(G, r)$, see Section 2.7. We do not develop an r-local version of tangles, as the need did not arise in this project. Still, we think it would be exciting to have an r-local version of tangles.

2.3.1 Bottlenecks

A $\top$-*star* of G (with $\top$ pronounced like the letter T) is a set $\sigma = \{ (A_i, B_i) : i \in [3] \}$ of three distinct oriented separations (A_i, B_i) of G with separators X_i such that

(Y1) the sets $A_i \smallsetminus B_i$ for $i \in [3]$ form a near-partition of the vertex-set $V(G) \smallsetminus (\bigcup_{i\in[3]} X_i)$, and

(Y2) every vertex in some X_i is contained in at least one other X_j.

Note that $\top$-stars are stars. The unoriented separations $\{A_i, B_i\}$ are the *constituents* of σ. The $\top$-star σ has *order* $\leqslant k$ if all its elements have order $\leqslant k$. We remark that $\bigcup_{i\in[3]} X_i$ has size $\leqslant \lfloor 3k/2 \rfloor$ if σ has order $\leqslant k$. If a $\top$-star σ is introduced as $\{ (A_1, B_1), (A_2, B_2), (A_3, B_3) \}$, then we denote by X_i the separator $A_i \cap B_i$ of $\{A_i, B_i\}$, let X be their union $\bigcup_{i\in[3]} X_i$, let Z be the *centre* $\bigcap_{i\in[3]} X_i$ of σ, and denote the *ij-link* $(X_i \cap X_j) \smallsetminus Z$ by X_{ij} for every two distinct indices $i, j \in [3]$.

These definitions are inspired by their usual occurrence as halves, or T, of a corner diagram:

Example 2.3.1. This example is supported by Figure 2.9. Let $\{A, B\}$ and $\{C, D\}$ be crossing separations of a graph G with separators X and Y, respectively. Let $(A_1, B_1) := (A, B)$ and let $(A_2, B_2) := (B \cap C, A \cup D)$ and $(A_3, B_3) := (B \cap D, A \cup C)$ be the two corners on the B-side. Then $\sigma := \{ (A_i, B_i) : i \in [3] \}$ is a $\top$-star of G. The 12-link of σ is equal to the X-link for C. Then 13-link of σ is equal to the X-link for D. The 23-link of σ is equal to the Y-link for B. The centre of σ is equal to the centre $X \cap Y$ of the two crossing separations.

We remark that (Y2) implies

$$X_i = X_{ij} \sqcup Z \sqcup X_{ik} \text{ for every } i \in [3], \text{ where } j \text{ and } k \text{ are the other two indices in } [3] \quad (2.1)$$

Let $X, Y \subseteq V(G)$ be disjoint from $A \subseteq V(G)$. A $\top$-star $\sigma = \{ (A_1, B_1), (A_2, B_2), (A_3, B_3) \}$ is *relevant with base* $\{A_1, B_1\}$ if $\{A_1, B_1\}$ is tight and for every $x \in X_{23}$ and every $i \in \{2, 3\}$ there exists an x–X_{1i} edge in G or there exist $y \in X_{1i} \cup Z$ and a component K of $G[A_i \smallsetminus B_i]$ with $x, y \in N(K)$; see Figure 2.10.

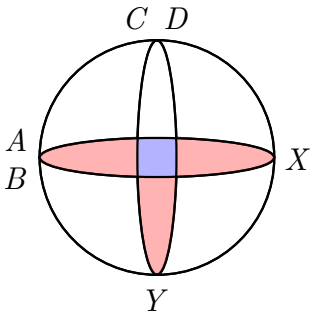

Figure 2.9: The situation in Example 2.3.1: the ij-links are red, the centre of σ is blue

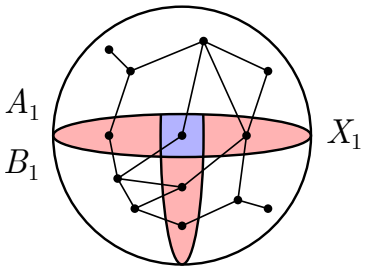

Figure 2.10: (A_1, B_1) is part of a $\top$-star that is relevant with base $\{A_1, B_1\}$. The links are red and the centre is blue. The red link on the right is the 23-link

Lemma 2.3.2. *Let $\{A, B\}$ and $\{C, D\}$ be crossing tight separations of a connected graph G, with separators X and Y, respectively. Assume that the X-link for C is empty while the Y-link for B is non-empty. Then the Y-link for A is empty while $X \cap Y$ is non-empty.*

Proof. Since $\{C, D\}$ is tight, there is a component $K \subseteq G[C \setminus D]$ with $N(K) = Y$. Since the X-link for C is empty, and since K has a neighbour in the non-empty Y-link for B, we get that K avoids A. So $Y = N(K)$ is included in $Y \cap B$, and in particular the Y-link for A is empty. Since $\{A, B\}$ and $\{C, D\}$ cross, and since the X-link for C and the Y-link for A are empty, the centre $X \cap Y$ is non-empty by Lemma 2.2.2. □

Lemma 2.3.3. *Let $\{A, B\}$ and $\{C, D\}$ be two tight separations of a connected graph G. Then the two corners $(B \cap C, A \cup D)$ and $(B \cap D, A \cup C)$ on the B-side together with (A, B) form a relevant $\top$-star with base $\{A, B\}$.*

Proof. Let $(A_1, B_1) := (A, B)$, $(A_2, B_2) := (B \cap C, A \cup D)$ and $(A_3, B_3) := (B \cap D, A \cup C)$. Then $\sigma := \{ (A_i, B_i) : i \in [3] \}$ is a $\top$-star by Example 2.3.1. Let $X_i := A_i \cap B_i$ for all i. We will implicitly use the correspondences between the links and centres provided by Example 2.3.1.

It remains to show that σ is relevant with base $\{A_1, B_1\}$. We may assume that X_{23} is non-empty, as we are done otherwise. Let $x \in X_{23}$ and $i \in \{2, 3\}$, say $i = 2$, as the other case follows by symmetry. We have to show that G contains either an x–X_{12} edge or that there are a vertex $y \in X_{12} \cup Z$ and a component K of $G[A_2 \setminus B_2]$ with $x, y \in N(K)$.

Since $\{C, D\}$ is tight, there exists a component K of $G - (C \cap D)$ with $V(K) \subseteq C \setminus D$ and $N(K) = C \cap D$. Recall that $Z = \bigcap_{i \in [3]} X_i$, and note that $Z = A \cap B \cap C \cap D$ by the definition

of σ. Hence, $N(K) = X_{23} \cup Z \cup Y$, where Y is the link $(C \cap D) \smallsetminus B$. We distinguish three cases.

First, we assume that X_{12} is empty. Since X_{23} is non-empty, Lemma 2.3.2 gives that Z is non-empty while Y is empty. Since K is included in $G[C \smallsetminus D]$ and has a neighbour $x \in X_{23}$, it must also be included in $G[A_2 \smallsetminus B_2]$. Thus, K is a component of $G[A_2 \smallsetminus B_2]$ with a neighbour in Z.

Second, we assume that X_{12} is non-empty and Y is non-empty. Since $x \in X_{23} \subseteq B \smallsetminus A$ and $Y \subseteq A \smallsetminus B$, there is an x–Y path Q through $K \subseteq G[C \smallsetminus D]$ and Q meets $(A \cap B) \smallsetminus D = X_{12}$. Let P be the x–X_{12} path in Q. Then P either is an x–X_{12} path through $A_2 \smallsetminus B_2$ or contains an x–X_{12} edge. In the former case, the internal vertices of P are contained in a desired component of $G[A_2 \smallsetminus B_2]$.

Finally, we assume that X_{12} is non-empty while Y is empty. Then Z is non-empty by Lemma 2.3.2. Consider an x–Z path Q through the component $K \subseteq G[C \smallsetminus D]$. Let P be the x–$(X_{12} \cup Z)$ subpath of Q. If P has no internal vertex, then P ends in X_{12} because Q goes through K. Hence, P either contains an x–X_{12} edge or is an x–$(X_{12} \cup Z)$ path through $A_2 \smallsetminus B_2$. In the latter case, the internal vertices of P are contained in a desired component of $G[A_2 \smallsetminus B_2]$. □

A *bottleneck of order* $k \in \mathbb{N}$ in G is a non-empty set β of tight k-separations of G such that β satisfies (B):

(B) Whenever β contains a base of a relevant $\top$-star σ of G of order $\leqslant k$, at least one of the other two constituents of σ is contained in β as well.

A bottleneck of order k is also called a *k-bottleneck* for short.

Recall that a separation *distinguishes* two tangles if the two tangles orient the separation to different sides. A separation distinguishes two given tangles *efficiently* if it distinguishes them with minimal order. Two tangles in G are *($\leqslant k$)-distinguishable* if there is a separation of G of order $\leqslant k$ that distinguishes the two tangles. Two tangles in G are *distinguishable* if some separation distinguishes them.

Example 2.3.4. Every pair of distinguishable tangles τ, τ' in a graph G induces a bottleneck $\beta = \beta(\tau, \tau')$ in G, which consists of the separations that efficiently distinguish τ and τ'; see [Die25, §12.5] for definitions regarding tangles.

Proof. As τ and τ' are distinguishable, there is a k-separation of G that efficiently distinguishes τ and τ'. Hence, β is non-empty, and every separation in β has order k. We claim that β is k-bottleneck. By [EKT22, Lemma 6.1 and its proof], every separation in β is tight. To show that β satisifes (B), let $\sigma = \{ (A_i, B_i) : i \in [3] \}$ be a relevant $\top$-star of G of order $\leqslant k$ with base $\{A_1, B_1\} \in \beta$. Then τ and τ' orient $\{A_1, B_1\}$ differently; we may assume that τ towards its A_1-side, and τ' towards its B_1-side, say. Both tangles τ, τ' also orient the other constituents of σ. By the tangle-property, there is $j \in \{2, 3\}$ such that τ' orients $\{A_j, B_j\}$ towards its A_j-side. Moreover, τ orients $\{A_j, B_j\}$ towards its B_j-side by consistency. Hence,

$\{A_j, B_j\}$ distinguishes τ and τ', and it must do so efficiently since $\{A_j, B_j\}$ has order $\leqslant k$. Therefore, $\{A_j, B_j\} \in \beta$. □

Recall that a separation $\{A, B\}$ of G *distinguishes* two sets $Y, Z \subseteq V(G)$ if, up to renaming the sides A and B, it satisfies that $Y \subseteq A$, Y meets $A \setminus B$, $Z \subseteq B$ and Z meets $B \setminus A$. A separation distinguishes two given sets *efficiently* if it distinguishes them with minimal order.

Example 2.3.5. Let Y and Z be two cliques in a graph G which are distinguished by a separation of G of order $< \min\{\,|Y|, |Z|\,\}$. Then Y and Z induce a bottleneck $\beta = \beta(Y, Z)$ in G, which consists of the separations of G that efficiently distinguish Y and Z.

Proof. The set β is non-empty by assumption. Denote by k the minimal order of a separation distinguishing Y and Z. Then all separations in β have order k. We claim that every separation in β is tight. Indeed, every separation $\{A, B\}$ in β separates each vertex in $Y \cap (A \setminus B)$ from each vertex in $Z \cap (B \setminus A)$ order-minimally, and both these sets are non-empty by assumption. Thus, every $\{A, B\}$ in β is tight.

It remains to show that β satisfies (B). For this, let $\sigma = \{\,(A_i, B_i) : i \in [3]\,\}$ be a relevant $\top$-star of G of order $\leqslant k$ with base $\{A_1, B_1\} \in \beta$. As the base $\{A_1, B_1\}$ is in β, we may assume by symmetry that $Y \subseteq A_1$, Y meets $A_1 \setminus B_1$, $Z \subseteq B_1$ and Z meets $B_1 \setminus A_1$. Then, for $i = 2, 3$, (Y1) ensures that $Y \subseteq B_i$, and thus Y meets $B_i \setminus A_i$ as $k < |Y|$. Note that for $i = 2, 3$, either $Z \subseteq A_i$ or $Z \subseteq B_i$ because Z is a clique.

We are done if $Z \subseteq A_i$ holds for some $i \in \{2, 3\}$. Indeed, as the order of (A_i, B_i) is $\leqslant k < |Z|$, the clique Z then meets $A_i \setminus B_i$, and thus $\{A_i, B_i\}$ efficiently separates Y and Z and is hence contained in β.

So suppose for a contradiction that $Z \subseteq B_i$ holds for both $i = 2, 3$. Then the union $\bigcup_{i=1}^{3} X_i$ of the separators $X_i := A_i \cap B_i$ of the separations (A_i, B_i) in σ contains Z. As there are no edges between any two links of σ and Z is a clique, at most one of the links of σ meets Z. Moreover, by (Y2), $\bigcup_{i=1}^{3} X_i$ equals the union of the $\top$-star's centre and its links. Hence, Z is contained in some X_i, which contradicts $|X_i| \leqslant k < |Z|$. □

2.3.2 A tree of bottlenecks

The following construction transfers the construction introduced in the proof of [EKT22, Theorem 1.2] to bottlenecks:

Construction 2.3.6 (Nested separations from bottlenecks)**.** Let G be a graph. Let $\mathcal{B}^k(G)$ be the set of all k-bottlenecks in G and $X^k := \bigcup \mathcal{B}^k(G)$ for every $k \in \mathbb{N}$. We recursively define for every $k \in \mathbb{N}$ a set $N^k(G)$, as follows. Let $x^k \colon X^k \to \mathbb{N}$ assign to each separation $s \in X^k$ the number of separations in X^k that s crosses, and assume that x^k is well-defined. Write $N^{<k}(G) := \bigcup_{j<k} N^j(G)$. For each $\beta \in \mathcal{B}^k(G)$, let $N^k(G, \beta)$ consist of the separations $s \in \beta$ which are nested with $N^{<k}(G)$ and among those have minimal $x^k(s)$. Finally, set $N^k(G) := \bigcup \{\, N^k(G, \beta) : \beta \in \mathcal{B}^k(G)\,\}$. We further write $N(G) := \bigcup_{k \in \mathbb{N}} N^k(G)$.

Theorem 2.3.7. *Let G be a locally finite, connected graph. Then the sets $N(G)$ and $N^{<k}(G)$ (for $k \in \mathbb{N}$) are well-defined and nested, and $N^k(G, \beta)$ is non-empty for every $k \in \mathbb{N}$ and every k-bottleneck β in G.*

Lemma 2.3.8. *Let G be a locally finite, connected graph and $k \in \mathbb{N}$. Then every tight finite-order separation of G is crossed by only finitely many tight separations of G of order at most k.*

Proof. This follows from a lemma of Halin [Hal91a, p. 2.4] as in the proof of [EKT22, Proposition 6.2]. □

Lemma 2.3.9. *Let β_i be k_i-bottlenecks for $i = 1, 2$ in a locally finite, connected graph G (possibly $\beta_1 = \beta_2$). If $s_1 \in \beta_1$ crosses $s_2 \in \beta_2$, then at least three corners have order $\leqslant \max\{k_1, k_2\}$.*

Proof. The short proof is analogue to that of [CK24, Sublemma 3.6], but we include it for convenience. Without loss of generality we have $k_1 \leqslant k_2$. Suppose for a contradiction that at most two corners of s_1 and s_2 have order $\leqslant k_2$. By submodularity (†), at least one of every two opposite corners has order $\leqslant k_2$. So there are exactly two corners c, d of order $\leqslant k_2$, and these are adjacent. Let c' and d' be the corners opposite of c and d, respectively. By assumption, c' and d' have order $> k_2$. Then c and d have order $< k_1$ by submodularity (†). Let i be the index for which c and d lie on the same side of s_i. Then c, d and s_i are the constituents of a relevant $\top$-star with base s_i of order $\leqslant k_i$ by Lemma 2.3.3. Since $s_i \in \beta_i$, (B) yields that either c or d is contained in β_i, say $c \in \beta_i$. But then c has order equal to k_i, contradicting that we have showed above that c has order $< k_1$. □

Lemma 2.3.10. *For every k-bottleneck β in a locally finite, connected graph G, the set $N^k(G, \beta)$ is non-empty.*

Proof. It suffices to show that β contains separations that are nested with $N^{<k}(G)$ for $k > 0$. Let $s \in \beta$ cross as few separations in $N^{<k}(G)$ as possible. Lemma 2.3.8 ensures that these are only finitely many. Assume for a contradiction that s crosses some $s' \in N^{<k}(G)$. Then at least three corners have order $\leqslant k$ by Lemma 2.3.9. Thus, two corners on the same side of s have order $\leqslant k$. By Lemma 2.3.3, these two corners and s form a relevant $\top$-star σ based at s. Hence, (B) applied to σ yields a corner $c \in \beta$. The corner c crosses strictly fewer separations in $N^{<k}(G)$ than s by (F). □

Lemma 2.3.11. *Let G be a locally finite, connected graph. Let s, s' be two crossing separations in tight k-bottlenecks in G, and let c, d be two opposite corners of s, s' that are contained in $\bigcup \mathcal{B}^k(G)$. Then*

$$x^k(c) + x^k(d) < x^k(s) + x^k(s'). \tag{2.2}$$

Proof. This follows from (F) like in the proof of [CK24, Corollary 3.3]. □

Proof of Theorem 2.3.7. Lemma 2.3.8 ensures that x^k in Construction 2.3.6 is well-defined. The sets $N^k(G, \beta)$ are non-empty by Lemma 2.3.10. To show that $N(G)$ and the sets $N^{<k}(G)$

are nested, it suffices to show that $N^k(G)$ is nested for every $k \in \mathbb{N}$. Suppose for a contradiction that there are two separations $s, s' \in N^k(G)$ that cross. Then $s \in N^k(G, \beta)$ and $s' \in N^k(G, \beta')$ for some k-bottlenecks β, β' in G. Since s and s' are nested with $N^{<k}(G)$, (F1) implies that

$$\textit{all four corners of } s, s' \textit{ are nested with } N^{<k}(G). \tag{2.3}$$

By Lemma 2.3.9, at least three corners have order $\leqslant k$. Without loss of generality, the corners c, c', d' depicted in Figure 2.11 have order $\leqslant k$.

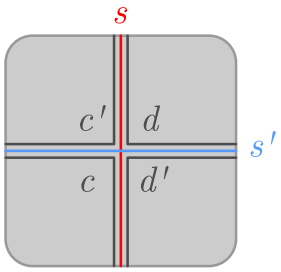

Figure 2.11: The four corners in the proof of Theorem 2.3.7

Case 1: The corner d has order $> k$. By Lemma 2.3.3 and (B), β meets $\{c, c'\}$ and β' meets $\{c, d'\}$. If c is contained in β or β', then c has order k, and then d has order $\leqslant k$ by submodularity (†) contradicting our assumption; hence, c lies in neither β nor β'. We therefore get $c' \in \beta$ and $d' \in \beta'$. But $x^k(c') + x^k(d') < x^k(s) + x^k(s')$ from (2.2) implies that $x^k(c') < x^k(s)$ or $x^k(d') < x^k(s')$, say $x^k(c') < x^k(s)$. This together with (2.3) shows that c' contradicts the fact that s lies in $N^k(G, \beta)$.

Case 2: The corner d has order $\leqslant k$. We apply Lemma 2.3.3 and (B) to both sides of both $s \in \beta$ and $s' \in \beta'$. If for a pair of opposite corners one is in β and the other in β', then we obtain a contradiction from the strict submodularity of x^k (2.2) as in Case 1. So we may assume that this does not happen. Without loss of generality, $c \in \beta$. Then $c' \in \beta'$, which in turn gives $d \in \beta$, and finally $d' \in \beta'$. Without loss of generality, we have $x^k(s) \leqslant x^k(s')$. Then $x^k(c') + x^k(d') < x^k(s) + x^k(s')$ from (2.2) implies that $x^k(c') < x^k(s')$ or $x^k(d') < x^k(s')$. This in combination with (2.3) shows that one of $c', d' \in \beta'$ contradicts the fact that s' lies in $N^k(G, \beta')$. □

A tree-decomposition $\mathcal{T}$ of G is a *tree-of-tangles* of G of *resolution* k if, for every pair of $(\leqslant k)$-distinguishable tangles τ_1, τ_2 in G, there is an induced separation of $\mathcal{T}$ that efficiently distinguishes τ_1 and τ_2.

Theorem 2.3.12. *Let G be a locally finite connected graph, and let $k \in \mathbb{N}$. Then $\mathcal{T}(N^{\leqslant k}(G))$ is a tree-of-tangles of G of resolution k.*

Proof. We combine Theorem 2.3.7 with Lemma 2.2.4 and Example 2.3.4. □

A tree-of-tangles that has resolution k for every $k \in \mathbb{N}$ has *infinite* resolution.

Theorem 2.3.13. *Let G be finite connected graph. Then $\mathcal{T}(N(G_r)) =: \mathcal{T}$ is a tree-of-tangles of the r-local covering G_r of infinite resolution.*

Proof. To see that $\mathcal{T}$ is indeed a tree-decomposition, we combine Lemma 2.2.4 with the following results of [DJKK22]: Theorem 3, Theorem 6.2, Lemma 6.3 and Lemma 8.1. By Example 2.3.4, $\mathcal{T}$ is even a tree-of-tangles. □

2.4 Local separations

In this section, we define *r-local separations* formally, along with notions of *r-local separators* and *r-local components*. We then show that these definitions are compatible with the perspective of r-local coverings: If the binary cycle space of a graph is generated by cycles of length $\leqslant r$, such as in r-local coverings, then there is a direct 1-to-1 correspondence between the global separators/components/separations and their r-local counterparts (Correspondences 2.4.2, 2.4.5 and 2.4.6). Finally, we study the interplay of these concepts with the covering map p_r and prove lifting and projection lemmas for r-local components (Projection 2.4.12 and Lift 2.4.14) and r-local separations (Projection 2.4.16 and Lift 2.4.17).

2.4.1 Local separations and correspondence in the local cover

Let G be a connected graph and $X \subseteq V(G)$. Recall that $\partial_G(X)$ or ∂X denotes the set of all edges in G with precisely one end in X. The cuts of the form $E_G(C, X)$ for a component C of $G - X$ are the *componental cuts at X* in G. The set of components C of $G - X$ bijectively corresponds to the partition of ∂X into the componental cuts $E_G(C, X)$, as G is connected. This partition is also defined by the equivalence relation on ∂X where we declare two edges $e, f \in \partial X$ to be equivalent if there exists an X-walk in G which starts with e and ends with f. The notion of *r-local components* below is motivated by the perspective of this equivalence relation.

An *X-edge* means an edge with both endvertices in X. An *r-local X-walk* in G is an X-walk in some short cycle in G or an X-walk consisting of a single X-edge; an r-local X-walk which is a path is called an *r-local X-path*. We remark that a non-trivial reduced r-local X-walk is either an r-local X-path or a walk once around a short cycle that meets X in precisely its base vertex. The *r-local components at X* in G are the equivalence classes of the transitive closure of the relation on ∂X defined by letting

$$e \sim_r f \text{ at } X \text{ in } G :\Leftrightarrow e \text{ and } f \text{ are the first and last edge on the same } r\text{-local } X\text{-walk in } G, \text{ or } e = f.$$

It is immediate from the definition that every r-local component at X is included in some componental cut at X. We further note that every r-local X-walk meets every r-local component at X in precisely 0 or 2 edges, and thus every short cycle meets every r-local component at X in an even number of edges.

We call $X \subseteq V(G)$ an *r-local separator* of G if there are at least two r-local components at X. Every separator of G is an r-local separator. A new example is depicted in Figure 2.12. The size of an r-local separator is its *order*. An r-local separator of order k is also called an *r-local k-separator.*

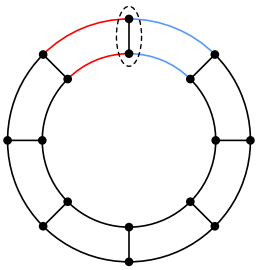

Figure 2.12: There are two 4-local components (red and blue) at X (circled)

An r-local component F at X is *tight* if every vertex in X is incident to some edge in F. We say that an r-local separator is *tight* if there are at least two tight r-local components at X.

Example 2.4.1. Even though we have noted above that a separator X is also an r-local separator, X being tight as a separator need not imply that X is tight as an r-local separator, see Figure 2.13.

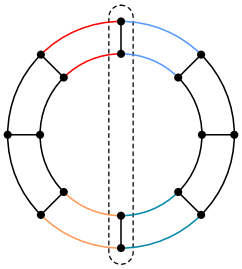

Figure 2.13: The circled vertex-set X is a tight separator of G. However, there are four 4-local components at X, none of which is tight. So while X is a 4-local separator, it is not a tight one.

Correspondence 2.4.2 (Tight separator)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$. A set $X \subseteq V(G)$ is a tight separator of G if and only if it is a tight r-local separator of G.*

The proof requires some preparation. For a vertex-set $X \subseteq V(G)$ we denote by $\mathcal{W}_r(X)$ the set of all X-walks that are concatenations of r-local X-paths and that repeat no vertices in X. The *r-local atoms* of X, or *r-toms* for short, are the $\subseteq$-maximal subsets $X' \subseteq X$ such that every two vertices in X' are linked by a walk in $\mathcal{W}_r(X)$. Then X is partitioned by its r-toms. We say that X is *r-locally atomic*, or *r-tomic* for short, if X has only one r-local atom. If X is r-tomic, then all r-local atoms of $p_r^{-1}(X)$ in G_r are isomorphic via deck transformations of p_r.

Lemma 2.4.3. *Every tight r-local separator of G is r-tomic.* □

Lemma 2.4.4. *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$. Then every tight separator of G is r-tomic.*

Proof. Suppose for a contradiction that some tight separator X of G is not r-tomic. Fix some y in an r-local atom Y of X and some $x \in X \setminus Y$. Since X is a tight separator of G, there are two tight components C and C' of $G - X$. Let P and P' be x–y paths through C and C', respectively. The path P meets the edge-set $E := E_G(C, Y)$ precisely once while P' avoids it. Thus, the cycle $P \cup P'$ meets E precisely once; in particular, oddly. Since the binary cycle space of G is generated by short cycles, there exists a short cycle $O \subseteq G$ that meets E oddly; in particular, O meets Y. As the cycle O meets the cut $E_G(C, X)$ evenly and $E_G(C, X)$ is the disjoint union of $E = E_G(C, Y)$ and $E_G(C, X \setminus Y)$, the cycle O has an edge $f \in E_G(C, X \setminus Y)$; in particular, O meets $X \setminus Y$. Hence, the cycle O of length $\leqslant r$ meets both Y and $X \setminus Y$, which contradicts the fact that Y is an r-local atom of X. □

Correspondence 2.4.5 (Local components)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$. The r-local components at an r-tomic $X \subseteq V(G)$ are precisely the non-empty componental cuts at X.*

Proof. By definition, every r-local component at X is included in a componental cut at X. Conversely, let C be a component of $G - X$ with non-empty componental cut. Let e, f be C–X edges in G and let E be the r-local component at X containing e. If E contains f as well, we are done; so we assume for a contradiction that $f \notin E$. Using the component C, we find an X-walk P that starts in X and traverses the edge e first, then uses vertices and edges in C, and finally traverses the edge f to end in X; in particular, P passes through E precisely once, and thus oddly often. Since X is r-tomic, there is a walk Q in G which joins the endvertices of e, f in X and is a concatenation of r-local X-paths. Since r-local X-paths traverse r-local components at X evenly, so do concatenations of them, such as the walk Q. Now concatenating the walks P and Q yields a closed walk which passes through E oddly often. Since the binary cycle space of G is generated by short cycles, there exists a short cycle $O \subseteq G$ which meets E oddly, and therefore contradicts that E is an r-local component at X. □

Proof of Correspondence 2.4.2. Let $X \subseteq V(G)$. If X is a tight separator, or if X is a tight r-local separator, then X is r-tomic by Lemma 2.4.4 or Lemma 2.4.3, respectively. Hence, Correspondence 2.4.5 completes the proof. □

A separation could alternatively be defined as a triple $\{U, X, W\}$ such that $U \sqcup X \sqcup W = V(G)$ and every component of $G - X$ has its vertex-set included in U or in W (see for example [Gro16a]). Indeed, the map $\{A, B\} \mapsto \{A \setminus B, A \cap B, B \setminus A\}$ would then be a bijection between the two perspectives on separations. We use the latter perspective to introduce local separations.

An *r-local separation* of G is a triple $\{E_1, X, E_2\}$ where X is a set of vertices of G and E_1, E_2 near-bipartition ∂X so that every r-local component at X is included in E_1 or in E_2. We refer to X as its *separator*, to the two classes E_1, E_2 as its *sides* and to $|X|$ as its *order*. Also we call $\{E_1, X, E_2\}$ an *r-local k-separation* for $k = |X|$. The ordered triples (E_2, X, E_1)

and (E_1, X, E_2) are the *orientations* of $\{E_1, X, E_2\}$ *towards* E_1 and *towards* E_2, respectively. We also call these *oriented r-local separations*, or just *r-local separations* by a slight abuse of notation. The *inverse* (E_2, X, E_1) of an oriented r-local separation (E_1, X, E_2) is denoted by $(E_1, X, E_2)^*$. An r-local separation with separator X is *tight* if each side includes a tight r-local component at X.

Every oriented separation (A_1, A_2) of G *induces* an oriented r-local separation (E_1, X, E_2), namely the one given by letting $X := A_1 \cap A_2$ and $E_i := E_G(A_i \setminus X, X)$. Then we also say that $\{A_1, A_2\}$ *induces* $\{E_1, X, E_2\}$. If (A_1, A_2) or $\{A_1, A_2\}$ is introduced as s, then we denote the (oriented) r-local separation that s induces by s_r. Moreover, if an r-local separation is introduced as $\{E_1, X, E_2\} := s_r$ given $\{A_1, A_2\} = s$, we tacitly assume of course that $E_i = E_G(A_i \setminus X, X)$ and not $E_i = E_G(A_{3-i} \setminus X, X)$.

Correspondence 2.4.6 (Separations)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. The map $s \mapsto s_r$ is a bijection between the (oriented) separations of G with r-tomic separator and the (oriented) r-local separations of G with r-tomic separator.*

Proof. This is immediate from Correspondence 2.4.5. □

Example 2.4.7. The assumption in Correspondence 2.4.6 that the separators are r-tomic is necessary. For example, let G be the graph depicted in Figure 2.14 and $r := 4$. The binary cycle space of G is generated by cycles of length $\leqslant r$. Let X denote the set of vertices circled in the figure, which is not r-tomic. Let E_1, E_2 consist of the red edges and the blue edges, respectively. Then $\{E_1, X, E_2\}$ is an r-local separation of G that is not induced by a separation of G.

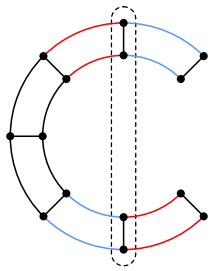

Figure 2.14: An r-local separation that is not induced by a separation, as the sides of the r-local separation restricted at the r-local atoms of its separator are flipped.

Correspondence 2.4.8 (Tight separations)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Then the map $s \mapsto s_r$ restricts to a bijection between the tight (oriented) separations of G and the tight (oriented) r-local separations of G.*

Proof. If s or s_r is tight, then their separator is r-tomic by Lemma 2.4.4 or Lemma 2.4.3, so the other of s, s_r is tight as well by Correspondence 2.4.5. Thus, $s \mapsto s_r$ is a bijection from the tight (oriented) separations to the tight (oriented) r-local separations by Correspondence 2.4.6. □

2.4.2 Interplay with the covering map

Projection 2.4.9 (Locally atomic)**.** *Let $\hat{X} \subseteq V(G_r)$ be r-tomic. Then $p_r(\hat{X})$ is r-tomic in G.*

Proof. Let x, y be two vertices in $p_r(\hat{X})$. Consider lifts $\hat{x}, \hat{y} \in \hat{X}$ of x, y. Since $\hat{X}$ is r-tomic, there exists an $\hat{x}$–$\hat{y}$ walk $\hat{W} \in \mathcal{W}_r(\hat{X})$, i.e. $\hat{W}$ is a concatenation of r-local $\hat{X}$-paths $\hat{P}_1, \dots, \hat{P}_n$. As p_r maps short cycles to short cycles, the projections $P_i := p_r(\hat{P}_i)$ are r-local X-paths or concatenations of r-local X-paths. Hence, their concatenation $P_1 \dots P_n = p_r(\hat{W})$ witnesses that x and y are in the same r-tom of X, as desired. □

Let $\hat{X}$ be a set of vertices of G_r and $X := p_r(\hat{X})$. The vertex-set $\hat{X}$ is *r-locally clopen* (with respect to p_r) if no r-local $p_r^{-1}(X)$-path links a vertex in $\hat{X}$ to a vertex in $p_r^{-1}(X) \setminus \hat{X}$. We remark that if $\hat{X}$ is r-locally clopen, then p_r restricts to a surjection $\partial\hat{X} \to \partial X$.

Example 2.4.10. For every $X \subseteq V(G)$, the r-local atoms of $p_r^{-1}(X)$ are the non-empty $\subseteq$-minimal r-locally clopen. More generally, $\hat{X} \subseteq V(G_r)$ is r-locally clopen if and only if it is a union of r-local atoms of $p_r^{-1}(p_r(\hat{X}))$.

Lift 2.4.11 ($e \sim_r f$ at X in G)**.** *Let $\hat{X}$ be an r-locally clopen set of vertices of G_r, and let $X := p_r(\hat{X})$. The following assertions are equivalent for every two edges $e, f \in \partial X$:*

(i) *$e \sim_r f$ at X in G;*

(ii) *for every lift $\hat{e} \in \partial\hat{X}$ of e there exists a lift $\hat{f} \in \partial\hat{X}$ of f such that $\hat{e} \sim_r \hat{f}$ at $\hat{X}$ in G_r;*

(iii) *there exist lifts $\hat{e}, \hat{f} \in \partial\hat{X}$ of e and f, respectively, such that $\hat{e} \sim_r \hat{f}$ at $\hat{X}$ in G_r.*

Proof. (i)→(ii). Let $e \sim_r f$ at X in G and let $\hat{e} \in \partial\hat{X}$ be a lift of e. If $e = f$, then we set $\hat{f} := \hat{e}$. Else there is an r-local X-walk W that traverses e first and f last. Let O be a short cycle in G that contains W. Let $\hat{O}$ be the lift of O that contains $\hat{e}$, and let $\hat{W}$ be the lift of W that starts with $\hat{e}$, which is thus a walk in $\hat{O}$. Let $\hat{f}$ be the lift of f that lies on $\hat{O}$, and is thus also the last edge of $\hat{W}$. As $\hat{X}$ is r-locally clopen, $\hat{W}$ is an r-local $\hat{X}$-walk that traverses $\hat{e}$ first and $\hat{f}$ last. Thus, $\hat{e} \sim_r \hat{f}$ at $\hat{X}$ in G_r.

(ii)→(iii) is straightforward, as every edge $e \in \partial X$ lifts to some edge $\hat{e} \in \partial\hat{X}$.

(iii)→(i). Similarly as in the proof of (i)→(ii), the r-local $\hat{X}$-walk witnessing $\hat{e} \sim_r \hat{f}$ at $\hat{X}$ in G_r projects to an r-local X-walk witnessing $e \sim_r f$ at X in G, as $\hat{X}$ is r-locally clopen. □

Projection 2.4.12 (Local components)**.** *Let $\hat{X}$ be an r-locally clopen set of vertices of G_r. Then every r-local component at $\hat{X}$ projects onto an r-local component at $X := p_r(\hat{X})$.*

Proof. Let $\hat{C}$ be an r-local component at $\hat{X}$ and put $C := p_r(\hat{C})$. Since $\hat{X}$ is r-locally clopen, all of C is included in ∂X. Then C is included in an r-local component C' at X by Lift 2.4.11 (iii)→(i). Conversely, C' is included in C by Lift 2.4.11 (i)→(ii), so $C = C'$ as desired. □

Example 2.4.13. The assumption in Projection 2.4.12 that $\hat{X}$ is r-locally clopen is necessary, see Figure 2.15.

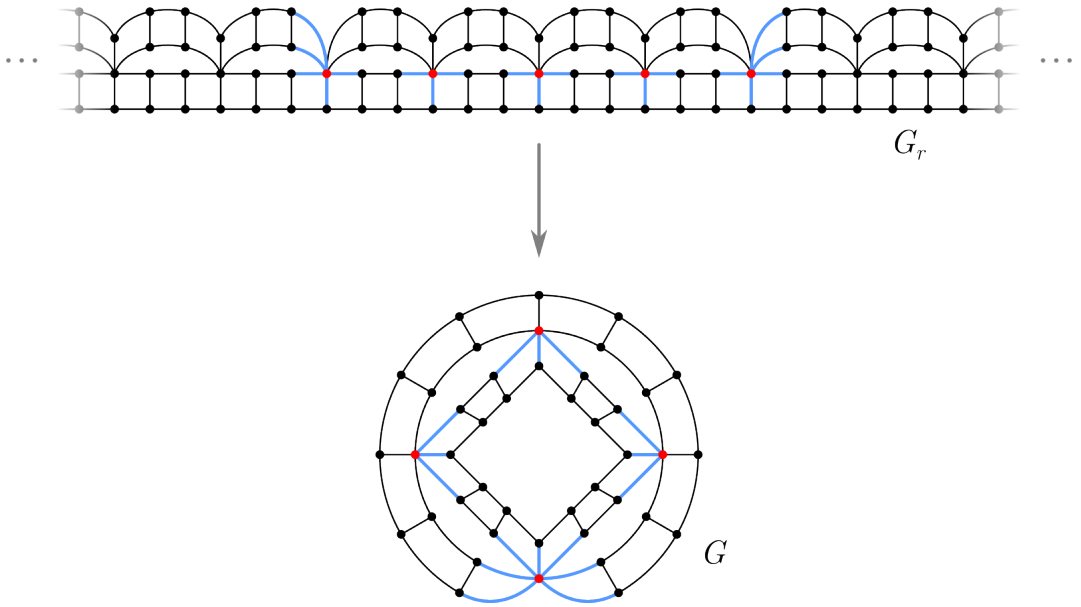

Figure 2.15: Let $r := 6$. The figure shows an r-local component (blue, top) at a non-r-locally-clopen vertex-set $\hat{X}$ (red) in G_r whose projection to G (blue, bottom) fails to be an r-local component at $p_r(\hat{X})$ (red)

Recall that a lift of $X \subseteq V(G)$ to G_r means a vertex-set $\hat{X} \subseteq V(G_r)$ such that p_r restricts to a bijection $\hat{X} \to X$. The following lemmas consider lifts that are r-locally clopen. However, not every $X \subseteq V(G)$ has a lift that is r-locally clopen. Sufficient conditions will be discussed in Section 2.5.

Lift 2.4.14 (Local components)**.** *Let X be a set of vertices in a connected graph G, and let $\hat{X}$ be an r-locally clopen lift of X to G_r.*

(i) *For every r-local component $\hat{C}$ at $\hat{X}$ in G_r, the covering p_r restricts to a bijection $\hat{C} \to C$ to an r-local component C at X in G.*

(ii) *The above map $\hat{C} \mapsto C$ is a bijection between the set of r-local components at $\hat{X}$ in G_r and the set of r-local components at X in G.*

(iii) *$\hat{C}$ is tight (with regard to $\hat{X}$) if and only if C is tight (with regard to X), for all $\hat{C}$.*

Proof. (i) and (ii). Let $\hat{C}$ be an r-local component at $\hat{X}$ in G_r. As $\hat{X}$ is r-locally clopen, p_r restricts to a surjection $\partial\hat{X} \to \partial X$. Thus, Projection 2.4.12 yields that C is an r-local component at X in G. The assumption that $\hat{X}$ is a lift of X implies that the restriction of p_r to $\hat{C}$ is a bijection $\hat{C} \to C$, giving (i), and it also implies (ii).

(iii) follows from (i) and (ii) as $\hat{X}$ is a lift of X. □

Example 2.4.15. The assumption in Lift 2.4.14 that $\hat{X}$ is a lift of X to G_r is necessary. For example, let $r := 6$ and consider the graph G in Figure 2.16. Let X be the red set of vertices

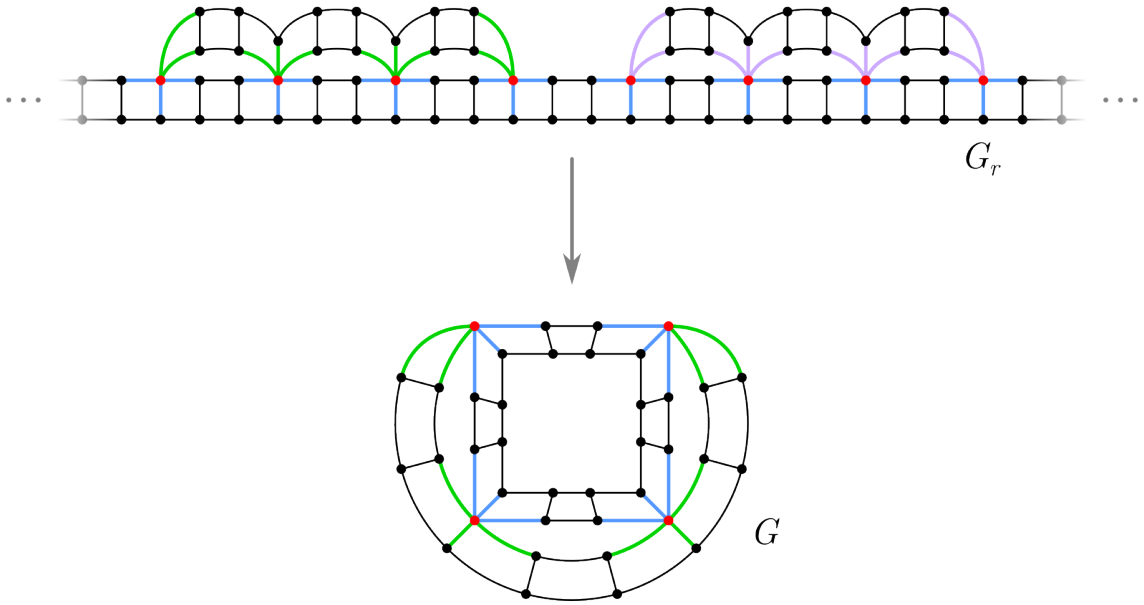

Figure 2.16: Let $r := 6$. The figure shows an r-locally clopen vertex-set $\hat{X}$ (red, top) in G_r which is not a lift of $X = p_r(\hat{X})$ (red, bottom). The coloured edge-sets in G_r are r-local components at $\hat{X}$, and the coloured edge-sets in G are r-local components at X.

in G. Then $\hat{X} := p_r^{-1}(X)$ is the only non-empty r-locally clopen subset of $p_r^{-1}(X)$, but $\hat{X}$ is not a lift of X. The blue r-local component at $\hat{X}$ projects to the blue r-local component at X, but not injectively, so (i) of Lift 2.4.14 fails. Both the green and the purple r-local component at $\hat{X}$ project to the green r-local component at X, so (ii) fails. The green r-local component at $\hat{X}$ is not tight, but its green projection at X is, so (iii) fails.

Projection 2.4.16 (Local separations)**.** *Let X be a set of vertices in a connected graph G, and let $\hat{X}$ be an r-locally clopen lift of X to G_r. Let $\{\hat{E}_1, \hat{X}, \hat{E}_2\}$ be an r-local k-separation of G_r. Then $\{E_1, X, E_2\}$ for $E_i := p_r(\hat{E}_i)$ and $X := p_r(\hat{X})$ is an r-local k-separation of G. Moreover, $\{E_1, X, E_2\}$ is tight if and only if $\{\hat{E}_1, \hat{X}, \hat{E}_2\}$ is tight.*

Proof. This follows immediately from Lift 2.4.14. □

In the context of Projection 2.4.16, when $\{\hat{E}_1, \hat{X}, \hat{E}_2\}$ is introduced as $\hat{s}$, then we often denote $\{E_1, X, E_2\}$ by $p_r(\hat{s})$ and refer to it as the *projection* of $\hat{s}$.

Lift 2.4.17 (Local separations)**.** *Let (E_1, X, E_2) be an r-local k-separation of a connected graph G. Let $\hat{X}$ be an r-locally clopen lift of X to G_r. For $i = 1, 2$ let $\hat{E}_i := p_r^{-1}(E_i) \cap \partial\hat{X}$. Then the following assertions hold:*

(i) *$(\hat{E}_1, \hat{X}, \hat{E}_2)$ is an r-local k-separation of G_r.*

(ii) *p_r restricts to a bijection $\hat{E}_i \to E_i$ for both $i = 1, 2$.*

(iii) *(E_1, X, E_2) is tight if and only if $(\hat{E}_1, \hat{X}, \hat{E}_2)$ is tight.*

Proof. This follows immediately from Lift 2.4.14. □

In the context of Lift 2.4.17, we call $(\hat{E}_1, \hat{X}, \hat{E}_2)$ the *$\hat{X}$-lift* of (E_1, X, E_2) to G_r. A *lift* of (E_1, X, E_2) means an $\hat{X}$-lift of (E_1, X, E_2) for some $\hat{X}$ as in Lift 2.4.17. We emphasise that a lift $\hat{s}$ of an r-local separation s of G is an r-local separation of G_r, even though it is also induced by a separation of G_r by Correspondence 2.4.6, if $\hat{X}$ is not only r-locally clopen but also r-tomic.

We remark that Projection 2.4.16 and Lift 2.4.17 are inverses of each other in the following sense: Let X be a set of vertices in a connected graph G and let $\hat{X}$ be an r-locally clopen lift of X to G_r. Then an r-local separation of G_r with separator $\hat{X}$ equals the $\hat{X}$-lift of its projection. Conversely, an r-local separation of G equals the projection of any of its lifts.

2.5 High displacement: a guarantee for the existence of lifts

While Projection 2.4.16 and Lift 2.4.17 explain the interplay of r-local separations with the covering map p_r, they both rest on the assumption that the separator $\hat{X}$ of G_r is an r-locally clopen lift of the separator X of G. Not all separators X admit such r-locally clopen lifts $\hat{X}$ (see Figures 2.15 and 2.16). We now introduce the *displacement $\Delta_r(G)$ of p_r*, a graph-invariant with which we formulate a guarantee for the existence of such lifts: If X is r-tomic and $|X| < \frac{2\Delta_r(G)}{r}$, then X has an r-locally clopen lift $\hat{X}$ (Lemma 2.5.7). This then yields the desired lifting/projection lemmas for r-local separations, Projection 2.5.8 and Lift 2.5.9. We also give various characterisations of the displacement (Lemma 2.5.2) as well as a purely combinatorial lower bound (Lemma 2.5.3). Finally, we provide a construction-scheme for graphs (Examples 2.5.10 and 2.5.13) which have high displacement and on which our theory of r-local separation is hence guaranteed to capture much of their r-local connectivity-structure.

2.5.1 Displacement as a guarantee

The *displacement* of p_r is defined as

$$\Delta_r(G) := \min \{ \operatorname{dist}_{G_r}(\hat{u}, \hat{v}) : \hat{u} \text{ and } \hat{v} \text{ are distinct lifts of the same vertex of } G \}$$

where we follow the convention that $\min(\emptyset) = \infty$. We remark that $\Delta_r(G) = \infty$ if and only if $G = G_r$.

Lemma 2.5.1. *$\Delta_r(G) > r$ for every connected graph G and $r \in \mathbb{N}$.*

Proof. The proof of [DJKK22, Lemma 4.3] shows this. □

Let ϱ be a non-negative integer. The *(combinatorial) $(\varrho/2)$-ball* $B_G(v, \varrho/2)$ around a vertex v in G is formed by all vertices and edges traversed by closed walks in G of length at most ϱ based at v. We remark that if $\varrho/2 \notin \mathbb{N}$, then the $(\varrho/2)$-ball around v is precisely the subgraph induced on all vertices of distance at most $\varrho/2$ from v; this, however, is in general not the case for $\varrho/2 \in \mathbb{N}$, where the edges between vertices of distance $\varrho/2$ to v may be

missing. A covering $p\colon C \to G$ of G *preserves* $(\varrho/2)$*-balls* if p restricts to an isomorphism $B_C(v, \varrho/2) \to B_G(p(v), \varrho/2)$ for every vertex v of C.[7]

Lemma 2.5.2. *Let G be a connected graph, $\varrho \geqslant r \geqslant 0$ integers, and x_0 a vertex of G. Then the following assertions are equivalent:*

(i) $\Delta_r(G) > \varrho$*;*

(ii) p_r *preserves* $(\varrho/2)$*-balls;*

(iii) $\pi_1^r(G, x_0) \supseteq \pi_1^\varrho(G, x_0)$*;*

(iv) p_r *and* p_ϱ *are isomorphic as coverings of* G*;*

In particular, either $\Delta_r(G) = \infty$ and (i) *through* (iv) *are satisfied for every integer $\varrho \geqslant r$, or $\Delta_r(G) - 1$ is the maximum $\varrho \geqslant r$ that satisfies* (i) *through* (iv)*.*

Proof. (i)↔(ii) is [DJKK22, Lemma 4.2].

(ii)→(iii). Since the covering p_r preserves $(\varrho/2)$-balls, its characteristic subgroup $\pi_1^r(G, x_0)$ contains all closed walks at x_0 that stem from cycles of length $\leqslant \varrho$. Thus, we also have $\pi_1^r(G, x_0) \supseteq \pi_1^\varrho(G, x_0)$.

(iii)→(iv). $\varrho \geqslant r$ yields $\pi_1^r(G, x_0) \subseteq \pi_1^\varrho(G, x_0)$ by definition of the local subgroups of the fundamental group of G. By (iii), the inclusion is an equality. Now (iv) follows directly from the fact that coverings are isomorphic if and only if their characteristic subgroups are [Hat02, Proposition 1.37].

(iv)→(i). By Lemma 2.5.1, every two lifts of the same vertex of G to G_ϱ have distance greater than ϱ in G_ϱ. Since p_r and p_ϱ are isomorphic, the same is true for p_r, that is, $\Delta_r(G) > \varrho$. □

Lemma 2.5.3. *Let G be a connected graph and $r \geqslant 0$ an integer. Then $\Delta_r(G)$ is at least the minimum of the lengths of induced cycles in G longer than r.*

Proof. Let ℓ be the minimum of all lengths of induced cycles in G longer than r (letting $\ell := \infty$ if no such cycle exists), and fix a vertex x_0 of G. By Lemma 2.5.2 (iii), it suffices to show that every closed walk W in G based at x_0 that stems from a cycle O of length $< \ell$ is in $\pi_1^r(G, x_0)$. We proceed by induction on $|O|$. If O has length $\leqslant r$, the walk W is in $\pi_1^r(G, x_0)$ by definition. Thus, we may assume that $|O| > r$. Since W stems from O, we may write $W = W_0 Q W_0^-$ where $Q = v_0 e_1 v_1 \dots v_n$ is a closed walk once around O. The definition of ℓ implies that O is not induced, i.e. there exists a chord $e = v_i v_j$ in G with $i < j$. Now W is homotopic to the concatenation of two closed walks $R = W_0 Q_R W_0^-$ and $S = S_0 Q_S S_0^-$ stemming from shorter cycles where $Q_R := v_0 Q v_i e v_j Q v_n$, $S_0 := W_0 v_n Q^- v_j$ and $Q_S := v_j e v_i Q v_j$. By the induction assumption, we have $R, S \in \pi_1^r(G, x_0)$, and thus $W \in \pi_1^r(G, x_0)$. □

Lemma 2.5.4. *Assume $r \geqslant 2$. Let $X \subseteq V(G)$ and let $W \in \mathcal{W}_r(X)$ use exactly ℓ vertices of X. Then the endvertices of W have distance $\leqslant (\ell - 1)r/2$ in G.*

[7] A notion that could be regarded as a precursor of the preservation of $(\varrho/2)$-balls has been studied by Georgakopoulos [Geo17].

Proof. Since W is contained in $\mathcal{W}_r(X)$, we can write W as a concatenation of r-local X-paths $P_1, \dots, P_n$. Let $i \in [n]$. If P_i has a single edge, set $P_i' := P_i$. Otherwise, P_i is contained in a short cycle O_i, and we let P_i' be a shortest path in O_i that joins the ends of P_i. In both cases, P_i' has length $\leqslant r/2$ (using the assumption that $r \geqslant 2$). As W repeats no vertices in X by definition of $\mathcal{W}_r(X)$, we have $n \leqslant \ell - 1$. Hence, the concatenation of all P_i' is a walk that links the same vertices as W does, and has length $\leqslant (\ell - 1)r/2$. □

Due to the assumption $r \geqslant 2$ in Lemma 2.5.4, in statements we will often assume

$$r \geqslant 2 \text{ or } k = 1 \tag{$*$}$$

by adding * to the respective statement.

Lemma 2.5.5. *Let G be connected and $r \leqslant 1$. Then every tight r-local separator of G has size 1.*

Proof. Since G is simple, G contains no cycles of length $\leqslant 1$. But a tight r-local separator of size $\geqslant 2$ would witness the existence of such a cycle. Thus, tight r-local separators have size 1. □

A set X of vertices of G is *r-locally 1-sheeted* if it is r-tomic and additionally the covering map p_r restricts to a bijection $\hat{X} \to X$ for some (equivalently: every) r-local atom $\hat{X}$ of $p_r^{-1}(X)$ in G_r.

Observation 2.5.6. *If $X \subseteq V(G)$ is r-locally 1-sheeted, then the action of the group $\Gamma_r(G)$ of deck transformations of p_r on the r-local atoms of $p_r^{-1}(X)$ is transitive and fixed-point free.*

Lemma 2.5.7.* *Assume that $kr/2 < \Delta_r(G)$ for a $k \in \mathbb{N}$. Then the following assertions are equivalent for all $(\leqslant k)$-sized vertex sets $X \subseteq V(G)$:*

(i) *X is r-tomic;*

(ii) *X is r-locally 1-sheeted.*

Proof. (i)→(ii). Let $\hat{X}$ be an arbitrary r-tom of $p_r^{-1}(X)$. As X is r-tomic by (i), p_r maps $\hat{X}$ onto X. Thus, it remains to show that p_r is injective on $\hat{X}$. So suppose for a contradiction that there are two distinct vertices $\hat{x}_1, \hat{x}_2 \in \hat{X}$ with $p_r(\hat{x}_1) = p_r(\hat{x}_2)$. Since $\hat{X}$ is r-tomic, there is an $\hat{x}_1$–$\hat{x}_2$ walk $\hat{W} \in \mathcal{W}_r(\hat{X})$. We choose the pair of vertices $\hat{x}_i$ and $\hat{W}$ such that $\hat{W}$ has minimum length. Then p_r is injective on $(V(\hat{W}) \cap \hat{X}) \setminus \{\hat{x}_2\}$. Hence, $|V(\hat{W}) \cap \hat{X}| =: \ell \leqslant k + 1$ because X has size $\leqslant k$. If $r \geqslant 2$, then the $\hat{x}_i$ have distance $\leqslant (\ell - 1)r/2 \leqslant kr/2$ by Lemma 2.5.4. This contradicts our assumption that $kr/2 < \Delta_r(G)$. Otherwise we may assume $k = 1$ and $r \leqslant 1$. Then the $\hat{x}_i$ have distance at most 1 but are contained in the same fibre of p_r. This contradicts that G has no loops.

(ii)→(i) is immediate from the definition of r-locally 1-sheeted. □

Projection 2.5.8.* *Assume that G is connected. Let $\hat{s}$ be an r-local separation of G_r with r-tomic separator $\hat{X}$ of size $\leqslant k$ such that $kr/2 < \Delta_r(G)$. Then $X := p_r(\hat{X})$ is r-locally 1-sheeted, $\hat{X}$ is an r-tom of $p_r^{-1}(X)$ and the projection $s := p_r(\hat{s})$ is well-defined.*

Proof. The projection $X = p_r(\hat{X})$ has size $\leqslant k$ and is r-tomic by Projection 2.4.9. Hence, X is r-locally 1-sheeted by Lemma 2.5.7, and so $\hat{X}$ is an r-tom of $p_r^{-1}(X)$. So the projection of $\hat{s}$ is well-defined by Projection 2.4.16. □

Lift 2.5.9.* *Assume that G is connected. Let s be an r-local separation of G with r-tomic separator X of size $\leqslant k$ such that $kr/2 < \Delta_r(G)$. Then X is r-locally 1-sheeted and s can be lifted.*

Proof. The separator X is r-locally 1-sheeted by Lemma 2.5.7. Hence, s can be lifted by Lift 2.4.17. □

2.5.2 A class of examples with high displacement

Recall that the *girth* $g(G)$ of a graph G is the minimal length of a cycle in G (and $g(G) := \infty$ if G is a forest). Let $(H, \mathcal{G})$ be a graph-decomposition of G. We associate with every edge $e = h_0h_1$ of H its *adhesion graph* $G_e := G_{h_0} \cap G_{h_1}$.

Assume now that the adhesion graphs of $(H, \mathcal{G})$ are pairwise disjoint. For every edge $e = hh'$ of H, we define $s^e := \{E_h^e, X^e, E_{h'}^e\}$ where $X^e := V(G_e)$, $E_h^e := \partial_{G_h} X^e$ and $E_{h'}^e := \partial_{G_{h'}} X^e$. The adhesion graphs being disjoint implies that E_h^e and $E_{h'}^e$ partition $\partial_G X^e$. In this context, a *traversal* of s^e is a walk in G such that its first edge is in E_h^e, its last edge is in $E_{h'}^e$ and all internal vertices are in X^e. A closed walk $W = v_1e_1v_2 \dots v_ne_nv_1$ in G *traverses* s^e with $e \in E(H)$ if (a cyclic shift of) W contains a traversal of s^e. For a closed walk W in G, we let H^W denote the graph that is formed by all the edges $e \in E(H)$ (and their endvertices) such that W traverses s^e.

Example 2.5.10. Assume that G is a connected graph. Let $(H, \mathcal{G})$ be a graph-decomposition of G such that

(i) every part G_h is complete, and

(ii) the adhesion graphs G_e are pairwise disjoint.

Then all induced cycles of G of length > 3 have length $\geqslant g(H)$, so $\Delta_r(G) \geqslant g(H)$ for every $r \geqslant 3$.

Proof. Let O be an induced cycle in G. In the following, we interpret O also as a closed walk once around O. If H^O is empty, then O is contained in a part G_h; thus, O is a triangle because the parts are complete by (i). Assume now that H^O is not empty. Note that for every traversal P of s^e for an edge $e \in E(H)$, the two endvertices of P send edges of G to all internal vertices of P. Since O is induced and the parts of $(H, \mathcal{G})$ are complete, for every traversed edge e of H^O the intersection $O \cap G_e$ consists of a single vertex, and for every node h of H^O the

intersection $O \cap G_h$ consists of a single edge joining different adhesion graphs. Thus, H^O is a cycle $h_1 \dots h_n h_1$. The unique edges in $O \cap G_{h_i}$ are pairwise distinct by (ii). Hence, the length of O is at least $g(H)$. □

Let $(H, \mathcal{G})$ be a graph-decomposition of G whose adhesion graphs are pairwise disjoint. We say that a closed walk W traverses s^e *k times* (where $e \in E(H)$) if the number of cyclic shifts of W which contain a traversal of s^e as initial segment equals k.

Lemma 2.5.11. *Let $(H, \mathcal{G})$ be a graph-decomposition of a finite connected graph G with pairwise disjoint adhesion graphs. Let W be a closed walk in G such that W traverses s^e at most once for every edge e of H and such that W traverses at least one s^e. Then:*

(i) *H^W is non-empty and every node of H^W has degree $\geqslant 2$; in particular, H^W contains a cyle.*

(ii) *W contains a subwalk in G_h between two distinct adhesion graphs of $(H, \mathcal{G})$ for every $h \in V(H^W)$.*

Proof. Since W traverses at least one s^e, the graph H^W is non-empty. Let h be an arbitrary node of H^W. By the definition of H^W, the closed walk W traverses s^e for some $e \in E(H)$ incident to h. But as the adhesion graphs of $(H, \mathcal{G})$ are disjoint and W traverses s^e no more than once, W leaves G_h by traversing another s^f with $f \in E(H)$ incident to h. So h has degree $\geqslant 2$ in H^W, and (ii) follows. As G is finite and the adhesion graphs are disjoint, H and H^W are also finite. Hence, (i) implies that H^W contains a cycle. □

We denote by $\ell(H, \mathcal{G})$ the length of a shortest closed walk W in G such that W traverses s^e at most once for every edge e of H and such that H^W contains a cycle. If no such walk exists, we let $\ell(H, \mathcal{G}) := \infty$. Further, we write $d(H, \mathcal{G})$ for the minimum distance $d_{G_h}(G_e, G_f)$ over all nodes h of H and all edges e, f of H incident to h. In edge-cases we let $d(H, \mathcal{G}) := \infty$.

Corollary 2.5.12. *Let $(H, \mathcal{G})$ be a graph-decomposition of a finite connected graph G with pairwise disjoint adhesion graphs. Then $\ell(H, \mathcal{G}) \geqslant g(H) \cdot d(H, \mathcal{G})$.*

Proof. If $\ell(H, \mathcal{G}) = \infty$ we are done, so we may assume that there exists a closed walk W in G that witnesses $\ell(H, \mathcal{G}) < \infty$. By Lemma 2.5.11, H^W contains a cycle O. For every node h of O, the walk W contains a subwalk W_h in G_h between two distinct adhesion graphs of $(H, \mathcal{G})$, as W traverses every s^e at most once. We can and do choose the W_h so that their internal vertices are not contained in adhesion graphs. Since the adhesion graphs are pairwise disjoint, the subwalks W_h are pairwise edge-disjoint. Hence, W has length at least $g(H) \cdot d(H, \mathcal{G})$. □

Example 2.5.13. Let $(H, \mathcal{G})$ be a graph-decomposition of a finite connected graph G such that

(i) every part G_h is connected and $(G_h)_r = G_h$,

(ii) its adhesion graphs G_e are pairwise disjoint,

(iii) every adhesion graph G_e is connected.

Then $\Delta_r(G) \geqslant \ell(H,\mathcal{G}) \geqslant g(H) \cdot d(H,\mathcal{G}) \geqslant g(H)$.

Proof. The inequality $\ell(H,\mathcal{G}) \geqslant g(H) \cdot d(H,\mathcal{G})$ was proved in Lemma 2.5.11. The disjointness of the adhesion graphs ensure that $d(H,\mathcal{G}) \geqslant 1$, and thus $g(H) \cdot d(H,\mathcal{G}) \geqslant g(H)$. So it remains to show $\Delta_r(G) \geqslant \ell(H,\mathcal{G})$.

We introduce some terminology first. For every node h of H, let $\mathring{G}_h$ denote the graph obtained from G_h by deleting all edges that are contained adhesion graphs. Since the adhesion graphs are pairwise disjoint, the graphs $\mathring{G}_h$ are pairwise edge-disjoint. For a closed walk $W = v_1e_1v_2 \dots v_ne_nv_1$ in G, we denote by $f(W)$ the number of indices i such that the edge e_i lies in some $\mathring{G}_h$. We say that W is *generated by short cycles* if, for some (equivalently: every) vertex x_0 of G and for some (equivalently: every) closed walk W' based at x_0 which stems from W, we have $W' \in \pi_1^r(G,x_0)$.

To show $\Delta_r(G) \geqslant \ell(H,\mathcal{G})$, by Lemma 2.5.2 (iii) it suffices to show that every closed walk W in G of length $< \ell(H,\mathcal{G})$ is generated by short cycles. Since $f(W)$ is no larger than the length of W, we may prove the stronger statement that

$$\textit{every closed walk } W \textit{ in } G \textit{ with } f(W) < \ell(H,\mathcal{G}) \textit{ is generated by short cycles.} \tag{2.4}$$

We prove (2.4) by induction on $f(W)$. For this, we analyse how W traverses the s^e of $(H,\mathcal{G})$.

If W traverses no s^e (which is the case if $f(W) = 0$), then a single part G_h contains W, and thus the short cycles in G_h generate W by (i). So we may assume that W traverses some s^e of $(H,\mathcal{G})$ and that the statement holds for all closed walks W' with $f(W') < f(W)$.

First, assume that W traverses every s^e of $(H,\mathcal{G})$ at most once. By Lemma 2.5.11 (i), the graph $H^W \subseteq H$ contains a cycle. Thus, $f(W) \geqslant \ell(H,\mathcal{G})$ by Lemma 2.5.11 (ii), which contradicts the choice of W.

Now assume that $W =: v_1e_1v_2 \dots v_ne_nv_1$ traverses some s^e of $(H,\mathcal{G})$ at least twice. Let P_i and P_j be two such traversals where P_i meets G_e in a vertex v_i and P_j meets G_e in a vertex v_j with $i < j$. We may pick a v_i–v_j path P in G_e by (iii). Let $W_1 := v_1Wv_iPv_jWv_1$ and $W_2 := v_jP^-v_iWv_j$. As W is homotopic to $W_1(v_1W^-v_j)W_2(v_jWv_1)$, it suffices to show that W_1 and W_2 are generated by short cycles. It is immediate form the construction that $f(W_1) + f(W_2) = f(W)$, as $P \subseteq G_e$. Moreover, $f(W_i) > 0$ for both i, since by construction each W_i contains one of the edges in $\partial X^e \subseteq \mathring{G}_h \cup \mathring{G}_{h'}$, where $e = hh'$. Hence, $f(W_i) < f(W)$ for both i. So by the induction hypothesis both W_i are generated by short cycles, as desired. □

2.5.3 Local separations whose separator is a clique

Lift 2.5.14. *Let G be a connected graph and $r \geqslant 3$ an integer. The r-local separations of G_r whose separator is a clique are precisely the lifts of the r-local separations of G whose separator is a clique.*

Lemma 2.5.15. *Let G be a connected graph and $r \geqslant 3$ an integer. Then:*

(i) *For every clique $\hat{X}$ of G_r, its projection $p_r(\hat{X})$ is a clique of G, and p_r restricts to a bijection from $\hat{X}$ to $p_r(\hat{X})$.*

(ii) *For every clique X of G, there exists a clique $\hat{X}$ of G_r such that p_r restricts to a bijection from $\hat{X}$ to X.*

(iii) $(\bigcup_{K\in K_{G_r}(\hat{X})} K) \cap (\bigcup_{K\in K_{G_r}(\gamma(\hat{X}))} K) = \emptyset$ *for every clique $\hat{X}$ of G_r and every $\gamma \in \Gamma(p_r) \setminus \{\mathrm{id}_{G_r}\}$.*

For a given clique X of G, we refer to all cliques $\hat{X}$ of G_r such that p_r restricts to a bijection from $\hat{X}$ to X as *lifts of X to G_r*. We remark that Lemma 2.5.15 ensures that every clique of G_r is a lift of a clique in G, lifts of cliques of G exist, and two distinct lifts of a common clique X are not only disjoint but r-toms of $p_r^- 1(X)$ and thus r-locally clopen lifts of X.

Proof of Lemma 2.5.15. (i) and (ii). Every clique is contained in the 3/2-ball around each of its vertices. Since $r \geqslant 3$, the r-local covering p_r is in particular 3/2-ball-preserving. From this, (i) and (ii) follow immediately.

(iii). Suppose for a contradiction that there is $\gamma \in \Gamma(p_r) \setminus \{\mathrm{id}_{G_r}\}$ and a clique $\hat{X}$ of G_r such that $\bigcup_{K\in K_{G_r}(\hat{X})} K$and $\bigcup_{K\in K_{G_r}(\gamma(\hat{X}))} K$ share a vertex $\hat{y}$. Then there is a vertex $\hat{x}_0 \in \hat{X}$ and a vertex $\hat{x}_1 \in \gamma(\hat{X})$ both with distance $\leqslant 1$ to $\hat{y}$. As $\hat{x}_0' \coloneqq \gamma(\hat{x}_0)$ has distance $\leqslant 1$ from $\hat{x}_1$, the two vertices $\hat{x}_0, \hat{x}_0'$ in the same fibre of p_r have distance $\leqslant 3 \leqslant r$. This contradicts Lemma 2.5.1. □

Proof of Lift 2.5.14. By Lemma 2.5.15 (i) and (ii), p_r induces a surjection $\hat{X} \mapsto p_r(\hat{X})$ from the cliques of G_r to the cliques of G which is injective up to deck transformation. Lemma 2.5.15 (iii) yields that the lifts $\hat{X}$ of X to G_r are precisely the r-toms of $p_r^{-1}(X)$ in G_r; in particular, the lifts $\hat{X}$ are precisely the r-locally clopen lifts of X. Thus, the statement follows immediately from Projection 2.4.16 and Lift 2.4.17. □

2.6 Links and crossing

As the next step in our finite combinatorial theory of r-local separations, we develop notions of *crossing* and *nested* for r-local separations. We show that they agree with the respective notions for global separations on r-local coverings (Correspondences 2.6.3 and 2.6.4). We then prove, assuming $k < K(G,r) = \frac{\Delta_r(G)}{r} + 1$, that two $(\leqslant k)$-separations s and t of G are nested if and only if all their lifts are nested (Lifts 2.6.6 and 2.6.7). Moreover, if s and t cross, then the crossing-relation defines a matching on the set of all lifts of s and t (Lift 2.6.8).

2.6.1 From global to local, and correspondence in the local cover

Let $\{E_1, X, E_2\}$ and $\{F_1, Y, F_2\}$ be two r-local separations of G. The *X-link for F_i* is the set of all vertices $x \in X \setminus Y$ such that there is a walk in $\mathcal{W}_r(X)$ that starts at x, visits Y, and its first edge in ∂Y is contained in F_i. The *Y-link for E_i* is defined analogously, by replacing F_i with E_i and interchanging the roles of X and Y.

A cycle O *alternates* between two disjoint vertex-sets X and Y if there are four distinct vertices x, x', y, y' on O such that $x, x' \in X$ and $y, y' \in Y$ and O visits these four vertices in the order x, y, x', y'. We emphasise that O may intersect $X \cup Y$ in more than four vertices. This definition is inspired from [Car22].

Two vertex-sets $X, Y \subseteq V(G)$ are *r-coupled* in G if either X and Y share a vertex, or X and Y are disjoint and there is a short cycle in G that alternates between X and Y.

Let $\{E_1, X, E_2\}$ and $\{F_1, Y, F_2\}$ be two r-local separations with r-tomic separators. Inspired by Lemma 2.2.2, we say that $\{E_1, X, E_2\}$ and $\{F_1, Y, F_2\}$ *cross* if

- X and Y are r-coupled, and
- for every pair of indices $i, j \in [2]$ at least one of the following three sets is non-empty: the X-link for F_j, the Y-link for E_i, or the edge-set $E_i \cap F_j \cap \partial(X \cap Y)$.

If $\{E_1, X, E_2\}$ and $\{F_1, Y, F_2\}$ do not cross, then they are *nested*.

Lemma 2.6.1. *If two separations of a connected graph with empty centre cross, then some two opposite links are non-empty.*

Proof. Otherwise, the separator of some corner is empty, which contradicts Lemma 2.2.2 (iii). □

Lemma 2.6.2. *Let $\{A_1, A_2\}$ and $\{C_1, C_2\}$ be crossing separations of a connected graph G with r-tomic separators X and Y, respectively. Then X and Y are r-coupled.*

Proof. We are done if X and Y intersect. Thus, we may assume that $X \cap Y = \emptyset$. Then Lemma 2.6.1 yields that two opposite links, say both X-links, are non-empty. Since X is r-tomic and $X \cap Y$ is empty, there is a short cycle $O \subseteq G$ or a edge e that meets both X-links, one in a vertex x and the other one in a vertex x'. The edge e would contradict the fact that $\{C_1, C_2\}$ is a separation. Thus, O exists. Then both x–x' paths in O must meet Y, and therefore O alternates between X and Y. □

Correspondence 2.6.3 (Links)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Let $\{A_1, A_2\}$ and $\{C_1, C_2\}$ be separations of G with r-tomic separators X and Y, respectively, such that X and Y are r-coupled. Let $\{E_1, X, E_2\}$ and $\{F_1, Y, F_2\}$ be the r-local separations induced by $\{A_1, A_2\}$ and $\{C_1, C_2\}$, respectively. Then:*

(i) *the X-link for C_i is equal to the X-link for F_i, for both $i = 1, 2$;*

(ii) *the Y-link for A_i is equal to the Y-link for E_i, for both $i = 1, 2$;*

(iii) $E\big(X \cap Y, (A_i \smallsetminus A_{3-i}) \cap (C_j \smallsetminus C_{3-j})\big) = E_i \cap F_j \cap \partial(X \cap Y)$ *for all $i, j \in \{1, 2\}$.*

Proof. Since X and Y are r-coupled, there is a vertex $v \in X \cap Y$ or there is a short cycle $O \subseteq G$ that alternates between X and Y. To show (i) and (ii), by symmetry it suffices to show that the X-link for C_1 is equal to the X-link for F_1. The backward inclusion is straightforward.

To see the forward inclusion, let x be a vertex in the X-link for C_1. If $v \in X \cap Y$ exists, we use that X is r-tomic to find a walk $W \in \mathcal{W}_r(X)$ that starts with x and visits $v \in Y$. Otherwise, we use that X is r-tomic and the existence of O to find a walk $W \in \mathcal{W}_r(X)$ that starts in X and visits Y. Let y denote the first vertex of Y on W. Since x lies in the X-link for C_1, all vertices of W that precede y are contained in $C_1 \setminus C_2$. Hence, y is preceded by an edge in F_1 on W. In particular, the first edge of the walk W in ∂Y lies in F_1. So x lies in the X-link for F_1.

(iii) is immediate from the assumption that $\{A_1, A_2\}$ induces $\{E_1, X, E_2\}$ and $\{C_1, C_2\}$ induces $\{F_1, Y, F_2\}$. □

Correspondence 2.6.4 (Crossing)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Let s and t be separations of G with r-tomic separators. The separations s and t cross if and only if their r-local counterparts s_r and t_r cross.*

Proof. Let X and Y denote the separators of s and t, respectively. If s and t cross, then X and Y are r-coupled by Lemma 2.6.2. If s_r and t_r cross, then X and Y are r-coupled by the definition of crossing. In either case, we may use Correspondence 2.6.3 to find that s_r and t_r or s and t also cross. □

2.6.2 Interplay with the covering map

Lemma 2.6.5. *$(k-1)r < \Delta_r(G)$ implies $kr/2 < \Delta_r(G)$, for all integers $k, r \geqslant 0$.*

Proof. If $k \leqslant 1$, then $kr/2 \leqslant r/2 < r < \Delta_r(G)$, where the last inequality holds by Lemma 2.5.1. Otherwise, $k \geqslant 2$ gives $(k-1)r - kr/2 = (k/2 - 1)r \geqslant 0$, which yields the desired $kr/2 < (k-1)r < \Delta_r(G)$. □

Lift 2.6.6.* *Let X and Y be r-tomic vertex-sets of size $\leqslant k$ in a connected graph G, such that X and Y are r-coupled and $(k-1)r < \Delta_r(G)$. Let $\hat{X}$ be an r-tom of $p_r^{-1}(X)$. Let $\hat{H}$ denote the union of all short cycles in G_r that have at least two vertices in $\hat{X}$. Then there exists a unique r-tom $\hat{Y}$ of $p_r^{-1}(Y)$ such that $\hat{Y}$ meets $\hat{X} \cup V(\hat{H})$. Moreover, $\hat{Y}$ is the unique r-tom of $p_r^{-1}(Y)$ that is r-coupled with $\hat{X}$.*

Proof. First, assume that $r \leqslant 1$. Then $k = 1$. Since there are no cycles of length at most 1 and as X and Y are r-coupled, we thus have $X = Y$. Also $\hat{H}$ consists precisely of $\hat{X}$, and thus $\hat{Y}$ is unique, as it equals $\hat{X}$, as desired. Hence, we may assume that $r \geqslant 2$.

Existence. Recall that X and Y are r-locally 1-sheeted by Lemma 2.6.5 and Lemma 2.5.7. Since X and Y are r-coupled, there either is a vertex $v \in X \cap Y$, or X and Y are disjoint and there is a short cycle $O \subseteq G$ that alternates between X and Y. Let an r-tom $\hat{X}$ of $p_r^{-1}(X)$ be given. If v exists, then we denote by $\hat{v}$ the unique lift of v that is contained in $\hat{X}$. If O exists, then we let $\hat{O}$ denote the unique lift of O for which p_r restricts to a bijection between $V(\hat{O}) \cap \hat{X}$ and $V(O) \cap X$. Now let $\hat{Y}$ denote the unique r-tom of $p_r^{-1}(Y)$ that contains $\hat{v}$ or for which p_r restricts to a bijection between $\hat{Y} \cap V(\hat{O})$ and $Y \cap V(O)$. Then $\hat{v} \in \hat{X} \cap \hat{Y}$ or $\hat{O}$ is a

short cycle in G_r that alternates between the disjoint vertex-sets $\hat{X}$ and $\hat{Y}$. So $\hat{X}$ and $\hat{Y}$ are r-coupled, and in particular $\hat{Y}$ meets $\hat{X} \cup V(\hat{H})$.

Uniqueness. Assume for a contradiction that there are two distinct r-toms $\hat{Y}_1, \hat{Y}_2$ of $p_r^{-1}(Y)$ that both meet $\hat{X} \cup V(\hat{H})$ in vertices $\hat{x}_1, \hat{x}_2$, respectively. If $\hat{x}_i \in V(\hat{H}) \setminus \hat{X}$, then let $\hat{O}_i$ be a short cycle in G_r that contains $\hat{x}_i$ and intersects $\hat{X}$ in some two vertices $\hat{u}_i, \hat{v}_i$. Further, let $\hat{x}'_i$ denote the unique vertex in $\hat{Y}_{3-i}$ with $p_r(\hat{x}'_i) = p_r(\hat{x}_i)$.

Suppose first that both $\hat{x}_i$ lie in $\hat{X}$. Since $\hat{X}$ and $\hat{Y}_2$ are r-tomic, there are an $\hat{x}_1$–$\hat{x}_2$ path and an $\hat{x}_2$–$\hat{x}'_1$ path in G_r, both of length $\leqslant (k-1)r/2$, by Lemma 2.5.4. So $\hat{x}_1$ and $\hat{x}'_1$ have distance $\leqslant (k-1)r$, contradicting that $(k-1)r < \Delta_r(G)$.

Suppose second that $\hat{x}_1$ lies in $\hat{X}$ and that $\hat{x}_2$ lies on $\hat{O}_2$. Since $\hat{X}$ is r-tomic, there is a walk $Q(\hat{X}) \in \mathcal{W}_r(\hat{X})$ that starts in $\hat{x}_1$ and ends either in $\hat{u}_2$ or in $\hat{v}_2$. By choosing $Q(\hat{X})$ of shortest length, we get that $Q(\hat{X})$ ends in $\hat{u}_2$ and avoids $\hat{v}_2$, say. Since $Q(\hat{X})$ avoids $\hat{v}_2$, there is a $\hat{x}_1$–$\hat{u}_2$ path $Q'(\hat{X})$ in G_r of length $\leqslant (k-2)r/2$ by Lemma 2.5.4. Let $Q(\hat{O}_2)$ be a path of length $\leqslant r/2$ that links $\hat{u}_2$ to $\hat{x}_2$. Since $\hat{Y}_2$ is r-tomic, there is an $\hat{x}_2$–$\hat{x}'_1$ path $Q(\hat{Y}_2)$ in G_r of length $\leqslant (k-1)r/2$, by Lemma 2.5.4. Then the concatenation of $Q'(\hat{X})$, $Q(\hat{O}_2)$ and $Q(\hat{Y}_2)$ links $\hat{x}_1$ to $\hat{x}'_1$ and has length $\leqslant (k-1)r$, contradicting that $(k-1)r < \Delta_r(G)$.

The case that $\hat{x}_1$ lies on $\hat{O}_1$ while $\hat{x}_2$ lies in $\hat{X}$ follows from the previous case by symmetry.

Suppose finally that both $\hat{x}_i$ lie on $\hat{O}_i$. Since $\hat{X}$ is r-tomic, there is a path $Q(\hat{X}) \in \mathcal{W}_r(\hat{X})$ that starts in $\hat{u}_1$ or in $\hat{v}_1$ and that ends either in $\hat{u}_2$ or in $\hat{v}_2$. We choose $Q(\hat{X})$ of shortest length. Possibly after renaming the vertices, we have that $Q(\hat{X})$ links $\hat{u}_1$ to $\hat{u}_2$. If $\{\hat{u}_1, \hat{v}_1\}$ and $\{\hat{u}_2, \hat{v}_2\}$ share a vertex, then $Q(\hat{X})$ has length 0, and we let $Q'(\hat{X}) := Q(\hat{X})$. Otherwise, $Q(\hat{X})$ avoids two vertices in $\hat{X}$, so there is a $\hat{u}_1$–$\hat{u}_2$ path $Q'(\hat{X})$ in G_r of length $\leqslant (k-3)r/2$ by Lemma 2.5.4. For both i let $Q(\hat{O}_i)$ be a $\hat{x}_i$–$\hat{u}_i$ path in $\hat{O}_i$ of length $\leqslant r/2$. Using that $\hat{Y}_2$ is r-tomic and Lemma 2.5.4, let $Q(\hat{Y}_2)$ be a $\hat{x}_2$–$\hat{x}'_1$ path in G_r of length $\leqslant (k-1)r/2$. Then the concatenation of $Q(\hat{O}_1)$, $Q'(\hat{X})$, $Q(\hat{O}_2)$ and $Q(\hat{Y}_2)$ links $\hat{x}_1$ to $\hat{x}'_1$ and has length $\leqslant (k-1)r$, contradicting that $(k-1)r < \Delta_r(G)$. □

Lift 2.6.7 (Links).* *Let G be a connected graph. Let $s = \{E_1, X, E_2\}$ and $t = \{F_1, Y, F_2\}$ be r-local separations of G with r-tomic separators of size $\leqslant k$ such that X and Y are r-coupled and $(k-1)r < \Delta_r(G)$. Let $\hat{s} = \{\hat{E}_1, \hat{X}, \hat{E}_2\}$ be a lift of s, and let $\hat{t} = \{\hat{F}_1, \hat{Y}, \hat{F}_2\}$ be the unique lift of t such that $\hat{X}$ and $\hat{Y}$ are r-coupled, which exists by Lift 2.6.6. Then p_r restricts to bijections between*

(i) *the $\hat{X}$-link for $\hat{F}_i$ and the X-link for F_i, for both $i = 1, 2$;*

(ii) *the $\hat{Y}$-link for $\hat{E}_i$ and the Y-link for E_i, for both $i = 1, 2$;*

(iii) *$\hat{X} \cap \hat{Y}$ and $X \cap Y$;*

(iv) *$\hat{E}_i \cap \hat{F}_j \cap \partial(\hat{X} \cap \hat{Y})$ and $E_i \cap F_j \cap \partial(X \cap Y)$, for all $i, j \in \{1, 2\}$.*

Moreover, s and t cross if and only if $\hat{s}$ and $\hat{t}$ cross.

In the context of Lift 2.6.7, we recall that s and t can be lifted by Lift 2.5.9 and Lemma 2.6.5.

Proof. Since X and Y are r-coupled, there either is a vertex $v \in X \cap Y$ or X and Y are disjoint and there is a short cycle $O \subseteq G$ that alternates between X and Y. If v exists, then we denote by $\hat{v}$ the unique lift of v in $\hat{X} \cap \hat{Y}$. If O exists, then we let $\hat{O}$ denote the unique lift of O that alternates between $\hat{X}$ and $\hat{Y}$.

The separators X and Y of s and t are r-locally 1-sheeted by Lift 2.5.9. The separators $\hat{X}$ and $\hat{Y}$ of $\hat{s}$ and $\hat{t}$ are r-toms of $p_r^{-1}(X)$ and $p_r^{-1}(Y)$, respectively, by the definition of lift. Hence, p_r restricts to bijections $\hat{X} \to X$ and $\hat{Y} \to Y$. Since $\hat{Y}$ is the unique r-tom of $p_r^{-1}(Y)$ which is r-coupled with $\hat{X}$, this shows (iii). (iv) then follows from (iii).

We now claim that

$$\text{the lift to } \hat{X} \text{ of the } X\text{-link for } F_1 \text{ is included in the } \hat{X}\text{-link for } \hat{F}_1. \tag{2.5}$$

To see this, let x be a vertex in the X-link for F_1, and let $\hat{x}$ denote its unique lift in $\hat{X}$. Let $W \in \mathcal{W}_r(X)$ witness that x lies in the X-link for F_1; that is, W starts at x, W is a concatenation of r-local X-paths $P_1, \dots, P_n$, and the first edge e of W in ∂Y exists and lies in F_1. Let $\hat{W}$ denote the lift of W to G_r that starts at $\hat{x}$. This defines lifts $\hat{P}_i$ of all P_i, and a lift $\hat{e}$ of e. As $\hat{X}$ is an r-tom and thus r-locally closed, all $\hat{P}_i$ are r-local $\hat{X}$-paths; hence, $\hat{W} \in \mathcal{W}_r(\hat{X})$.

It suffices to show that $\hat{e}$ lies in $\partial\hat{Y}$. Let $\hat{Z}$ be the r-tom of $p_r^{-1}(Y)$ with $\hat{e} \in \partial\hat{Z}$. Since $\hat{e}$ lies on some r-local $\hat{X}$-path $\hat{P}_i$, we find that the endvertex of $\hat{e}$ in $\hat{Z}$ either lies on a short cycle in G_r that has two vertices in $\hat{X}$ or is contained in $\hat{X}$ (which occurs if $\hat{e}$ is an X-edge). Since $\hat{Y}$ and $\hat{X}$ are r-coupled, this yields $\hat{Z} = \hat{Y}$ by the uniqueness in Lift 2.6.6. Therefore, $\hat{e} \in \partial Y$. Then $\hat{e} \in \hat{F}_1$ as well. This completes the proof of (2.5).

Next, we claim that

$$\text{the projection of the } \hat{X}\text{-link for } \hat{F}_1 \text{ is included in the } X\text{-link for } F_1. \tag{2.6}$$

So let $\hat{x}$ be a vertex in the $\hat{X}$-link for $\hat{F}_1$, and let $\hat{W} \in \mathcal{W}_r(\hat{X})$ and $\hat{e} \in \hat{F}_1 \cap E(\hat{W})$ witness this. Let $\hat{Q}$ denote the maximal initial subwalk of $\hat{W}$ that does not use $\hat{e}$, so $\hat{Q}$ ends in the endvertex of $\hat{e}$ outside of $\hat{Y}$. Let x, e, W and Q denote the projections of $\hat{x}, \hat{e}, \hat{W}$ and $\hat{Q}$. It suffices to show that Q avoids Y. Let us assume for a contradiction that Q has a vertex y in Y, and let $\hat{y}$ denote the lift of y on $\hat{Q}$. Then $\hat{y}$ lies in an r-tom $\hat{Z}$ of $p_r^{-1}(Y)$ other than $\hat{Y}$. Since $\hat{Q}$ is a subwalk of $\hat{W} \in \mathcal{W}_r(\hat{X})$, we get that $\hat{y}$ either lies on some short cycle in G_r that meets $\hat{X}$ in two vertices or is contained in $\hat{X}$ (which occurs if y lies on an X-edge). Since $\hat{X}$ and $\hat{Y}$ are r-coupled and $\hat{Z} \neq \hat{Y}$, this contradicts the uniqueness in Lift 2.6.6. Hence, x lies in the X-link for F_1.

Both Equations (2.5) and (2.6) extend to all four links by analogue arguments, so Items (i) and (ii) hold. The 'moreover'-part is immediate from Items (i), (ii) and (iv). □

Lift 2.6.8 (Crossing).* *Let s and t be two r-local separations of a connected graph G with r-tomic separators X and Y of size $\leqslant k$, respectively, such that $(k-1)r < \Delta_r(G)$. The following assertions are equivalent:*

(i) *s crosses t;*

(ii) *some lift of s crosses some lift of t;*

(iii) *the crossing-relation defines a matching between the set of lifts of s and the set of lifts of t.*

In the context of Lift 2.6.8, we recall that s and t can be lifted by Lift 2.5.9 and Lemma 2.6.5.

Proof. (i)→(ii). This follows directly from the 'moreover'-part of Lift 2.6.7.

(ii)→(iii). As the lifts of s and the lifts of t each form an orbit under the action of the group of deck transformations, it suffices to show that some lift of s is crossed by precisely one lift of t. Let $\hat{s}$ and $\hat{t}$ be the respective lifts of s and t that cross by (ii). By the definition of crossing, their respective separators $\hat{X}$ and $\hat{Y}$ are r-coupled. Thus, Lift 2.6.6 yields that $\hat{t}$ is the unique lift of t which might even cross $\hat{s}$, which shows the desired.

(iii)→(i). Let $\hat{s}, \hat{t}$ be crossing lifts of s, t with separators $\hat{X}, \hat{Y}$, respectively. Then $\hat{X}$ and $\hat{Y}$ share a vertex $\hat{v}$, or they are disjoint and there is a short cycle $\hat{O} \subseteq G_r$ that alternates between $\hat{X}$ and $\hat{Y}$. This implies that X and Y share the vertex $p_r(\hat{v})$, or they are disjoint and the short cycle $p_r(\hat{O})$ alternates between X and Y, so they are r-coupled. Thus, Lift 2.6.7 provides a lift of t given $\hat{s}$ that coincides with $\hat{t}$ by the uniqueness in Lift 2.6.6. Since $\hat{s}$ and $\hat{t}$ cross by assumption, the 'moreover'-part of Lift 2.6.7 also tells us that s and t cross. □

We will not use the following lemma, but have included it for completeness.

Lemma 2.6.9.* *Let s and t be r-local separations of a connected graph G with r-tomic separators X and Y of size $\leqslant k$, respectively, such that X and Y are r-coupled and $(k-1)r < \Delta_r(G)$. Then the four links and $X \cap Y$ together form a near-partition of $X \cup Y$, and there are no edges in G between opposite links.*

Proof. Recall that s and t can be lifted by Lift 2.5.9 and Lemma 2.6.5. Let $\hat{s}$ be an arbitrary lift of s with separator $\hat{X}$. By Lift 2.6.6, there is a unique lift $\hat{t}$ of t whose separator $\hat{Y}$ is r-coupled with $\hat{X}$. Then by Lift 2.6.7, p_r restricts to bijections between the respective pairs of links in G_r and G and between $\hat{X} \cap \hat{Y}$ and $X \cap Y$. Now let $\{\hat{A}_1, \hat{A}_2\}$ and $\{\hat{C}_1, \hat{C}_2\}$ be the separations of G_r that induce $\hat{s}$ and $\hat{t}$, respectively, by Correspondence 2.4.6. The links with regard to $\{\hat{A}_1, \hat{A}_2\}$ and $\{\hat{C}_1, \hat{C}_2\}$ together with $\hat{X} \cap \hat{Y}$ form a near-partition of $\hat{X} \cup \hat{Y}$ by definition. By Correspondence 2.6.3, the links with regard to $\hat{s}$ and $\hat{t}$ together with $\hat{X} \cap \hat{Y}$ form the same near-partition of $\hat{X} \cup \hat{Y}$. Finally, since by Lift 2.6.7 the covering map p_r restricts to bijections between the pairs of links in G_r and G and between $\hat{X} \cap \hat{Y}$ and $X \cap Y$, also the links and $X \cap Y$ together form a near-partition of $X \cup Y$.

Finally, assume for a contradiction that there is an edge x_1x_2 such that the x_i are contained in opposite X-links; the Y-links-case then follows by symmetry. Let $\hat{x}_i$ denote the lifts of the x_i in $\hat{X}$. Then $\hat{x}_1\hat{x}_2$ is an edge in G_r between opposite $\hat{X}$-links. But then $\hat{x}_1\hat{x}_2$ is an edge between $\hat{C}_1 \smallsetminus \hat{C}_2$ and $\hat{C}_2 \smallsetminus \hat{C}_1$, contradicting that $\{\hat{C}_1, \hat{C}_2\}$ is a separation. □

2.7 Local bottlenecks

In this section, we introduce and study one of the centrepieces of our theory of r-local separations: *r-local bottlenecks*. For all relevant concepts, we here follow a recurring scheme: we define the local notions, show that they agree with the respective global counterpart on r-local coverings and prove lifting/projection lemmas under a suitable assumption on the displacement of p_r. We work along this three-step process first for *r-local $\top$-stars* (Correspondence 2.7.3, Projection 2.7.11, and Lift 2.7.12), then for *r-local bottlenecks* (Correspondence 2.7.13, Projection 2.7.14, and Lift 2.7.15) and finally for a *construction* of nested sets of r-local separations representing the r-local bottlenecks (Construction 2.7.16, Correspondence 2.7.17, and Theorems 2.7.18 and 2.7.19).

2.7.1 Local $\top$-stars and correspondence in the local cover

An *r-local $\top$-star* in G is a set $\sigma = \{ (E_i, X_i, F_i) : i \in [3] \}$ of three r-local oriented separations such that

(LY1) the E_i for $i \in [3]$ form a near-partition of the edge-set $\partial(\bigcup_{i\in[3]} X_i)$,

(LY2) every vertex in some X_i is contained in at least one other X_j, and

(LY3) every r-local component at $\bigcup_{i\in[3]} X_i$ is included in some E_i.

The unoriented separations $\{E_i, X_i, F_i\}$ are the *constituents* of σ. We say that σ has *order* $\leqslant k$ if all its elements have order $\leqslant k$. If an r-local $\top$-star σ is introduced as $\{ (E_i, X_i, F_i) : i \in [3] \}$, then we denote the union of the separators X_i by X, the *centre* $\bigcap_{i\in[3]} X_i$ by Z and the *ij-link* $(X_i \cap X_j) \setminus Z$ by X_{ij} for every two distinct $i, j \in [3]$. We remark that $\{E_i, X, \partial X \setminus E_i\}$ is an r-local separation by (LY1) and (LY3).

An r-local $\top$-star $\sigma = \{ (E_i, X_i, F_i) : i \in [3] \}$ is *relevant with base* $\{E_1, X_1, F_1\}$ if $\{E_1, X_1, F_1\}$ is tight and for every $x \in X_{23}$ and for every $i \in \{2, 3\}$ there exists an x–X_{1i} edge in G or there exists an r-local component $F \subseteq E_i$ at X which contains an edge incident to x and one incident to $X_{1i} \cup Z$.

Lemma 2.7.1. *Let $\sigma = \{ (E_i, X_i, F_i) : i \in [3] \}$ be a relevant r-local $\top$-star in a connected graph G with base $\{E_1, X_1, F_1\}$. Then $X := \bigcup_{i\in[3]} X_i$ is r-tomic.*

Proof. Since $\{E_1, X_1, F_1\}$ is tight, the separator X_1 is r-tomic by Lemma 2.4.3. Now let x be a vertex in the 23-link. Since σ is relevant with base $\{E_1, X_1, F_1\}$, there is an x–X_{12} edge e or there is an r-local component $F \subseteq E_2$ at X which contains an edge incident to x and an edge incident to $X_{12} \cup Z$, where Z is the centre $\bigcap_{i\in[3]} X_i$. If e exists, then its two endvertices, i.e. x and some vertex in X_1, lie in the same r-tom by the definition of r-tom, as desired. The same is ensured by the definition of r-local component, if F exists. □

Lemma 2.7.2. *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Let $\sigma = \{ (A_i, B_i) : i \in [3] \}$ be a relevant $\top$-star in G with base $\{A_1, B_1\}$. Then $X := \bigcup_{i\in[3]} X_i$ is r-tomic.*

Proof. By Lemma 2.4.4, X_1 is r-tomic. Let Y be the r-tom of X which contains X_1. Suppose for a contradiction that there is some $x \in Y' := X \setminus Y$. By (Y2) and $X_1 \subseteq Y$, we have $Y' = X_{23} \setminus Y$. Consider $F := E(Y', X_{12} \cup (A_2 \setminus B_2))$.

We claim that there is a (not necessarily short) cycle O' in G which meets F precisely once. For $i \in \{2, 3\}$, let P_i be an x–$(X_{1i} \cup Z)$ path through $A_i \setminus B_i$ or let P_i be a path formed by a single x–X_{1i} edge, given by the relevance of σ. Since $\{A_1, B_1\}$ is tight, we may join the endvertices that the P_i have in X_1 by a path through $A_1 \setminus B_1$ to obtain the desired cycle O' in G.

As the binary cycle space of G is generated by short cycles, we infer from the existence of O' that there exists a short cycle $O \subseteq G$ which meets F oddly. If O meets F in some X_{23}–X_{12} edge, then O witnesses that some vertex in Y' is in a common r-tom with some vertex in $X_{12} \subseteq Y$, which contradicts the fact that Y is an r-tom. Thus, we may assume that O does not meet F in an X_{23}–X_{12} edge. Hence, O meets $E(Y', A_2 \setminus B_2)$ oddly. Since $X_2 = (X_2 \cap Y) \sqcup Y'$ and X_2 is the separator of the separation (A_2, B_2), we have that the cut $E(X_2, A_2 \setminus B_2)$ is bipartitioned by $E(X_2 \cap Y, A_2 \setminus B_2)$ and $E(Y', A_2 \setminus B_2)$. Since O meets the cut $E(X_2, A_2 \setminus B_2)$ evenly, it must meet $E(X_2 \cap Y, A_2 \setminus B_2)$ oddly. Thus, O witnesses that a vertex in $X_2 \cap Y$ and a vertex in $X_2 \setminus Y$ are in a common r-tom, which contradicts the fact that Y is an r-tom of X. □

Given a set σ of (possibly oriented) separations of G, we write $\sigma_r := \{s_r : s \in \sigma\}$ for the corresponding set of r-local separations of G induced by the separations in σ.

Correspondence 2.7.3 ($\top$-star)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. The map $(\sigma, s) \mapsto (\sigma_r, s_r)$ defines a bijection between the set of all pairs (σ, s) of relevant $\top$-stars σ in G with base s and the set of all pairs (τ, t) of relevant r-local $\top$-stars τ in G with base t.*

Proof. Well-defined. Let $\sigma = \{ (A_i, B_i) : i \in [3] \}$ be a relevant $\top$-star in G with base $\{A_1, B_1\}$. Let (E_i, X_i, F_i) denote the r-local separation induced by (A_i, B_i) for all i. Then $\{E_1, X_1, F_1\}$ is tight by Correspondence 2.4.8, and $X := \bigcup_{i \in [3]} X_i$ is r-tomic by Lemma 2.7.2.

Now we claim that $\sigma_r = \{ (E_i, X_i, F_i) : i \in [3] \}$ is an r-local $\top$-star. Since the (E_i, X_i, F_i) are induced by the (A_i, B_i), (Y1) yields (LY1). (LY2) follows immediately from (Y2). For (LY3), let F be an r-local component at X. Since X is r-tomic, there is a component K of $G - X$ such that $F = E(K, X)$ by Correspondence 2.4.5. As the (A_i, B_i) are separations of G which satisfy (Y1), every component of $G - X$, in particular K, has its vertex-set contained in some $A_i \setminus B_i$. Thus, $F \subseteq E_i$ as (E_i, X_i, F_i) is induced by (A_i, B_i). Hence, (LY3) holds.

Since σ is relevant with base $\{A_1, B_1\}$ and X is r-tomic, it follows immediately from Correspondence 2.4.5 that also σ_r is relevant with base $\{E_1, X_1, F_1\}$.

Injective. This is immediate from Correspondence 2.4.6.

Surjective. Let $\tau = \{ (E_i, X_i, F_i) : i \in [3] \}$ be a relevant r-local $\top$-star in G with base $t = \{E_1, X_1, F_1\}$. We have to find a pair (σ, s) with $(\sigma_r, s_r) = (\tau, t)$.

First, we show that there are separations (A_i, B_i) of G that induce the (E_i, X_i, F_i). The vertex-set $X := \bigcup_{i\in[3]} X_i$ is r-tomic by Lemma 2.7.1. Let $i \in [3]$. By Correspondence 2.4.6, the r-local separation $(E_i, X, \partial X \setminus E_i)$ is hence induced by some separation (A'_i, B'_i) of G; in particular, $E(A'_i \setminus B'_i, X) = E_i$. Now set $A_i := (A'_i \setminus B'_i) \cup X_i$ and $B_i := B'_i$. It is immediate that $A_i \setminus B_i = A'_i \setminus B'_i$, and in particular $E_G(A_i \setminus B_i, X_i) = E_i$. Thus, $E(B_i \setminus A_i, X_i) = \partial X_i \setminus E_i = F_i$. Hence, (A_i, B_i) induces (E_i, X_i, F_i).

Second, we claim that $\sigma := \{\, (A_i, B_i) : i \in [3] \,\}$ is a $\top$-star. The $(E_i, X, \partial X \setminus E_i)$ are r-local separations such that the edge-sets E_i near-partition ∂X by (LY1). Hence, the corresponding separations (A'_i, B'_i) have the vertex-sets $A'_i \setminus B'_i$ near-partition the components of $G - X$ by Correspondence 2.4.5. Since $A'_i \setminus B'_i = A_i \setminus B_i$, the vertex-sets $A_i \setminus B_i$ near-partition $V(G) \setminus X$. (LY2) immediately implies (Y2).

Finally, it remains to show that σ is relevant with base $\{A_1, B_1\}$. Since X is r-tomic, this follows from the fact that τ is relevant with base $\{E_1, X_1, F_1\}$ by Correspondence 2.4.5. We have $(\sigma, \{A_1, B_1\}) \mapsto (\tau, \{E_1, X_1, F_1\})$ by construction. □

2.7.2 Interplay of local $\top$-stars with the covering map

We define a function $f\colon \mathbb{N} \to \mathbb{N}$ such that $f(k) \cdot r/2 < \Delta_r(G)$ is a sufficient condition for the union of the separators of a relevant r-local $\top$-star to be r-locally 1-sheeted, but also satisfies $f(k) \leqslant 2(k-1)$ for every positive $k \in \mathbb{N}$:

$$f(k) := \lfloor 3k/2 \rfloor \text{ for } k \geqslant 3 \text{ and } f(k) := 0 \text{ for } k \leqslant 2. \tag{2.7}$$

Lemma 2.7.4. *$f(k) \leqslant 2(k-1)$ for every positive $k \in \mathbb{N}$.*

Proof. For $k \leqslant 2$, the statement is trivial. For $k = 3$, we have $\lfloor 3k/2 \rfloor = 4 = 2(k-1)$, and for $k \geqslant 4$ we have $2(k-1) - \lfloor 3k/2 \rfloor \geqslant k/2 - 2 \geqslant 0$. □

Lemma 2.7.5. *$f(k) \cdot r/2 < \Delta_r(G)$ implies $kr/2 < \Delta_r(G)$, for all $k, r \in \mathbb{N}$.*

Proof. If $k \leqslant 2$, then $kr/2 \leqslant r < \Delta_r(G)$, where the last inequality holds by Lemma 2.5.1. Otherwise, $k \geqslant 3$ gives $\Delta_r(G) > f(k) \cdot r/2 = \lfloor 3k/2 \rfloor \cdot r/2 \geqslant kr/2$, as desired. □

The following Lemma 2.7.6 is of the same flavour as [Car22, Lemma 3.10].

Lemma 2.7.6. *Let $\{E_1, X, E_2\}$ be a tight r-local 2-separation of a connected graph G. There exists a short cycle $O \subseteq G$ that is a union of two X-paths P_1, P_2 such that P_1 starts and ends in edges in E_1, while P_2 starts and ends in edges in E_2 or P_2 is formed by an edge that spans X.*

Proof. Let $\{x, y\} := X$. Let F_i be a tight component in E_i for $i = 1, 2$. Since being incident to x or to y induces a bipartition of F_i, there are edges $e_i^x, e_i^y \in F_i$ such that e_i^x is incident to x, e_i^y is incident to y, and there is a short cycle $O_i \subseteq G$ that includes an X-path that starts in e_i^x and ends in e_i^y. If O_1 uses an edge in E_2, or if the edge xy exists and is used by O_1, then $O := O_1$ is formed by the union of the desired paths. If O_2 exhibits the analogue property, we are done as well. Hence, we may assume that both X-paths in O_i start and end in edges in E_i,

for both i. Let P_i be a shortest X-path in O_i for both i, so the P_i have length $\leqslant r/2$. Then P_1, P_2 and $O := P_1 \cup P_2$ are as desired. □

Lemma 2.7.7. *Let σ be a relevant r-local $\top$-star of order $\leqslant k$ for $k \in [2]$ in a connected graph G. Then the union X of the separators of the elements of σ has size $\leqslant 2$ or some short cycle in G contains X.*

Proof. Let $\{ (E_i, X_i, F_i) : i \in [3] \} := \sigma$ and let $\{E_1, X_1, F_1\}$ be its base. Assume that X has size $\geqslant 3$.

We claim that $k = 2$, all X_i have size two, all links have size one and the centre is empty. If an X_i has size one, then it is included in another X_j by (LY2), and since the third X_k must also be included in $X_i \cup X_j = X_j$ by (LY2), this gives $|X| \leqslant 2$, a contradiction. So we have that $k = 2$ and all X_i have size two.

If two separators are equal, then also the third separator must be equal to these two by (LY2), contradicting that $|X| \geqslant 3$. Thus, all the X_i are pairwise distinct, and by (LY2) every two distinct X_i share an element. That means, the centre of the $\top$-star is empty, and every ij-link consists of a unique vertex which we denote by x_{ij}. So $X = \{x_{12}, x_{23}, x_{13}\}$.

Since $\{E_1, X_1, F_1\}$ is tight, we may use Lemma 2.7.6 (applied with $\{E_1, X, E_2\}$ chosen as $\{F_1, X_1, E_1\}$) to find a short cycle $O \subseteq G$ that is a union of two X_1-paths P, Q such that P starts and ends in edges in E_1 or is the path $x_{12}x_{13}$, and such that Q starts and ends in edges in F_1; in particular, O meets both x_{12} and x_{13}. Let e denote the edge on P incident to x_{12}. Let f denote the edge on Q incident to x_{12}.

If f is incident to x_{23}, then O contains X as desired. Hence, we may assume that f is not incident to x_{23}. So f is incident to x_{12}, but not to x_{13} or x_{23}, which leaves $f \in E_2$ as the only option by (LY1). We have $e \in F_2$: If the edge e is of the form $x_{12}x_{13}$, then it does not lie in ∂X, but in ∂X_2, and it hence is in F_2 by (LY1). Otherwise, the edge e lies in $E_1 \subseteq \partial X$ and is incident to $x_{12} \in X_2$, so $e \in F_2$ by (LY1). Altogether, $e \in F_2$ and $f \in E_2$ lie in different sides of the r-local separation $\{E_2, X_2, F_2\}$. So as the short cycle O contains both e and f, it must meet X_2 in both its elements and especially in x_{23}. Hence, O contains X. □

Lemma 2.7.8. *Let X be an r-tomic set of vertices in a connected graph G. If X has size $\leqslant 2$ or X is contained in a short cycle, then X is r-locally 1-sheeted.*

Proof. First, assume that X has size $\leqslant 2$. By Lemma 2.5.1, we have $\Delta_r(G) > r \geqslant k \cdot r/2$ for $k \leqslant 2$. Thus, X is r-locally 1-sheeted by Lemma 2.5.7.

Second, we may assume that X is contained in some short cycle $O \subseteq G$. Let $\hat{O}$ be a component of $p_r^{-1}(O)$, and let $\hat{X}$ consist of the lifts of the elements of X in $\hat{O}$. Since O is a short cycle, p_r restricts to an isomorphism $\hat{O} \to O$. In particular, the distance in G_r between every two elements of $\hat{X}$ is at most $r/2$, and p_r restricts to a bijection $\hat{X} \to X$. It remains to show that $\hat{X}$ is an r-tom of $p_r^{-1}(X)$. Suppose for a contradiction that there exists some short cycle which meets some $\hat{y} \in p_r^{-1}(X)$ but also $\hat{x} \in \hat{X}$. In particular, $d_{G_r}(\hat{y}, \hat{x}) \leqslant r/2$. Let $\hat{y}' \in \hat{X}$ be in the same fibre as $\hat{y}$. Then we have

$$d_{G_r}(\hat{y}, \hat{y}') \leqslant d_{G_r}(\hat{y}, \hat{x}) + d_{G_r}(\hat{x}, \hat{y}') \leqslant r/2 + r/2 = r.$$

This contradicts Lemma 2.5.1. □

Lemma 2.7.9.* *Let σ be a relevant r-local $\top$-star of order $\leqslant k$ in a connected graph G such that $f(k) \cdot r/2 < \Delta_r(G)$. Then the union X of the separators of the elements of σ is r-locally 1-sheeted.*

Proof. The vertex-set X is r-tomic by Lemma 2.7.1. If $k \leqslant 2$, then X is r-locally 1-sheeted by Lemma 2.7.7 and Lemma 2.7.8. So we may assume that $k \geqslant 3$. By (LY2), X has size $\leqslant \lfloor 3k/2 \rfloor$. So since $\lfloor 3k/2 \rfloor \cdot r/2 = f(k) \cdot r/2 < \Delta_r(G)$, X is r-locally 1-sheeted by Lemma 2.5.7. □

Lemma 2.7.10.* *Let G be a connected graph. Let $\hat{\sigma}$ be a relevant r-local $\top$-star of order $\leqslant k$ in G_r such that $f(k) \cdot r/2 < \Delta_r(G)$. Then the projection $X := p_r(\hat{X})$ of the union $\hat{X}$ of the separators of the elements of $\hat{\sigma}$ is r-locally 1-sheeted. Moreover, $\hat{X}$ is an r-tom of $p_r^{-1}(X)$.*

Proof. The vertex-set $\hat{X}$ is r-tomic in G_r by Lemma 2.7.1. Hence, X is r-tomic by Projection 2.4.9.

If $k \leqslant 2$, then $\hat{X}$ and thus also X either has size $\leqslant 2$ or is contained in a short cycle by Lemma 2.7.7. Hence, X is r-locally 1-sheeted by Lemma 2.7.8.

So we may assume $k \geqslant 3$. By (LY2), $\hat{X}$ has size $\leqslant \lfloor 3k/2 \rfloor$, and so does its projection X. Thus, X is r-locally 1-sheeted by Lemma 2.5.7 and $\lfloor 3k/2 \rfloor \cdot r/2 = f(k) \cdot r/2 < \Delta_r(G)$.

In both cases, $\hat{X}$ is indeed an r-tom of $p_r^{-1}(X)$, since $\hat{X}$ is r-tomic and $X = p_r(\hat{X})$ is r-locally 1-sheeted. □

Projection 2.7.11 (Local $\top$-star)**.*** *Let G be a connected graph. Let $\hat{\sigma}$ be a relevant r-local $\top$-star in G_r of order $\leqslant k$ with base $\hat{s}$ such that $f(k) \cdot r/2 < \Delta_r(G)$. Then $p_r(\hat{\sigma})$ is a relevant r-local $\top$-star in G of order $\leqslant k$ with base $p_r(\hat{s})$.*

Proof. Let $\{ \hat{s}_i : i \in [3] \} := \hat{\sigma}$ where $\hat{s}_i =: (\hat{E}_i, \hat{X}_i, \hat{F}_i)$ so that $\hat{s} =: \hat{s}_1$. Let $\hat{X} := \bigcup_{i \in [3]} \hat{X}_i$. By Lemma 2.7.10, $X := p_r(\hat{X})$ is r-locally 1-sheeted and $\hat{X}$ is an r-tom of $p_r^{-1}(X)$. Hence, the subsets $\hat{X}_i$ of $\hat{X}$ are r-locally clopen lifts of the X_i. Thus, by Projection 2.4.16, each $\hat{s}_i$ projects to an r-local separation $s_i = (E_i, X_i, F_i)$ of G. Moreover, the projection $s := \{E_1, X_1, F_1\}$ of $\hat{s}$ is tight by Projection 2.4.16.

Now (LY1) and (LY2) for $p_r(\hat{\sigma}) = \{ s_i : i \in [3] \}$ follow immediately from (LY1) and (LY2) of $\hat{\sigma}$, as X is r-locally 1-sheeted. Moreover, (LY3) and the relevance of $\hat{\sigma}$ with base $\hat{s}$ imply, using Lift 2.4.14, that also $p_r(\hat{\sigma})$ satisfies (LY3) and is relevant with base s. □

Lift 2.7.12 (Local $\top$-star)**.*** *Let G be a connected graph. Let $\sigma = \{ (E_i, X_i, F_i) : i \in [3] \}$ be a relevant r-local $\top$-star in G of order $\leqslant k$ with base $\{E_1, X_1, F_1\}$ such that $f(k) \cdot r/2 < \Delta_r(G)$. Let $\hat{X}$ be an r-tom of $p_r^{-1}(X)$. Then there are unique lifts $(\hat{E}_i, \hat{X}_i, \hat{F}_i)$ of the (E_i, X_i, F_i) with $\hat{X}_i \subseteq \hat{X}$ via Lift 2.4.17, and $\hat{\sigma} := \{ (\hat{E}_i, \hat{X}_i, \hat{F}_i) : i \in [3] \}$ is a relevant r-local $\top$-star in G_r of order $\leqslant k$ with base $\{\hat{E}_1, \hat{X}_1, \hat{F}_1\}$.*

Proof. Recall that $X := \bigcup_{i \in [3]} X_i$ is r-locally 1-sheeted by Lemma 2.7.9. Thus, each X_i has a unique lift $\hat{X}_i \subseteq \hat{X}$, which is r-locally clopen since $\hat{X}$ is r-locally clopen as an r-tom of

$p_r^{-1}(X)$. Hence, the lifts $\hat{s}_i = (\hat{E}_i, \hat{X}_i, \hat{F}_i)$ exist by Lift 2.4.17 and are unique. As X is r-locally 1-sheeted, it is immediate that (LY1) and (LY2) transfer from σ to $\hat{\sigma} := \{\, \hat{s}_i : i \in [3] \,\}$. Note that Lift 2.4.17 additionally ensures that $\{\hat{E}_1, \hat{X}_1, \hat{F}_1\}$ is tight. Lift 2.4.14 yields that also (LY3) and the relevance transfers from σ to $\hat{\sigma}$. □

2.7.3 Local bottlenecks and nested sets of local separations

Let $k \in \mathbb{N}$. An *r-local bottleneck of order k* in G, or *r-local k-bottleneck*, is a non-empty set β of tight r-local k-separations of G such that β satisfies the property (B_r):

(B_r) Whenever β contains the base of a relevant r-local $\top$-star σ in G of order $\leqslant k$, then β contains at least one of the other two constituents of σ.

Correspondence 2.7.13 (Local bottlenecks)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. The map $\beta \mapsto \beta_r := \{s_r : s \in \beta\}$ is a bijection between the k-bottlenecks β in G and the r-local k-bottlenecks in G.*

Proof. Correspondence 2.7.3 ensures that β satisfies (B) if and only if β_r satisfies (B_r). □

Projection 2.7.14 (Local bottlenecks)**.*** *Let G be a connected graph. Let $\hat{\beta}$ be an r-local k-bottleneck in G_r such that $f(k) \cdot r/2 < \Delta_r(G)$. Then $p_r(\hat{\beta})$ is an r-local k-bottleneck in G.*

Proof. Since the r-local separations in $\hat{\beta}$ are tight, their separators are r-tomic by Lemma 2.4.3. Hence, by Projection 2.5.8 and Lemma 2.7.5, we may use Projection 2.4.16 to find that $\beta := p_r(\hat{\beta})$ is well-defined and consists of tight r-local k-separations of G. In particular, β is non-empty.

To show that β exhibits (B_r), let σ be a relevant r-local $\top$-star of order $\leqslant k$ with base $s \in \beta$. Take any $\hat{s} \in \hat{\beta}$ with $p_r(\hat{s}) = s$. Let X denote the union of the separators of the elements of σ, and let $\hat{X}$ be the r-tom of $p_r^{-1}(X)$ that includes the r-tomic separator of $\hat{s}$. Then by Lift 2.7.12 the unique lifts of the elements of σ to G_r whose separators are included in $\hat{X}$ form a relevant r-local $\top$-star $\hat{\sigma}$ in G_r of order $\leqslant k$ with base $\hat{s}$; Lemma 2.7.4 ensures the inequality required for Lift 2.7.12. By (B_r) for $\hat{\beta}$, some constituent $\hat{t}$ of $\hat{\sigma}$ other than $\hat{s}$ is contained in $\hat{\beta}$. So the corresponding constituent $p_r(\hat{t}\,)$ of σ is in β and distinct from s. Hence, β satisfies (B_r). □

Lift 2.7.15 (Local bottlenecks)**.*** *Let G be a connected graph. Let β be an r-local k-bottleneck in G such that $f(k) \cdot r/2 < \Delta_r(G)$. Then all elements of β can be lifted, and the set $\hat{\beta}$ of all possible lifts of all r-local separations in β to G_r is an r-local k-bottleneck in G_r.*

Proof. Since β is an r-local k-bottleneck, all elements of β are tight r-local k-separations, which thus have r-tomic separators by Lemma 2.4.3. By Lift 2.5.9 and Lemma 2.7.5, all elements of β have r-locally 1-sheeted separators and can be lifted to G_r. All these lifts, and hence all elements of $\hat{\beta}$, are then tight r-local k-separations of G_r by Lift 2.4.17; in particular, $\hat{\beta}$ is non-empty. To show that $\hat{\beta}$ exhibits (B_r), let $\hat{\sigma}$ be a relevant r-local $\top$-star in G_r of order $\leqslant k$ with base $\hat{s} \in \hat{\beta}$. Then $s := p_r(\hat{s})$ is in β by definition of $\hat{\beta}$. The set σ that consists of the projections of the elements of $\hat{\sigma}$ is a relevant r-local $\top$-star in G of order $\leqslant k$ with base s by

Projection 2.7.11. By (B_r) for β, we know that β contains a second constituent t of σ besides s. Then $\hat{\beta}$ contains all lifts of t to G_r; in particular, the one that is a constituent of $\hat{\sigma}$. Hence, $\hat{\beta}$ satisfies (B_r). □

Construction 2.7.16 (Nested local separations from local bottlenecks)**.** Let G be a graph and $r \geqslant 0$ an integer. Let $\mathcal{B}_r^k(G)$ be the set of all r-local k-bottlenecks in G and $X_r^k := \bigcup \mathcal{B}_r^k(G)$ for every $k \in \mathbb{N}$. We recursively define for every $k \in \mathbb{N}$ a set $N_r^k(G)$, as follows. Let $x_r^k \colon X_r^k \to \mathbb{N}$ assign to each r-local separation $s \in X_r^k$ the number of r-local separations in X_r^k that s crosses, and assume that x_r^k is well-defined. Write $N_r^{<k}(G) := \bigcup_{j<k} N_r^j(G)$. For each $\beta \in \mathcal{B}_r^k(G)$, let $N_r^k(G, \beta)$ consist of the r-local separations $s \in \beta$ which are nested with $N_r^{<k}(G)$ and among those have minimal $x^k(s)$. Finally, set $N_r^k(G) := \bigcup_{\beta \in \mathcal{B}_r^k} N_r^k(G, \beta)$.

Correspondence 2.7.17 (Nested set)**.** *Let G be a locally finite, connected graph. Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$. Then the sets $N_r^k(G)$ are well-defined and consist precisely of all r-local separations s_r induced by the separations s in $N^k(G)$ given by Construction 2.3.6.*

Proof. The map $\beta \mapsto \beta_r$ given by Correspondence 2.7.13 is a bijection from the set $\mathcal{B}^k$ of all k-bottlenecks in G to the set $\mathcal{B}_r^k$ of all r-local k-bottlenecks in G. In particular, the map $s \mapsto s_r$ from Correspondence 2.4.8 restricts to a bijection $\bigcup \mathcal{B}^k \to \bigcup \mathcal{B}_r^k$.

We prove by induction on $k \in \mathbb{N}$ that $s \mapsto s_r$ restricts to a bijection from $N^k(G, \beta)$ to $N_r^k(G, \beta_r)$ for every $\beta \in \mathcal{B}^k$. The base case $k = 0$ is trivial as there are no r-local 0-bottlenecks. Assume now that $k > 0$. Then it follows from the induction hypothesis that the map $s \mapsto s_r$ restricts to a bijection $N^{<k}(G) \to N_r^{<k}(G)$. Hence, an $s \in \beta \in \mathcal{B}^k$ is nested with $N^{<k}(G)$ if and only if its r-local counterpart $s_r \in \beta_r \in \mathcal{B}_r^k$ is nested with $N_r^{<k}(G)$, by Correspondence 2.6.4. Moreover, Correspondence 2.6.4 also yields that an $s \in \bigcup \mathcal{B}^k$ crosses a $t \in \bigcup \mathcal{B}^k$ if and only if $s_r \in \bigcup \mathcal{B}_r^k$ crosses $t_r \in \bigcup \mathcal{B}_r^k$. In particular, s_r crosses only finitely many such t_r by Lemma 2.3.8, and thus x_r^k is well-defined. Moreover, $s \mapsto s_r$ restricts to a bijection $N^k(G, \beta) \to N_r^k(G, \beta_r)$ for every $\beta \in \bigcup \mathcal{B}^k$. Hence, the desired statement holds. □

Theorem 2.7.18 (Nested set)**.*** *Let G be a locally finite, connected graph. Assume that $(k-1)r < \Delta_r(G)$ for a $k \in \mathbb{N}$. Then both $N_r^k(G)$ and $N_r^k(G_r)$ are well-defined, and $N_r^k(G_r)$ consists precisely of the lifts to G_r of the elements of $N_r^k(G)$.*

Proof. Recall that $(k-1)r < \Delta_r(G)$ implies $f(k) \cdot r/2 < \Delta_r(G)$ by Lemma 2.7.4. For every $j \leqslant k$, p_r restricts to a surjection $\hat{\beta} \mapsto p_r(\hat{\beta})$ from the set $\hat{\mathcal{B}}_r^j$ of all r-local j-bottlenecks in G_r to the set $\mathcal{B}_r^j$ of all r-local j-bottlenecks in G; the map is well-defined by Projection 2.7.14 and is surjective by Lift 2.7.15. Moreover, Lift 2.7.15 ensures that $\bigcup \hat{\mathcal{B}}_r^j$ is the set of all lifts of the elements of $\bigcup \mathcal{B}_r^j$ for every $j \leqslant k$.

We prove the claim by induction on k. The base case $k = 0$ is trivial because there are no r-local 0-bottlenecks. So suppose that $k > 0$. It follows from the induction hypothesis that $N_r^{<k}(G_r)$ is the set of all lifts of the elements of $N_r^{<k}(G)$. Hence, $\hat{s} \in \hat{\beta} \in \hat{\mathcal{B}}_r^k$ is nested with $N_r^{<k}(G_r)$ if and only if $p_r(\hat{s}) \in p_r(\hat{\beta}) \in \mathcal{B}_r^k$ is nested with $N_r^{<k}(G)$ by Lift 2.6.8 (i) and (ii).

Moreover, Lift 2.6.8 (i) and (iii) yield that $s \in \bigcup \mathcal{B}_r^k$ crosses $t \in \bigcup \mathcal{B}_r^k$ if and only if every lift $\hat{s} \in \bigcup \hat{\mathcal{B}}_r^k$ of s crosses precisely one lift $\hat{t} \in \bigcup \hat{\mathcal{B}}_r^k$ of t. Thus, the numbers $x_r^k(s)$ and $x_r^k(\hat{s})$ in Construction 2.7.16 applied to G and G_r coincide. In particular, the former is well-defined, as the latter is well-defined by Correspondence 2.7.17. Moreover, $\hat{s} \mapsto p_r(\hat{s})$ restricts to a surjection $N_r^k(G_r, \hat{\beta}) \to N_r^k(G, p_r(\hat{\beta}))$ for every $\hat{\beta} \in \hat{\mathcal{B}}_r^k$ and, conversely, $N_r^k(G_r, \hat{\beta})$ is the set of all lifts of elements of $N_r^k(G_r, \beta)$ for every $\beta \in \mathcal{B}_r^k$ where $\hat{\beta} \in \hat{\mathcal{B}}_r^k$ is the set of all lifts of elements of β (which is an r-local k-bottleneck by Lift 2.7.15). All in all, $N_r^k(G_r)$ is the set of all lifts of elements of $N_r^k(G)$. □

Theorem 2.7.19.* *Let G be a locally finite, connected graph. Assume that $(k-1)r < \Delta_r(G)$ for a $k \in \mathbb{N}$. Then $N_r^{\leqslant k}(G)$ meets every r-local bottleneck in G of order $\leqslant k$.*

Proof. Let β_r be an r-local bottleneck in G of order $j \leqslant k$. By Lift 2.7.15, the set $\hat{\beta}_r$ of all lifts of all elements of β_r is an r-local j-bottleneck in G_r. There is a unique bottleneck $\hat{\beta}$ in G_r such that $s \mapsto s_r$ is a bijection $\hat{\beta} \to \hat{\beta}_r$, by Correspondence 2.7.13. Then there is $\hat{s} \in \hat{\beta} \cap N^{\leqslant k}(G_r)$ by Theorem 2.3.7. So $\hat{s}_r \in \hat{\beta}_r \cap N_r^{\leqslant k}(G_r)$ by Correspondence 2.7.17. Since $N_r^{\leqslant k}(G_r)$ is equal to the set of all lifts of all elements of $N_r^{\leqslant k}(G)$ by Theorem 2.7.18, we have $p_r(\hat{s}_r) \in \hat{\beta} \cap N_r^{\leqslant k}(G)$. □

2.8 Constructing graph-decompositions from nested sets of local separations

With the construction of nested sets of r-local separations representing the r-local bottlenecks of a graph at hand (see Section 2.7), we need just one more ingredient to obtain graphs of r-local bottlenecks: the construction of a graph-decomposition from a nested set of r-local separations. For this, we cannot simply imitate the construction of tree-decompositions from nested sets of global separations (Construction 2.2.3), since it heavily relies on the partial order on global separations, and it is not clear to us what a meaningful r-local version of this partial order could be.

To overcome this obstacle, we develop an alternative perspective on Construction 2.2.3 by introducing the concept of *cutouts* as an alternative to splitting stars (Section 2.8.1), which we then use to find a suitable r-local framework: We give the respective r-local definitions in Section 2.8.2, where we also show that they agree with their global counterparts on r-local coverings (Correspondence 2.8.12). We then prove the respective projection/lifting lemmas along the covering map p_r (Projection 2.8.14, Lift 2.8.15, and Proposition 2.8.16). Finally, we describe our construction of graph-decompositions from nested sets of r-local separations (Construction 2.8.17), which we show to agree with Construction 2.2.3 for nested sets of global separations on r-local coverings (Correspondence 2.8.18). Our construction also commutes with p_r (Theorem 2.8.21): The graph-decomposition from a nested set N of r-local separations is isomorphic to the one obtained by taking the set $\hat{N}$ of all lifts of the elements of N to G_r, then applying our construction to $\hat{N}$ and folding the resulting tree-decomposition back along p_r via Construction 2.2.6.

2.8.1 Splitting stars from the perspective of a relation

For a separation $s = (A_1, A_2)$ of G we denote its separator by $X(s)$ and define

$$L(s) := E_G(A_1 \setminus A_2, X) = E_G(A_1 \setminus A_2, A_2) \quad \text{and} \quad R(s) := E_G(X, A_2 \setminus A_1) = E_G(A_1, A_2 \setminus A_1).$$

We remark that $L(s)$ and $R(s)$ refer to the *left* and *right* sides of the local separations induced by the separation s, whose names are purely based on the reading direction. Let N be a nested set of separations of a graph G. We define a relation $\Sigma^{\vec{x}}_N$ on $\mathrm{ori}(N)$ by letting

$$\begin{aligned} s_1 \Sigma^{\vec{x}}_N s_2 :\Leftrightarrow\ & X(s_1) \text{ and } X(s_2) \text{ share a vertex } v \text{ and } s_1 > s_2^*, \\ & \text{and there is no } s \in \mathrm{ori}(N) \text{ with } v \in X(s) \text{ and } s_1 > s > s_2^*. \end{aligned}$$

The relation $\Sigma^{\vec{x}}_N$ is symmetric.

Lemma 2.8.1. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N)$ is a tree-decomposition. Let $s_1, s_2 \in \mathrm{ori}(N)$ with $s_1 \Sigma^{\vec{x}_N} s_2$. Then s_1 and s_2 lie in the same splitting star of N.*

Proof. By assumption, there is a vertex $v \in X(s_1) \cap X(s_2)$ and we have $s_1 > s_2^*$, say, and there is no $s \in \mathrm{ori}(N)$ with $v \in X(s)$ and $s_1 > s > s_2^*$. Suppose for a contradiction that s_1 and s_2 lie in distinct splitting stars of N. Then $s_1 > s_2^*$ with Lemma 2.2.5 yields that there is $s \in \mathrm{ori}(N)$ with $s_1 > s > s_2^*$. Then the vertex v is contained in $X(s)$, a contradiction. □

Let N be a nested set of separations of a graph G. For $s \in \mathrm{ori}(N)$ we write

$$R_N(s) := R(s) \setminus \bigcup \{\, R(s') : s' \in \mathrm{ori}(N) \text{ such that } X(s') \text{ meets } X(s) \text{ and } s > s' \,\}$$

Lemma 2.8.2. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N)$ is a tree-decomposition of G. Let σ be a splitting star of N and $s \in \sigma$. Then $R_N(s) = E(X(s), V_\sigma - X(s))$.*

Proof. Let $E := E(X(s), V_\sigma - X(s))$. The backward inclusion $R_N(s) \supseteq E$ is straightforward. For the forward inclusion $R_N(s) \subseteq E$, let $e = xy \in R(s) \setminus E$ with $x \in X(s)$; we have to show $e \notin R_N(s)$. Let $\tau \in V(T)$ be closest to σ with $y \in V_\tau$. As $\tau \neq \sigma$, by Lemma 2.2.5 there is $s' = (A, B) \in \sigma$ with $y \in A \setminus B$. Since $x \in X(s) \subseteq V_\sigma \subseteq B$, the fact that s' is a separation implies $x \in X(s')$, so $e \in L(s') = R((s')^*)$. Then $e \notin R_N(s)$ as $s, s' \in \sigma$ gives $s > (s')^*$. □

Let N be a nested set of separations of a graph G. We denote by $\mathcal{X} := \mathcal{X}(N)$ the set of all separators of elements of N. We denote by $\mathscr{K} := \mathscr{K}(N)$ the set of components of $G - \bigcup \mathcal{X}$ and define the relation $\Sigma^{\mathscr{K}}_N$ on $\mathrm{ori}(N)$ by letting

$$s_1 \Sigma^{\mathscr{K}}_N s_2 :\Leftrightarrow \text{there is a component } K \in \mathscr{K} \text{ such that } \partial K \text{ meets } R_N(s_i) \text{ for both } i = 1, 2.$$

Note that $\Sigma^{\mathscr{K}}_N$ is symmetric.

Lemma 2.8.3. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N) =: (T, \mathcal{V})$ is a tree-decomposition of G. Let $K \in \mathscr{K}(N)$. Then there is a unique node σ of T with $K \subseteq V_\sigma$. Every $s \in \mathrm{ori}(N)$ such that $R_N(s)$ meets ∂K lies in σ.*

Proof. As K is a component of $G - \bigcup \mathcal{X}(N)$ and the edges of T correspond to the separations in N, we can orient every edge of T in the direction of K to find a sink σ of T which satisfies $K \subseteq V_\sigma$. Then σ is unique with this property. Now assume for a contradiction that $s \notin \sigma$. As K meets $R_N(s) \subseteq R(s)$, we find with Lemma 2.2.5 a unique $t \in \sigma$ with $s > t$. The separator $X(t)$ separates $X(s)$ from K. As $R_N(s)$ and ∂K share an edge xk with $x \in X(s)$ and $k \in V(K)$, and since $X(t)$ is disjoint from K, it follows that the vertex x must lie in $X(t)$. So $X(s)$ and $X(t)$ share the vertex x and $s > t$. But $k \in V(K)$ means that the edge xk lies in $R(t)$, contradicting that $xk \in R_N(s)$. □

Corollary 2.8.4. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N)$ is a tree-decomposition. Let $s_1, s_2 \in \mathrm{ori}(N)$ with $s_1 \Sigma^{\mathscr{K}}_N s_2$. Then s_1 and s_2 lie in the same splitting star of N.* □

An edge e of G *crosses* a collection $\mathcal{Y}$ of vertex-sets of G if both ends of e are contained in $\bigcup \mathcal{Y}$ but every $Y \in \mathcal{Y}$ contains at most one end of e. Let N be a set of separations of G and $\mathcal{X} = \mathcal{X}(N)$. We define the relation $\Sigma^{\mathcal{X}}_N$ on $\mathrm{ori}(N)$ by letting

$$s_1 \Sigma^{\mathcal{X}}_N s_2 :\Leftrightarrow \text{there is an } \mathcal{X}\text{-crossing edge of } G \text{ that lies in } R_N(s_i) \text{ for both } i = 1, 2.$$

Note that $\Sigma^{\mathcal{X}}_N$ is symmetric.

Lemma 2.8.5. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N)$ is a tree-decomposition. Let e be an $\mathcal{X}(N)$-crossing edge of G. Then there is a unique node σ of T with $e \in E(V_\sigma)$.*

Proof. As e crosses $\mathcal{X}(N)$ and the edges of T correspond to the separations in N, we can orient every edge of T in the direction of e to find a sink σ of T which satisfies $e \in E(V_\sigma)$. Then σ is unique with this property. □

Corollary 2.8.6. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N)$ is a tree-decomposition. Let $s_1, s_2 \in \mathrm{ori}(N)$ with $s_1 \Sigma^{\mathcal{X}}_N s_2$. Then s_1 and s_2 lie in the same splitting star of N.*

Proof. This follows immediately from Lemma 2.8.2 and Lemma 2.8.5. □

We denote by Σ_N the transitive reflexive closure of the union of all three symmetric relations $\Sigma^{\vec{x}_N}$, $\Sigma^{\mathscr{K}}_N$ and $\Sigma^{\mathcal{X}}_N$. Then Σ_N is an equivalence relation on $\mathrm{ori}(N)$. The equivalence classes of Σ_N are the *cutouts* of N.

Proposition 2.8.7. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N)$ is a tree-decomposition of G. Then the cutouts of N are precisely the splitting stars of N.*

Proof. Every cutout of N is included in a splitting star of N by Lemma 2.8.1, Corollary 2.8.4 and Corollary 2.8.6.

For the converse, let σ be a splitting star of N, and assume for a contradiction that σ is not included in a cutout of N. Then, using that every cutout of N is included in a splitting star of N, we find a bipartition $\{\sigma_1, \sigma_2\}$ of σ such that the elements of σ_1 are not related to any elements of σ_2 (with regard to the relation Σ_N). Let X_i denote the union of all separators of separations in σ_i, for $i = 1, 2$. Let P be a shortest X_1–X_2 path in G with ends $x_i \in X_i$. For each i, choose $s_i \in \sigma_i$ with $x_i \in X(s_i)$. We distinguish three cases. For all of them, let $(T, \mathcal{V}) := \mathcal{T}(N)$ and $\mathcal{X} := \mathcal{X}(N)$.

First, assume that P has at least one internal vertex. Then X_1 and X_2 are disjoint and there is no X_1–X_2 edge. It follows from σ being a star that all internal vertices of P lie in $V_\sigma - (X_1 \cup X_2)$. So there is a unique component K of $G - \bigcup \mathcal{X}$ that contains all internal vertices of P, and $K \subseteq V_\sigma$. Let e_i denote the edge on P that is incident to x_i for both i. Then $e_i \in R_N(s_i)$ by Lemma 2.8.2. Hence, the component K witnesses $s_1 \Sigma_N^{\mathscr{K}} s_2$, which contradicts the choice of $\{\sigma_1, \sigma_2\}$.

Second, assume that P has a single edge x_1x_2. Then X_1 and X_2 are disjoint. Hence, x_1x_2 is $\mathcal{X}$-crossing and is an edge of V_σ. Then $x_1x_2 \in R_N(s_i)$ for both i by Lemma 2.8.2. So the edge x_1x_2 witnesses $s_1 \Sigma_N^{\mathcal{X}} s_2$, which contradicts the choice of $\{\sigma_1, \sigma_2\}$.

Finally, assume that P has only one vertex $x_1 = x_2$. Then $X(s_1)$ and $X(s_2)$ share this vertex, and since σ is a splitting star there is no $s \in \mathrm{ori}(N)$ with $s_1 > s > s_2^*$ by Lemma 2.2.5. So $s_1 \Sigma_N^{\vec{x}} s_2$ contradicts the choice of $\{\sigma_1, \sigma_2\}$.

As we have derived contradictions in all cases, the proof is complete. □

Let σ be a cutout of a set N of separations of G. The *part* of σ is the graph that is obtained from the union of the following ingredients:

(i) the induced subgraph $G[X(s)]$ for every $s \in \sigma$;

(ii) the components $K \in \mathscr{K}(N)$ such that ∂K shares an edge with $R_N(s)$ for some $s \in \sigma$; and

(iii) the edges in $R_N(s)$ for every $s \in \sigma$.

We denote the part of σ by $\mathrm{part}(\sigma)$.

For the next lemma, recall that the *interior* of a splitting star σ of N is defined as $\mathrm{int}(\sigma) := G[\bigcap_{(A,B)\in\sigma} B]$. In particular, $X(s) \subseteq V(\mathrm{int}(\sigma))$ for all $s \in \sigma$. In the context of $\mathcal{T}(N)$ we have $\mathrm{int}(\sigma) = V_\sigma$.

Proposition 2.8.8. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N)$ is a tree-decomposition of G. Let σ be both a splitting star and a cutout of N. Then* $\mathrm{int}(\sigma) = \mathrm{part}(\sigma)$.

Proof. Let $I := \mathrm{int}(\sigma)$ and $P := \mathrm{part}(\sigma)$.

We first show the inclusion $P \subseteq I$. All induced subgraphs $G[X(s)]$ for $s \in \sigma$ are included in I since $I = V_\sigma$. If K is a component in $\mathscr{K}(N)$ such that ∂K shares an edge with $R_N(s)$ for some $s \in \sigma$, then $K \subseteq V_\sigma = I$ by Lemma 2.8.3. All edges in $R_N(s)$ for all $s \in \sigma$ are edges in I by Lemma 2.8.2.

We now show the converse inclusion $I \subseteq P$. The induced subgraphs $G[X(s)]$ for $s \in \sigma$ are included in P by definition. Let $\mathscr{K}_\sigma$ denote the subset of $\mathscr{K}(N)$ formed by the components K with $K \subseteq V_\sigma$. Fix any $K \in \mathscr{K}_\sigma$. Since G is connected, there is an edge $e \in \partial K$. With the help of $\mathcal{T}(N)$ we find $s \in \sigma$ with $e \in R(s)$. Then $e \in R_N(s)$ by Lemma 2.8.2. So $K \subseteq P$ by definition. It remains to show that every edge e of I that is not contained in a $G[X(s)]$ or a $K \in \mathscr{K}_\sigma$ lies in P. So e either crosses $\mathcal{X}$, or e joins a vertex in some $X(s)$ to some $K \in \mathscr{K}_\sigma$. If e crosses $\mathcal{X}$, then e joins a vertex in $X(s_1)$ to a vertex in $X(s_2) \setminus X(s_1) \subseteq V_\sigma \setminus X(s_1)$ for some two distinct $s_1, s_2 \in \sigma$, so $e \in R_N(s_1)$ by Lemma 2.8.2 and hence $e \in E(P)$. If e joins a vertex in $X(s)$ for some $s \in \sigma$ to some $K \in \mathscr{K}_\sigma$, then $e \in E(X(s), V_\sigma - X(s))$ as $K \subseteq V_\sigma - X(s)$, so $e \in R_N(s)$ by Lemma 2.8.2 and hence $e \in E(P)$. □

2.8.2 Cutouts make possible a local version of splitting stars

Let $s = (E_1, X, E_2)$ and $t = (F_1, Y, F_2)$ be r-local separations of G with r-tomic separators. Inspired by Lemma 2.2.2, we set $s \succeq t$ if

- X and Y are r-coupled, and
- all three of the following sets are empty: the X-link for F_2, the Y-link for E_1, and the edge-set $E_1 \cap F_2 \cap \partial(X \cap Y)$.

We remark that $s \succeq t$ if and only if $s^* \preceq t^*$.

Correspondence 2.8.9 ($\geqslant$)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Let s and t be oriented separations of G with r-tomic separators X and Y, respectively, such that X and Y are r-coupled. Then $s \geqslant t$ if and only if $s_r \succeq t_r$.*

Proof. This follows from Correspondence 2.6.3 and the 'moreover'-part of Lemma 2.2.2. □

Lift 2.8.10 ($\geqslant$)**.*** *Let s and t be oriented r-local separations of a connected graph G with r-tomic separators X and Y, respectively, of size $\leqslant k$ such that X and Y are r-coupled and $(k-1)r < \Delta_r(G)$. Let $\hat{s}$ be a lift of s with separator $\hat{X}$, and let $\hat{t}$ be the unique lift of t with separator $\hat{Y}$ such that $\hat{X}$ and $\hat{Y}$ are r-coupled, which exists by Lift 2.6.6. Then $s \succeq t$ if and only if $\hat{s} \succeq \hat{t}$.*

Proof. This follows immediately from Lift 2.6.7. □

For an r-local separation $s = (E_1, X, E_2)$, we define $X(s) := X$, $L(s) := E_1$ and $R(s) := E_2$. Let N be a set of r-local separations of a graph G. We denote by $\mathrm{ori}(N)$ the set of all

orientations of r-local separations in N. We define a relation $\Sigma^{\vec{x}}_N$ on $\mathrm{ori}(N)$ by letting

$$s_1 \Sigma^{\vec{x}}_N s_2 :\Leftrightarrow X(s_1) \text{ and } X(s_2) \text{ share a vertex } v \text{ and } s_1 \succ s_2^*, \\ \text{and there is no } s \in \mathrm{ori}(N) \text{ with } v \in X(s) \text{ and } s_1 \succ s \succ s_2^*.$$

Note that $\Sigma^{\vec{x}}_N$ is symmetric. We refrain from adding the parameter r to the notation $\Sigma^{\vec{x}}_N$ as this notation already is overloaded. In practice, we will know whether the s_i are usual or r-local separations, and this will uniquely determine which of the two relations we refer to.

For $s \in \mathrm{ori}(N)$ we write

$$R_N(s) := R(s) \setminus \bigcup \{\, R(s') : s' \in \mathrm{ori}(N) \text{ such that } X(s') \text{ meets } X(s) \text{ and } s \succ s' \,\}$$

Correspondence 2.8.11. *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Let N be a set of oriented separations of G with r-tomic separators and write $N_r := \{s_r : s \in N\}$. For every $s \in \mathrm{ori}(N)$ we have $R_N(s) = R_{N_r}(s_r)$.*

Proof. This follows immediately from Correspondence 2.8.9. □

We denote by $\mathcal{X} = \mathcal{X}(N)$ the set of all separators of the r-local separations in N. We denote the set of components of $G - \bigcup \mathcal{X}$ by $\mathscr{K} := \mathscr{K}(N)$ and define the relation $\Sigma^{\mathscr{K}}_N$ on $\mathrm{ori}(N)$ by letting

$$s_1 \Sigma^{\mathscr{K}}_N s_2 :\Leftrightarrow \text{there is a component } K \in \mathscr{K} \text{ such that } \partial K \text{ meets } R_N(s_i) \text{ for both } i = 1, 2.$$

Note that $\Sigma^{\mathscr{K}}_N$ is symmetric. We define the relation $\Sigma^{\mathcal{X}}_N$ on $\mathrm{ori}(N)$ by letting

$$s_1 \Sigma^{\mathcal{X}}_N s_2 :\Leftrightarrow \text{there is an } \mathcal{X}\text{-crossing edge of } G \text{ that lies in } R_N(s_i) \text{ for both } i = 1, 2.$$

Note that $\Sigma^{\mathcal{X}}_N$ is symmetric. As for usual separations, we denote by Σ_N the transitive reflexive closure of the union of all three symmetric relations $\Sigma^{\vec{x}_N}$, $\Sigma^{\mathscr{K}}_N$ and $\Sigma^{\mathcal{X}}_N$. Then Σ_N is an equivalence relation on $\mathrm{ori}(N)$. We refer to the equivalence classes of Σ_N as *cutouts* of N. The *part* of a cutout σ of N is defined just like for a cutout of a set of usual separations, and we denote it by $\mathrm{part}(\sigma)$.

Correspondence 2.8.12. *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Let N be a set of proper separations of G with r-tomic separators.*

(i) $s \Sigma^*_N t \iff s_r \Sigma^*_{N_r} t_r$ *for all* $s, t \in \mathrm{ori}(N)$ *and* $* \in \{\vec{x}, \mathscr{K}, \mathcal{X}\}$.

(ii) $\sigma \mapsto \sigma_r = \{s_r : s \in \sigma\}$ *is a bijection between the cutouts of N and the cutouts of N_r.*

(iii) *The part of σ is equal to the part of σ_r.*

Proof. Note that $\mathcal{X}(N) = \mathcal{X}(N_r)$ and $\mathscr{K}(N) = \mathscr{K}(N_r)$. (i). For $* = \vec{x}$ this is immediate from Correspondence 2.8.9. For $* \in \{\mathcal{X}, \mathscr{K}\}$ this is immediate from Correspondence 2.8.11. (i) implies (ii). Then (iii) follows with the help of Correspondence 2.8.11. □

2.8.3 Interplay with the covering map

Projection 2.8.13.* *Let N be a nested set of r-local separations of a connected graph G with r-tomic separators of size $\leqslant k$ such that $(k-1)r < \Delta_r(G)$. The r-local covering p_r restricts to a bijection $R_{\hat{N}}(\hat{s}) \to R_N(s)$ for every $\hat{s} \in \mathrm{ori}(\hat{N})$, where $s = p_r(\hat{s})$.*

Proof. Let $\hat{s} \in \mathrm{ori}(\hat{N})$ be given and let $s := p_r(\hat{s})$. Let $\hat{S}$ be the set of all $\hat{t} \in \mathrm{ori}(\hat{N})$ such that $X(\hat{t})$ meets $X(\hat{s})$ and $\hat{s} \succ \hat{t}$. Define $S \subseteq \mathrm{ori}(N)$ analogously, with N in place of $\hat{N}$ and s in place of $\hat{s}$. So $\hat{S}$ and S contain precisely the local separations considered in the definition of $R_N(\hat{s})$ and $R_N(s)$. By Lift 2.6.6 and Lift 2.8.10, the r-local covering map p_r induces a bijection $\hat{S} \to S$, and for all $\hat{t} \in \hat{S} \cup \{\hat{s}\}$ it restricts to a bijection $R(\hat{t}) \to R(t)$. Hence, p_r restricts to a bijection $R_{\hat{N}}(\hat{s}) \to R_N(s)$. □

Projection 2.8.14 (Σ^*_N)**.*** *Let N be a nested set of r-local separations of a connected graph G with r-tomic separators of size $\leqslant k$ such that $(k-1)r < \Delta_r(G)$. Consider the set $\hat{N}$ of all lifts of all $s \in N$.*

(i) *$\hat{s}\Sigma^*_{\hat{N}}\hat{t} \Longrightarrow p_r(\hat{s})\Sigma^*_N p_r(\hat{t})$ for all $\hat{s}, \hat{t} \in \mathrm{ori}(\hat{N})$ and $* \in \{\vec{x}, \mathscr{H}, \mathcal{X}\}$.*

(ii) *For every cutout $\hat{\sigma}$ of $\hat{N}$, there exists a unique cutout σ of N such that $p_r(\hat{\sigma}) \subseteq \sigma$. Moreover, $p_r(\mathrm{part}(\hat{\sigma})) \subseteq \mathrm{part}(\sigma)$.*

Proof. (i). Let $\hat{s}, \hat{t} \in \mathrm{ori}(\hat{N})$ as in the statement and abbreviate $s := p_r(\hat{s})$ and $t := p_r(\hat{t})$. For $* = \vec{x}$ we are done by Lift 2.6.6 and Lift 2.8.10. For $* = \mathcal{X}$ this follows from Projection 2.8.13 and $\bigcup \mathcal{X}(\hat{N}) = p_r^{-1}(\bigcup \mathcal{X}(N))$.

For $* = \mathscr{H}$ we find a component $\hat{K}$ of $G_r - \bigcup \mathcal{X}(\hat{N})$ such that $\partial\hat{K}$ meets $R_{\hat{N}}(\hat{s})$ and $R_{\hat{N}}(\hat{t})$. The projection $K := p_r(\hat{K})$ is a component of $G - \bigcup \mathcal{X}(N)$. Then both $R_N(s)$ and $R_N(t)$ meet ∂K by Projection 2.8.13, giving $s\Sigma^{\mathscr{H}}_N t$.

(ii). The existence of σ given $\hat{\sigma}$ is immediate from (i). Every component $\hat{K} \in \mathscr{H}(\hat{N})$ projects to a component $p_r(\hat{K}) \in \mathscr{H}(N)$ as observed in the proof of (i). Hence, the 'moreover'-part follows with Projection 2.8.13. □

Lift 2.8.15 (Σ^*_N)**.*** *Let N be a nested set of r-local separations of a connected graph G with r-tomic separators of size $\leqslant k$ such that $(k-1)r < \Delta_r(G)$. Consider the set $\hat{N}$ of all lifts of all $s \in N$.*

(i) *Let $s, t \in \mathrm{ori}(N)$ and $* \in \{\vec{x}, K, \mathcal{X}\}$ such that $s\Sigma^*_N t$. For every lift $\hat{s}$ of s, there exists a lift $\hat{t}$ of t such that $\hat{s}\Sigma^*_{\hat{N}}\hat{t}$.*

(ii) *For every cutout σ of N and every lift $\hat{s}$ of some $s \in \sigma$, the cutout $\hat{\sigma}$ of $\hat{N}$ with $\hat{s} \in \hat{\sigma}$ satisfies $\sigma \subseteq p_r(\hat{\sigma})$. Moreover, $\mathrm{part}(\sigma) \subseteq p_r(\mathrm{part}(\hat{\sigma}))$.*

Proof. (i). Let s, t be as in the statement of (i), and let $\hat{s}$ be an arbitrary lift of s. For $* = \vec{x}$ we obtain a unique lift $\hat{t}$ of t such that $X(\hat{t})$ meets $X(\hat{s})$, by Lift 2.6.6. Then $\hat{s}\Sigma^{\vec{x}}_{\hat{N}}\hat{t}$ by Lift 2.6.6 and Lift 2.8.10.

For $* = \mathscr{K}$ we find $K \in \mathscr{K}(N)$ such that both $R_N(s)$ and $R_N(t)$ meet ∂K. Since p_r restricts to a bijection $R_{\hat{N}}(\hat{s}) \to R_N(s)$ by Projection 2.8.13, there is a component $\hat{K}$ of $p_r^{-1}(K)$ such that $R_{\hat{N}}(\hat{s})$ meets $\partial\hat{K}$. As K is connected, p_r restricts to a surjection $\partial\hat{K} \to \partial K$. Choose $e \in \partial K \cap R_N(t)$ and let $\hat{e} \in \partial\hat{K}$ be a lift of e. Let $\hat{t}$ be a lift of t with $\hat{e} \in \partial X(\hat{t})$. Then $\hat{e} \in R_{\hat{N}}(\hat{t})$ by Projection 2.8.13, so $\hat{K}$ witnesses that $\hat{s}\Sigma^{\mathscr{K}}_{\hat{N}}\hat{t}$.

For $* = \mathcal{X}$ we find an $\mathcal{X}(N)$-crossing edge e of G which is shared by $R_N(s)$ and $R_N(t)$. There is a lift $\hat{e}$ of e in $R_{\hat{N}}(\hat{s})$ by Projection 2.8.13. Note that $\hat{e}$ crosses $\mathcal{X}(\hat{N})$. Let $\hat{t}$ be a lift of t with $\hat{e} \in \partial X(\hat{t})$. Then $\hat{e} \in R_{\hat{N}}(\hat{t})$ by Projection 2.8.13. Hence, $\hat{s}\Sigma^{\mathcal{X}}_{\hat{N}}\hat{t}$ as desired.

(ii). Let $\sigma, s, \hat{\sigma}, \hat{s}$ be as in the statement of (ii). For every $t \in \sigma$, there are $n \in \mathbb{N}$ and $s_1, \dots, s_{n-1} \in \sigma$ such that for every $i \in [n]$ there is some $*(i) \in \{\vec{x}, \mathscr{K}, \mathcal{X}\}$ with $s_{i-1}\Sigma_N^{*(i)}s_i$, where $s_0 := s$ and $s_n := t$. Set $\hat{s}_0 := \hat{s}$. Iteratively applying (i) for $i = 1, \dots, n$ yields a lift $\hat{s}_i$ of s_i such that $\hat{s}_{i-1}\Sigma_N^{*(i)}\hat{s}_i$. Thus, $\hat{t} = \hat{s}_n \in \hat{\sigma}$, as desired. The 'moreover'-part follows with Projection 2.8.13 and the following observation from the proof of (i): For every component $K \in \mathscr{K}(N)$ and every component $\hat{K}$ of $p_r^{-1}(K)$, we have $\hat{K} \in \mathscr{K}(\hat{N})$ and p_r restricts to a surjection $\partial\hat{K} \to K$. □

Proposition 2.8.16 (Cutout)**.*** *Let N be a nested set of r-local separations of a connected graph G with r-tomic separators of size $\leqslant k$ such that $(k-1)r < \Delta_r(G)$. Consider the set $\hat{N}$ of all lifts of all $s \in N$. Then $\hat{\sigma} \mapsto p_r(\hat{\sigma})$ is a well-defined surjection from the cutouts of $\hat{N}$ to the cutouts of N. Moreover, $p_r(\mathrm{part}(\hat{\sigma})) = \mathrm{part}(p_r(\hat{\sigma}))$.*

Proof. Well-defined. Let $\hat{\sigma}$ be a cutout of $\hat{N}$. Then $p_r(\hat{\sigma}) \subseteq \sigma$ for a unique cutout σ of N by Projection 2.8.14 (ii). Pick $\hat{s} \in \hat{\sigma}$ arbitrarily. Then there exists a cutout $\hat{\tau}$ of $\hat{N}$ with $\hat{s} \in \hat{\tau}$ and $\sigma \subseteq p_r(\hat{\tau})$ by Lift 2.8.15 (ii). As $\hat{\sigma}$ and $\hat{\tau}$ are both cutouts of $\hat{N}$ and share $\hat{s}$, we get $\hat{\sigma} = \hat{\tau}$. So $p_r(\hat{\sigma}) \subseteq \sigma \subseteq p_r(\hat{\tau})$ yields $\sigma = p_r(\hat{\sigma})$. Hence, p_r defines a map from the cutouts of $\hat{N}$ to the cutouts of N.

Surjective. Let σ be a cutout of N. There exists a cutout $\hat{\sigma}$ of $\hat{N}$ with $\sigma \subseteq p_r(\hat{\sigma})$ by Lift 2.8.15 (ii). The inclusion is an equality as we have showed above that $p_r(\hat{\sigma})$ is a cutout of N.

Parts. We have '$\subseteq$' by Projection 2.8.14 (ii) and '$\supseteq$' by Lift 2.8.15 (ii). □

2.8.4 From nested local separations to graph-decompositions

Construction 2.8.17 (Graph-decomposition)**.** Let N be a nested set of r-local separations of G. We define the *candidate* $\mathcal{H}(N) = (H, \mathcal{G})$ for a graph-decomposition, as follows. If $N = \emptyset$, let $(H, \mathcal{G})$ be the trivial graph-decomposition whose decomposition graph H is the graph on a single vertex h and $G_h := G$. Assume now that N is non-empty. The nodes of H are the cutouts of N and the edges of H are the r-local separations in N, where a separation $\{E_1, X, E_2\} \in N$ connects the two (possibly equal) cutouts σ_1, σ_2 of N with $(E_{3-i}, X, E_i) \in \sigma_i$. We define the part G_h of a node h of H to be the part of h, with h viewed as a cutout of N.

We remark that this construction of $\mathcal{H}(N)$ is canonical; in particular, if N is canonical, so is $\mathcal{H}(N)$. Additional assumptions on N are required to ensure that $\mathcal{H}(N)$ is a graph-decomposition.

Correspondence 2.8.18 (Graph-decomposition)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Let N be a nested set of proper separations of G with r-tomic separators such that $\mathcal{T}(N) =: (T, \mathcal{V})$ is a tree-decomposition. Let $(H, \mathcal{G}) := \mathcal{H}(N_r)$. Then the map $s \mapsto s_r$ induces a canonical isomorphism $\psi\colon T \to H$ such that $V_\sigma = G_{\psi(\sigma)}$ for all nodes σ of T.*

Proof. Let σ be a node of T. Recall that σ is a splitting star of N by the definition of $\mathcal{T}(N)$. The splitting stars of N are precisely the cutouts of N by Proposition 2.8.7, and $s \mapsto s_r$ defines a bijection between the cutouts of N and the cutouts of N_r by Correspondence 2.8.12 (ii). So $\sigma \mapsto \sigma_r = \{s_r : s \in \sigma\}$ is a bijection $V(T) \to V(H)$ that respects the incidences of the edges for which $s \mapsto s_r$ is a bijection $E(T) \to E(H)$. We have $V_\sigma = \mathrm{part}(\sigma)$ by Proposition 2.8.8, and we have $\mathrm{part}(\sigma) = \mathrm{part}(\sigma_r)$ by Correspondence 2.8.12 (iii). □

To apply Correspondence 2.8.18, we need a sufficient condition for when a nested set of proper separations of G with r-tomic separators defines a tree-decomposition.

Lemma 2.8.19. *Let G be a locally finite and connected graph, let $k \in \mathbb{N}$, and let N be a nested set of proper separations of G with r-tomic separators of size $\leqslant k$.*

(i) *For every vertex v of G_r, there are only finitely many separations in N whose separators contain v.*

(ii) *$\mathcal{T}(N)$ is a tree-decomposition of G_r.*

Proof. (i). Let N_v consist of all $s \in N$ with $v \in X(s)$. Let $U := B_G(v, k(r+1))$. Note that U is finite. Since the separators of the separations in N are r-tomic, we have $X(s) \cup N_G(X(s)) \subseteq U$ for all $s \in N_v$. Hence, the map $\{A, B\} \mapsto \{A \cap V(U), B \cap V(U)\}$ that restricts each $\{A, B\} \in N_v$ to a proper separation of U is injective. Since U is finite, there are only finitely many separations of U, and so N_v is finite.

(ii). By [EKT22, Lemma 2.7] it suffices to show that N exhibits the following property: For every strictly decreasing sequence $(A_0, B_0) > (A_1, B_1) > \cdots$ of oriented separations in $\mathrm{ori}(N)$, we have $\bigcap_{i \in \mathbb{N}} B_i = \emptyset$. To show that N exhibits this property, we exactly follow the proof of [DJKK22, Lemma 7.5] but use (i) instead of [DJKK22, Lemma 7.4]. □

Lemma 2.8.20. *Let G be a locally finite and connected graph. Let $\hat{N}$ be a nested set of r-local separations of G_r with r-tomic separators all of size $\leqslant k$ for some $k \in \mathbb{N}$. Then $\mathcal{H}(\hat{N})$ is a tree-decomposition of G_r.*

Proof. By Correspondence 2.4.6, there exists a unique set $\hat{M}$ of separations of G_r with r-tomic separators such that $(\hat{M})_r = \hat{N}$. Since all separations in $\hat{M}$ have r-tomic separators of size $\leqslant k$, we have that $\mathcal{T}(\hat{M})$ is a tree-decomposition by Lemma 2.8.19. Hence, $\mathcal{H}(\hat{N})$ is a tree-decomposition by Correspondence 2.8.18. □

Recall that an isomorphism $\psi\colon G \to G'$ between multigraphs is a pair $\psi = (\psi_V, \psi_E)$ of bijections $\psi_V\colon V(G) \to V(G')$ and $\psi_E\colon E(G) \to E(G')$ that respect incidences of edges with vertices in both directions.

Theorem 2.8.21.* *Let G be a locally finite and connected graph. Let N be a nested set of r-local separations of G with r-tomic separators of size $\leqslant k$ such that $(k-1)r < \Delta_r(G)$. Let $(H, \mathcal{G}) := \mathcal{H}(N)$. Let $\hat{N}$ be the set of all lifts of all $s \in N$. Let $(\bar{H}, \bar{\mathcal{G}})$ be the graph-decomposition of G that is defined by the tree-decomposition $\mathcal{H}(\hat{N})$ via p_r; see Construction 2.2.6. For each cutout σ of N, let $\tau(\sigma)$ denote the set of all cutouts $\hat{\sigma}$ of $\hat{N}$ with $p_r(\hat{\sigma}) = \sigma$. For each $s \in N$, let $L(s)$ denote the set of all lifts of s. Then*

(i) *the pair of bijections $(\sigma \mapsto \tau(\sigma), s \mapsto L(s))$ is an isomorphism $H \to \bar{H}$, and*

(ii) *$G_\sigma = \bar{G}_{\tau(\sigma)}$.*

All in all, $\mathcal{H}(N)$ and $\bar{\mathcal{H}}(N)$ are canonically isomorphic.

In the context of Theorem 2.8.21, we recall that the lifts of the $s \in N$ exist by Lift 2.5.9 and Lemma 2.6.5, and $\mathcal{H}(\hat{N})$ is a tree-decomposition of G_r by Lemma 2.8.20.

Proof. (i). We first show that the map $\sigma \mapsto \tau(\sigma)$ is a bijection $V(H) \to V(\bar{H})$. Write $(\hat{H}, \hat{\mathcal{G}}) := \mathcal{H}(\hat{N})$. The vertices of $\bar{H}$ are the orbits of $V(\hat{H})$ under the action of the group $\Gamma_r := \Gamma_r(G)$ of deck transformations of p_r. The vertices of $V(\hat{H})$, in turn, are the cutouts of $\hat{N}$. The group Γ_r naturally acts on $\hat{N}$. Recall that p_r induces a well-defined surjection from the cutouts of $\hat{N}$ to the cutouts of N by Proposition 2.8.16. Since Γ_r acts transitively on the fibres of p_r, two cutouts of $\hat{N}$ lie in the same orbit if and only if they project to the same cutout of N. Hence, $V(\bar{H}) = \{\tau(\sigma) : \sigma \in V(H)\}$. So $\sigma \mapsto \tau(\sigma)$ is a surjection $V(H) \to V(\bar{H})$, and it is injective by definition.

Similarly, $E(\bar{H})$ consists of the Γ_r-orbits of $E(\hat{H})$, and $E(\hat{H}) = \hat{N}$. As the Γ_r-orbits of $\hat{N}$ are the sets $L(s)$ for $s \in N$, it follows that $s \mapsto L(s)$ is a bijection $E(H) \to E(\bar{H})$.

Incidences of edges and vertices are clearly preserved.

(ii). Given σ, we have to show that $G_\sigma = \bar{G}_{\tau(\sigma)}$. By definition we have $\bar{G}_{\tau(\sigma)} = p_r(\hat{G}_{\hat{\sigma}})$ for some (equivalently: every) cutout $\hat{\sigma} \in \tau(\sigma)$. We have $\hat{G}_{\hat{\sigma}} = \mathrm{part}(\hat{\sigma})$ and $G_\sigma = \mathrm{part}(\sigma)$ by definition. So it suffices to show $\mathrm{part}(\sigma) = p_r(\mathrm{part}(\hat{\sigma}))$. Indeed, this holds by Proposition 2.8.16. □

Corollary 2.8.22.* *Let G be a locally finite and connected graph. Let N be a nested set of r-local separations of G with r-tomic separators of size $\leqslant k$ such that $(k-1)r < \Delta_r(G)$. Then $\mathcal{H}(N)$ is a graph-decomposition of G.* □

Let $\mathcal{H} = (H, \mathcal{G})$ be a graph-decomposition of some graph G, and let F be a set of edges of H. We define $\mathcal{H}/F := (\tilde{H}, \tilde{\mathcal{G}})$ where $\tilde{H}$ is obtained from H by contracting the edges in F (keeping loops and parallel edges), and $\tilde{\mathcal{G}}$ contains for every vertex $\tilde{h}$ of $\tilde{H}$ the part $\tilde{G}_{\tilde{h}}$ that is the union of all parts G_h with h contained in the branch-set of $\tilde{h}$. We say that $\mathcal{H}$ *refines* $\mathcal{H}/F$. If $\mathcal{H} = \mathcal{H}(N)$ for a set of r-local separations and $N' \subseteq N$, then $\mathcal{H}/N'$ is well-defined as N' formally corresponds to a set of edges of $\mathcal{H}$.

Lemma 2.8.23.* *Let G be a locally finite and connected graph. Let $N \subseteq M$ be nested sets of r-local separations of G with r-tomic separators of size $\leqslant k$ such that $(k-1)r < \Delta_r(G)$. Then $\mathcal{H}(N) = \mathcal{H}(M)/(M \setminus N)$.*

Proof. Let $\hat{N}$ be the set of all lifts to G_r of the elements of N, which is possible by Lift 2.5.9 and Lemma 2.6.5. We define $\hat{M}$ analogously. Then $N \subseteq M$ implies $\hat{N} \subseteq \hat{M}$. By Lemma 2.8.20, both $\mathcal{H}(\hat{N})$ and $\mathcal{H}(\hat{M})$ are tree-decompositions of G_r. Hence, $\mathcal{H}(\hat{N}) = \mathcal{H}(\hat{M})/(\hat{M} \smallsetminus \hat{N})$ by standard arguments; see e.g. [CDHS14]. By Theorem 2.8.21, $\mathcal{H}(N)$ is defined by $\mathcal{H}(\hat{N})$ via p_r, and $\mathcal{H}(M)$ is defined by $\mathcal{H}(\hat{M})$ via p_r, giving $\mathcal{H}(N) = \mathcal{H}(M)/(M \smallsetminus N)$. □

2.9 Proof of Theorems 4 and 5

Recall that $N_r^{\leqslant k}(G)$ was defined in Construction 2.7.16. We define $\mathcal{H}_r^{\leqslant k}(G) := \mathcal{H}(N_r^{\leqslant k}(G))$, with the operator $\mathcal{H}$ from Construction 2.8.17. Recall that the edges of the decomposition graph of $\mathcal{H}_r^{\leqslant k}(G)$ are the r-local separations in $N_r^{\leqslant k}(G)$. We say that these edges *represent* all the r-local bottlenecks in G of order $\leqslant k$ if every r-local bottleneck in G of order $\leqslant k$ contains such an edge (viewed as the r-local separation that it formally is).

Theorem 4. *Let G be a finite connected graph and $r, k \in \mathbb{N}$ with $k < K(G, r)$.*

(i) *$\mathcal{H}_r^{\leqslant k}(G)$ is a canonical graph-decomposition of G that represents all the r-local bottlenecks in G of order $\leqslant k$.*

(ii) *$\mathcal{H}_r^{\leqslant k}(G_r)$ is a tree-decomposition induced by a tree-of-tangles of resolution k of the r-local covering G_r.*

(iii) *$\mathcal{H}_r^{\leqslant k}(G)$ is canonically isomorphic to the graph-decomposition of G that is obtained from the tree-decomposition $\mathcal{H}_r^{\leqslant k}(G_r)$ of the r-local covering G_r by folding via p_r.*

(iv) *$\mathcal{H}_r^{\leqslant k}(G)$ refines $\mathcal{H}_r^{\leqslant \ell}(G)$ for all $\ell \in \mathbb{N}$ with $\ell \leqslant k$.*

Proof. Let G be a finite connected graph and $r \in \mathbb{N}$ an integer. Let $k \in \mathbb{N}$ with $k < K(G, r)$. So $(k-1)r < \Delta_r(G)$.

We claim that we may assume without loss of generality that * is satisfied. To see this, suppose that $r \leqslant 1$. Then all tight r-local separators of G have size 1 by Lemma 2.5.5, and also all tight separators of G_r have size 1, as in this case G_r is the universal cover of G which is a tree. Thus, all r-local bottlenecks in G and all bottlenecks in G_r have order 1, and we may assume without loss of generality that $k = 1$. Hence, * is satisfied in any case.

(ii). The r-local separations in $N_r^{\leqslant k}(G_r)$ stem from r-local bottlenecks. So they are tight and, in particular, r-tomic by Lemma 2.4.3. Hence, $\mathcal{H}_r^{\leqslant k}(G_r)$ is a tree-decomposition by Lemma 2.8.20. Every pair of $(\leqslant k)$-distinguishable tangles in G_r induces a bottleneck in G_r by Example 2.3.4. The map $s \mapsto s_r$ induces a bijection between the bottlenecks in G_r and the r-local bottlenecks in G_r, by Correspondence 2.7.13. Hence, $\mathcal{H}_r^{\leqslant k}(G_r)$ is a tree-of-tangles of G_r of resolution k.

(iii). Recall that $\mathcal{H}_r^{\leqslant k}(G_r)$ is indeed a tree-of-tangles by (ii). Let $N := N_r^{\leqslant k}(G)$. The set $\hat{N}$ of all lifts of all $s \in N$ is equal to $N_r^{\leqslant k}(G_r)$ by Theorem 2.7.18. Hence, $\mathcal{H}(N)$ is canonically isomorphic to the graph-decomposition of G obtained from $\mathcal{H}(\hat{N})$ by folding via p_r, by Theorem 2.8.21. Since $\mathcal{H}_r^{\leqslant k}(G) = \mathcal{H}(N)$ and $\mathcal{H}(\hat{N}) = \mathcal{H}_r^{\leqslant k}(G_r)$, this yields (iii).

(i). We have already shown in (iii) that $\mathcal{H}_r^{\leqslant k}(G)$ is a graph-decomposition of G. It is canonical because N is canonical. So it remains to show that the edges of H represent all the r-local bottlenecks in G of order $\leqslant k$. Recall that $E(H) = N$. And indeed, every r-local bottleneck β_r in G of order $\leqslant k$ shares an r-local separation with N by Theorem 2.7.19.

(iv). Since $N_r^{\leqslant \ell}(G) \subseteq N_r^{\leqslant k}(G)$ by definition, this follows from Lemma 2.8.23. □

Proof of Theorem 5. Given G, r, k as in Theorem 4, the algorithm first performs Construction 2.7.16 to determine $N_r^{\leqslant k}(G)$, and in the second step applies Construction 2.8.17 to $N_r^{\leqslant k}(G)$ to determine $\mathcal{H}_r^{\leqslant k}(G)$. □

CHAPTER 3

Locally chordal graphs

This chapter is based on [AJKK25], which is joint work with Tara Abrishami, Raphael W. Jacobs and Jonas Kobler.

3.1 Introduction

Consider a graph G which is obtained by gluing cliques in the shape of a high-girth graph H in such a way that, locally at each vertex, the graph G looks as if it were composed of cliques glued together along a tree. An illustration is given in Figure 3.1. Graphs that are composed of cliques glued together along a tree are the well-known *chordal*[1] graphs. The global composition of the graph G along H described above ensures that, G still exhibits such tree-shape locally: that G is locally chordal.

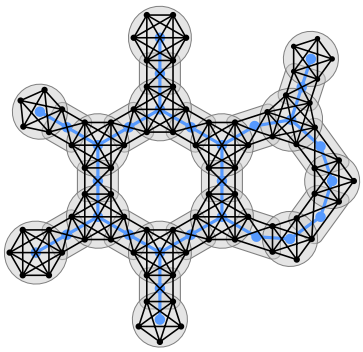

Figure 3.1: A locally chordal graph G, depicted in black, obtained by gluing cliques in the shape of H, depicted in blue.

In this context, one of the greatest hopes is that the converse from local to global also holds: Does a locally chordal graph G admit a global decomposition which simultaneously witnesses that G is locally chordal everywhere? Although locally at each vertex v the graph G, by assumption, consists of cliques glued along a tree T^v, it is generally not at all obvious that the local descriptions over the trees T^v agree in such a way that they can be combined into a common description over a common graph H. What is surprising, however, is that our main result confirms the local-to-global hope in the affirmative. Furthermore, our results are the first to make use of a new general theory designed to facilitate local-to-global analysis by constructing global decompositions from the local structure [DJKK22]. Our results therefore provide a possible roadmap for using this framework to study local-to-global structures more generally.

[1] Chordal graphs are also referred to as *rigid circuit* [Dir61] or *triangulated* [Ros70] graphs.

To state our results more precisely, we give some definitions. A graph is *chordal* if each of its cycles of length at least four has a chord. For an integer $r \geqslant 0$, the given *degree of locality*, a graph is *r-chordal* if it each of its cycles of length at least four but at most r has a chord, and a graph is *r-locally chordal* if each of its balls of radius $r/2$ is chordal.[2] Our main result not only shows that r-locally chordal graphs G admit a global decomposition which simultaneously witnesses that G is r-locally chordal, but indeed admit many different characterizations in terms of forbidden induced subgraphs, separators and intersection graphs.

Characterization. *Let G be a locally finite graph and $r \geqslant 3$ an integer. The following are equivalent to G being r-locally chordal:*

(i) *G is r-chordal and wheel-free (Theorem 6).*

(ii) *Every minimal r-local separator of G is a clique of G (Theorem 14).*

(iii) *G is an r-acyclic region intersection graph (Theorem 9).*

(iv) *G admits a graph-decomposition into cliques which trivially induces r-local separations of G (Theorem 12).*

Next, we describe in more detail our characterizations of locally chordal graphs and also their relation to chordal graphs.

3.1.1 When the degree of locality goes to infinity

Since trees are chordal and every ball of radius $r/2$ for $r \leqslant 2$ is a tree (specifically, a star), every graph is 0-, 1-, and 2-locally chordal. As we shall see in a moment, every graph which is r-locally chordal is r'-locally chordal for all $0 \leqslant r' \leqslant r$. Moreover, the graphs which are r-locally chordal for every $r \in \mathbb{N}$ are exactly the chordal graphs. Therefore, the sequence of r-locally chordal graphs for $r \in \mathbb{N}$ is an inclusion-wise decreasing sequence of graph classes, starting from the class of all graphs and ending with the class of chordal graphs. When r goes to infinity, our characterization of locally chordal graphs indeed recovers the well-known characterizations of chordal graphs [Dir61; Gav74; HL18].

Well-known characterizations of chordal graphs. *Let G be a locally finite graph. The following are equivalent:*

(1) G has no induced cycle of length at least four.

(2) Every minimal separator of G is a clique of G.

(3) G is a subtree intersection graph.

[2] For a vertex v of a graph G, we define the *ball of radius $r/2$ around v*, denoted $B_G(v, r/2)$ or, for short, $B_{r/2}(v)$, as the subgraph of G given by the set of all vertices of distance at most $r/2$ from v and the set of all edges of G for which the sum of the endpoints' distances to v is strictly less than r, see Section 3.2.2 for a comparison to the *depth-d closed neighborhood*.

(4) G admits a tree-decomposition into cliques.

The graph G is *chordal* if it satisfies these conditions.

Chordal graphs are widely studied and have many theoretical and algorithmic applications. One reason for this is that chordal graphs have so many different characterizations. A disadvantage of this strong structure is that only few graphs are chordal. Although locally chordal graphs are a much larger class of graphs, for a given $r \geqslant 0$ they range from all graphs to wheel-free graphs to chordal graphs, our characterization ensures that these classes still retain many of the properties that make chordal graphs so nice.

3.1.2 Characterization via induced subgraphs

At first glance it might seem that r-locally chordal is equivalent to r-chordal. But a wheel with a rim of length at least 4 is not chordal but contained in the 3/2-ball around its hub, and thus is not r-locally chordal for $r \geqslant 3$. Formally, a *wheel* W_n for $n \geqslant 4$ consists of a cycle of length n and an additional vertex complete to the cycle. A graph is *wheel-free* if it does not contain a wheel W_n for $n \geqslant 4$ as an induced subgraph. Surprisingly, Characterization (i) states that induced wheels are the only obstruction for an r-chordal graph to r-locally chordal.

When r goes to infinity, r-chordal yields that every cycle of length at least four has a chord, so characterization (i) recovers (1).

3.1.3 Characterization via coverings

Recently, Diestel, Jacobs, Knappe and Kurkofka [DJKK22] introduced a new framework to discriminate between the local structure of a graph and its global structure. Given a (finite) graph G and a parameter $r \geqslant 0$, they construct a (possibly infinite) graph G_r, called the *r-local cover of G*, that represents only the local structure of G. The construction of G_r is based on the idea that short cycles, of length at most r, are local, and that the "local structure" of a graph consists of exactly the structure generated by its local cycles. This is made precise using the notion of covering maps from topology. The graph G_r is equipped with a surjective homomorphism $p_r \colon G_r \mapsto G$, its covering map, which satisfies the property that it *preserves $r/2$-balls*, i.e. $B_{G_r}(\hat{v}, r/2))$ is isomorphic to $B_G(v, r/2)$ for every vertex v of $\hat{G}$ and $v = p(\hat{v})$ of G. The r-local cover G_r can be characterised as the cover which is universal among those that preserve $r/2$-balls. See Section 3.2.3 for a precise definition of the local cover.

In the local cover perspective, if the r-local cover G_r of a graph G is chordal, then the r-local structure of G is chordal. Since p_r preserves $r/2$-balls, every $r/2$-ball of G appears as an $r/2$-ball of G_r (possibly more than once). Since $r/2$-balls of chordal graphs are chordal (Lemma 3.2.2), every graph whose r-local cover G_r is chordal is r-locally chordal. It turns out that the converse is also true: for every $r \in \mathbb{N}$, a graph G is r-locally chordal if and only if G_r is chordal (Theorem 6).

The fact that r-locally chordal graphs are precisely the graphs whose r-local covers are chordal allows us to make use of local covers to study locally chordal graphs. Local covers

are used to prove several of our results: specifically, Characterization (i), i.e. r-locally chordal graphs are precisely those graphs that are r-chordal and wheel-free (Theorem 6) and to prove our Characterization (iii) and (iv) not only for finite but also locally finite graphs. Moreover, the framework introduced in [CJKK25; DJKK22] also yields for every integer $r \geqslant 0$ a canonical graph-decomposition $\mathcal{H}_r(G)$ of G (see Section 3.14, and Section 3.13, for a recap). We prove that the r-locally chordal graphs are precisely those for which the decomposition $\mathcal{H}_r(G)$ is into cliques (Theorem 19). Moreover, $\mathcal{H}_r(G)$ is used in the construction of (iii) and (iv) (Proof of Theorem 10).

3.1.4 Characterization via local separators

Local components at an arbitrarily sized vertex-subset X of a graph G were introduced in [CJKK25] as a local analog to *componental cuts*, the sets of edges between a component of $G - X$ and X. Based on this notion, they introduced *(tight) r-local separators*, a local analog of separators. When r goes to infinity, these become the usual (tight) separators. Unfortunately, tight local separators of locally chordal graphs are not in general cliques (Example 3.9.1). This is due to the fact that, in contrast to the usual separators, tight local separators are not necessarily minimal local separators, which we introduce as a local analogue to minimal separators in this chapter. When r goes to infinity, minimal r-local separators become the usual minimal (equivalently: tight) separators. Hence, when r goes to infinity, Characterization (ii) yields Dirac's characterization of chordal graphs as those graphs whose minimal separators are cliques [Dir61, Theorem 1].

3.1.5 Characterization via region representations

Let us next discuss Characterization of (i) by (iii). A *region* of a graph H is a connected subgraph of H. An map $v \mapsto H_v$ assigning to the vertices v of a graph G regions H_v of a graph H *represents* G if an edge uv is present in G if and only if H_u and H_v intersect. A region representation $v \mapsto H_v$ is *r-acyclic* if for every set $X \subseteq V(G)$ of at most r vertices the union $\bigcup_{x \in X} H_x$ is an acyclic subgraph of H. If G has such a representation, we then call it an *r-acyclic region intersection graph*.

Every graph is a 2-acyclic region intersection graph; see Example 3.4.4 for a proof. But it is already interesting to ask which graphs are 3-acyclic region intersection graphs. Wheels, and all graphs containing wheels as induced subgraphs, are not 3-acyclic region intersection graphs (Lemma 3.5.5), so 3-acyclic region intersection graphs must exclude wheels. In fact, Gavril [Gav96] characterized the graphs which are 3-acyclic region intersection graphs as precisely the wheel-free graphs (see Section 3.4.2). Wheel-free graphs are also precisely the 3-locally chordal graphs. At the other extreme, ∞-acyclic region intersection graphs are precisely the intersection graphs of subtrees of forest, which are well known to be exactly the chordal graphs. Characterization (iii) now fills in the spectrum for all other values of r, giving a characterization of r-acyclic region intersection graphs for all $r \in \mathbb{N}$, as exactly the r-locally chordal graphs.

3.1.6 Characterization via graph-decompositions

Let us finally turn to Characterization (iv). A *graph-decomposition* of a graph G, as introduced in [DJKK22], consists of a decomposition graph H and an assignment $h \mapsto V_h$ of a vertex-subset V_h of G to each node h of H satisfying conditions that formalize the idea that the graph-decomposition arranges the graph G in the shape of H. Specifically: (H1) the union of the subgraphs $G[V_h]$ is G, and (H2) for each vertex v of G the subgraph of H induced by the nodes h for which V_h contains v is connected (Section 3.3.1). A graph-decomposition is one *into cliques* if every set V_h is a clique of G. It turns out that graph-decompositions of a graph G into cliques correspond to its region representations (see Section 3.3 for details). More specifically, a graph-decomposition $(H, \mathcal{V})$ into cliques induces a region representation by mapping each vertex v of G to a spanning subgraph of the induced subgraph from (H2). Conversely, a region representation $v \mapsto H_v$ induces a graph-decomposition by assigning to each node h of H the vertex-subset $V_h \coloneqq \{v \in V(G) \mid h \in V(H_v)\}$. Notice that these two constructions are inverse to each other.

The r-acyclic region representations (equivalently: graph-decompositions into cliques), those from Characterization (iii), can be characterised as those graph-decompositions which "trivially induce" r-local separations (see Section 3.7 for details). *Local separations* were introduced in [CJKK25] simultaneously with local separators as a local analog to *separations*, which, when r goes to infinity, are precisely the usual separations. Hence, when r goes to infinity, Characterization (iv) gives that chordal graphs are exactly the graph that admit tree-decompositions into cliques.

For other values of r, we can ask what kinds of graph-decompositions $(H, \mathcal{V})$ we may choose for Characterization (iii) (equivalently: (iv)) for a given r-locally chordal graph G. For example, the r-acyclic condition in (iii) guarantees that their decomposition graphs H have girth greater than r (Lemma 3.8.1). This means that we can bound the number of edges of H by a subquadratic function in the vertices of H, by $|E(H)| \leqslant |V(H)|^{1+\frac{2}{r-2}}$ [Bol78, §III], so the decomposition graphs H are sparse. Moreover, the bound on the number of edges decreases with increasing r and approaches linear when r goes to infinity. Therefore, a qualitative way of describing Characterization (iv) is that the r-locally chordal graphs admit "well-structured" decompositions into cliques modeled on "r-sparse" graphs H. Furthermore, we may choose a graph-decomposition of G for (iv) *into the maximal cliques of* G, i.e. $h \mapsto V_h$ is a one-to-one correspondence between the nodes of H and the maximal cliques of G (Theorem 12 (ii)). There are at most $|V(G)|^2$ maximal cliques in a finite wheel-free (equivalently: 3-locally chordal) graph G (Section 3.12.1), which now bounds the number of nodes of the decomposition graph H. We conjecture that the number of maximal cliques in finite r-locally chordal graphs G approaches $O(|V(G)|)$ as r goes to infinity (Conjecture 3.12.1).

3.1.7 Further results: Locally deriving tree-decompositions

Recall that we started with the idea of gluing cliques together in the shape of a high-girth graph so that locally at each vertex the graph looks like being composed of cliques glued together along

a tree. Graph-decompositions from Characterization (iii) (equivalently: (iv)) formalize this idea. Indeed, let G be an r-locally chordal graph and $(H, \mathcal{V})$ be an r-acyclic graph-decomposition of G as in Characterization (iii). Now, roughly speaking, if we restrict the decomposition $(H, \mathcal{V})$ to just a ball of radius $r/2$ in G, we obtain a tree-decomposition into cliques. Essentially, this allows us to start from a collection of $|V(G)|$ tree-decompositions into cliques, one for each ball of radius $r/2$ of G, and find a single graph-decomposition of G which "puts them together" – in other words, which simultaneously witnesses each local tree-decomposition in a single structure. This property holds for every r-acyclic graph-decomposition of locally chordal graphs into cliques (Theorem 13).

3.1.8 Further results: Computing region representations

Given an r-locally chordal graph G on n vertices, we give an algorithm (Algorithm 16) to compute an r-acyclic region representation (Characterization (iii)). The algorithm does not depend on r. It is both efficient and conceptually simple: For every vertex v of G it runs Kruskal's algorithm to compute a maximum-weight spanning tree H^v of the intersection graph $\mathbf{K}(B_v)$ of the maximal cliques in $B_v := B_{3/2}(v) = G[N_G(v) \cup \{v\}]$. Then it sets H to be the union of the H^v, and returns the map $v \mapsto H^v$ as the desired r-acyclic region representation of G over H.

The (parallel) running time is at most $O(\Delta^2 \log \Delta)$ on $n = |V(G)|$ processors, where Δ is the maximum degree of G. And not only does the algorithm does not require r as an input (Theorem 15): given a wheel-free graph G, it returns an r-acyclic region representation for the maximum value of r such that G is r-locally chordal. In the extreme case that the graph G is chordal, the computed region representation is ∞-acyclic, thus over a forest H, recovering the familiar notion of *clique trees* (if G is additionally connected).

More generally, a graph-decomposition $(H, \mathcal{V})$ of a graph G into the maximal cliques of G corresponds to a spanning subgraph of the intersection graph $\mathbf{K}(G)$ of the maximal cliques in G by identifying the decomposition graph H with the obtained from the map $h \mapsto V_h$. Note that, for any fixed vertex v of G, the subgraph of H induced on the maximal cliques in G containing v is precisely H_v, and thus connected. Vice versa, any given spanning subgraph H of $\mathbf{K}(G)$ whose subgraphs H_v induced on the maximal cliques containing any fixed vertex v of G are connected yields a region representation of G over H by mapping a vertex v to H_v, which corresponds to a graph-decomposition of G into the maximal cliques of G. We call such H *clique graphs* of G, and it is *r-acylic*, if its corresponding region representation is. If H is a tree, this recovers the notion of *clique trees*. In the above algorithm the union H of the H^v indeed is a clique graph of G whose H_v are precisely the H^v.

Clique trees of a chordal graph G are known to be precisely the maximum-weight spanning trees of $\mathbf{K}(G)$, where we weight each edge by the size of the corresponding intersection of maximal cliques. We generalize this by proving that r-acyclic clique graphs of r-locally chordal graphs G are precisely the maximum-weight r-acyclic spanning subgraphs of $\mathbf{K}(G)$ (Theorem 17).

Since r-acyclic region representations can be efficiently computed, they can be used in algorithms. Clique trees are powerful algorithmic tools (see e.g. [WJ08]), and our algorithm to compute r-acyclic region representations of locally chordal graphs is an important step toward extending the algorithmic power of chordal graphs to locally chordal graphs. We outline a specific potential application, to the maximum independent set problem, in Section 3.12.2.

3.1.9 Further results: New characterizations of chordal graphs

In the course of studying locally chordal graphs, we also obtained several new characterizations of chordal graphs. First, we characterized chordal graphs as those graphs which are wheel-free and whose cycles are symmetric differences of triangles (Corollary 3.2.7). In fact, r-locally chordal graphs are precisely those wheel-free graphs whose cycles of length at most r are symmetric differences of triangles (Theorem 7). Second, every locally finite connected chordal graph G admits a representation as a subtree intersection graph which is canonical, i.e. which commutes with the automorphism of G (Theorem 8), which is a different kind of subtree representation of chordal graphs than the ones previously known (see Example 3.3.8). In fact, every locally finite r-locally chordal graph G admits a canonical r-acyclic region representation (Theorem 10). Third, we characterize chordal graphs as exactly those graphs G which are 3-acyclic region intersection graphs and such that every 3-acyclic region representation of G is in fact acyclic (Theorem 3.6.1). In fact, r-locally chordal graphs are precisely those graphs G which are 3-acyclic region intersection graphs and such that every 3-acyclic region representation of G is r-acyclic (Theorem 11).

3.1.10 An application: Locally interval graphs are circular-arc graphs

With Characterization (iii) in hand, we can ask which graphs have r-acyclic region representations over cycles. The graphs with 1-acyclic region representations over cycles form the well-known class of *circular-arc graphs*. By (iii), r-acyclic circular-arc graphs are a subclass of r-locally chordal graphs. Intuitively, zooming in around a vertex on a cycle yields a path, so we might guess that the chordal $r/2$-balls in this case have not only tree-decompositions but even path-decompositions into cliques. The subclass of chordal graphs which admit path-decompositions into cliques are the well-studied *interval graphs*. A correspondence between circular-arc graphs and locally interval graphs seems natural. Indeed, Rzążewski conjectured in private communication that locally interval graphs are precisely circular-arc graphs (Conjecture 3.11.3).

A graph is *r-locally interval* if each of its $r/2$-balls is interval. We prove that for $r \geqslant 4$, the r-locally interval graphs are exactly the r-acyclic circular-arc graphs (Theorem 18). The proof relies on the characterization of 3-acyclic circular arc graphs via forbidden induced subgraphs [CGS17] by Cao, Grippo and Safe. We also obtain a forbidden induced subgraph characterization of r-locally interval graphs for $r \geqslant 3$ (Proposition 3.11.5 and Theorem 18).

3.1.11 Notation & conventions

We follow standard graph theory notation and definitions from [Die25]. We also refer to [Hat02] for terminology on topology.

In this chapter we deal with simple graphs. Often, we will restrict our view to locally finite graphs; a graph is *locally finite* if it has countably many vertices and each vertex has finite degree. We specify when we deal with finite graphs, with locally finite, or with general (possibly infinite) graphs.

3.2 Basic characterizations of locally chordal graphs

In this section we prove the first set of characterizations of r-locally chordal graphs.

Theorem 6. *Let G be a (possibly infinite) graph and let $r \geqslant 3$ be an integer. The following are equivalent:*

(i) *G is r-locally chordal.*

(ii) *The r-local cover G_r of G is chordal*

(iii) *G is r-chordal and wheel-free.*

We remark that, for $r \leqslant 2$, every $r/2$-ball is chordal (and indeed a tree), so every graph is 0-, 1-, and 2-locally chordal. Also the r-local cover for $r \leqslant 2$ is a tree, and thus chordal, for every graph G. But G can contain some W_n with $n \geqslant 4$ as an induced subgraph, so the $r \geqslant 3$ condition is necessary in the statement of Theorem 6.

3.2.1 Background: Separators

To prove the main result of this section, we need to understand separators in chordal graphs. Let u, v be two vertices and X a vertex-subset of a graph G. Then X is an *u–v separator* of G, and *separates* u, v in G, if X meets every u–v path in G. If X separates some two vertices in G, then X is a *separator* of G. Moreover, X is a *tight* separator of G if there are at least two components C of $G - X$ are *tight at* X, i.e. $N_G(C) = X$.[3] We remark that X is a tight separator if and only if, for some two vertices u, v of G, the vertex-set X is inclusion-minimal among the u–v separators in G. Tight separators are also often called *minimal (vertex) separators*, e.g. [HL18].

Chordal graphs can be characterized by the structure of their tight separators:

Theorem 3.2.1 (Dirac [Dir61], Theorem 1)**.** *A (possibly infinite)[4] graph is chordal if and only if every minimal separator is a clique.*

[3] Tight components are also often called *full components*.

[4] It is immediate from the proof presented in [BP93, Theorem 2.1] that this also holds for infinite graphs.

3.2.2 Chordal graphs are locally chordal

Our first result of the section is the following.

Lemma 3.2.2. *Let G be a (possibly infinite) chordal graph and let $r \geqslant 3$ be an integer. If G is chordal, then G is r-locally chordal.*

We begin with some definitions. Let G be a graph. For a vertex v of G and an integer $d \geqslant 0$, the *depth-d closed neighborhood* $N_G^d[v]$ is the set of all vertices of G that have distance at most d to v, and the *depth-d (open) neighborhood* $N_G^d(v)$ is the set of all vertices of G that have distance precisely d to v. Note that, if $r = 2d+1$, then $B_G(v, r/2) = G[N^d[v]]$. Otherwise, if $r = 2d$, then $B_G(v, r/2)$ is obtained from $G[N^d[v]]$ by deleting the edges between vertices of distance precisely $r/2$ to v.

Let X and Y be two vertex-subsets of G. An *X–Y path* is a path such that only its first vertex is in X and only its last vertex is in Y. A path P in G is *through X* if all internal vertices of P are in X and P has at least one internal vertex.

Lemma 3.2.3. *Let $d = \lfloor r/2 \rfloor$ with $r \geqslant 0$ an integer. Let v be a vertex of a (possibly infinite) graph G such that $B_G(v, r/2)$ is chordal. Let u, w be two distinct vertices in $N^{d-1}(v)$. If there is a u–w path in $B_G(v, r/2)$ through $N^d(v)$, then uw is an edge of $B_G(v, r/2)$.*

Proof. For clarity in the proof, we remark that the assumption that $N^{d-1}(v)$ contains two distinct vertices u, w yields $d \geqslant 2$, i.e. $r \geqslant 4$. The assumption on u, w guarantees that there is a u–w path P through $N^d(v)$ in $B_G(v, r/2)$. Fix a vertex z of P in $N^d(v)$. Note that $N^{d-1}(v)$ is a v–z separator in G and thus also in $B_G(v, r/2)$. Let $S \subseteq N^{d-1}(v)$ be an inclusion-minimal v–z separator. The u–w path P through $N^d(v)$ together with $S \subseteq N^{d-1}(v)$ yields that $u, w \in S$. By Theorem 3.2.1, the minimal separator S of the chordal graph $B_G(v, r/2)$ is a clique; in particular, uw is an edge of $B_G(v, r/2)$. □

Now we prove Lemma 3.2.2.

Proof of Lemma 3.2.2. We want to show that every $r/2$-ball of G is chordal. If r is odd, then the $r/2$-balls in G are induced subgraphs and thus chordal, as G is chordal. So assume that r is even. Let O be a cycle of length at least 4 in an $r/2$-ball B around a vertex v in G. Since $r - 1$ is odd, the $(r-1)/2$-ball around v in G is chordal. Thus, O has a chord or contains a vertex z of distance $r/2 =: d$ to v. If O contains a chord, then we are done, so we may assume that O contains a vertex z of distance d to v. Let u, w be the two (distinct) neighbors of z on the cycle O. Since r is even, it follows that $N^d(v)$ is independent in $B_G(v, r/2)$, so every neighbor of z in $B_G(v, r/2)$ is in $N^{d-1}(v)$. Now, $u, w \in N^{d-1}(v)$ and uzw is a u–w path through $N^d(v)$, so by Lemma 3.2.3, $uw \in E(G)$. It follows that O has a chord. Since O and B were chosen arbitrarily, we conclude that G is r-locally chordal. □

3.2.3 Background: local cover & cliques

Our major tool to study the local structure of graphs involves coverings from topology. Let G be a graph. A *covering* of G is a surjective homomorphism $p\colon \hat{G} \to G$ such that p restricts to a bijection from the edges incident to any given vertex $\hat{v}$ of $\hat{G}$ to the edges incident to its projection $v \coloneqq p_r(\hat{v})$. In this chapter, we restrict our view to those coverings p of graphs G whose preimage $p^{-1}(C)$ of a component C of G is connected.

Let $r \in \mathbb{N}$. A homomorphism $p\colon \hat{G} \to G$ is *$r/2$-ball-preserving* if p restricts to an isomorphism from $B_{\hat{G}}(\hat{v}, r/2)$ to $B_G(v, r/2))$ for every vertex $\hat{v}$ of $\hat{G}$ and $v \coloneqq p_r(\hat{v})$. Thus, a surjective homomorphism $p\colon \hat{G} \to G$ is a covering of G if and only if p is 1-ball preserving. The *r-local covering* $p_r\colon G_r \to G$ of G is the universal $r/2$-ball-preserving covering of G, i.e. for every $r/2$-ball-preserving covering $p\colon \hat{G} \to G$ there exists a covering $q\colon G_r \to \hat{G}$ such that $p_r = p \circ q$. We refer to G_r as the *r-local cover* of G. Note that the 0-, 1- and 2-local covers of a connected, simple graph G are trees. For experts, we remark that the 0-, 1- and 2-local coverings are all isomorphic to the universal covering of G. A *deck transformation* of a covering $p\colon \hat{G} \to G$ is an automorphism γ of $\hat{G}$ that commutes with p, i.e. $p = p \circ \gamma$. We denote the group of deck transformations of a covering p by $\Gamma(p)$.

Local covers of graphs interact well with cliques, as stated in the following lemma. Let G be a graph. Recall that a *clique* X in G is a vertex-set which induces a complete subgraph of G, and X is a *maximal clique* in G if it is inclusion-maximal among the cliques. For a vertex v of a graph G, we denote the set of the maximal cliques in G containing v by $K_G(v)$. Moreover, for a vertex-subset X of G, we define $K_G(X)$ to be the set $\bigcup_{x \in X} K_G(x)$ of all cliques meeting X. For a set $\mathcal{X}$ of sets, we denote $\bigcup_{X \in \mathcal{X}} X$ by $\bigcup \mathcal{X}$.

Lemma 3.2.4 ([CJKK25], Lemma 5.15)**.** *Let G be a (possibly infinite) connected graph and $r \geqslant 3$ an integer. Then:*

(i) *For every clique $\hat{X}$ of G_r, its projection $p_r(\hat{X})$ is a clique of G, and p_r restricts to a bijection from $\hat{X}$ to $p_r(\hat{X})$.*

(ii) *For every clique X of G, there exists a clique $\hat{X}$ of G_r such that p_r restricts to a bijection from $\hat{X}$ to X.*

(iii) $(\bigcup K_{G_r}(\hat{X})) \cap (\bigcup K_{G_r}(\gamma(\hat{X}))) = \emptyset$ *for every clique $\hat{X}$ of G_r and every $\gamma \in \Gamma(p_r) \setminus \{\mathrm{id}_{G_r}\}$.*

For a given clique X of G, we refer to all cliques $\hat{X}$ of G_r such that p_r restricts to a bijection from $\hat{X}$ to X as *lifts of X to G_r*. Lemma 3.2.4 ensures that every clique of G_r is a lift of a clique in G, lifts of cliques of G exist, and two distinct lifts of a common clique are disjoint.

3.2.4 Binary cycle space of chordal graphs

Let $\mathcal{E}(G)$ be the $\mathbb{Z}_2$-vector space whose elements are the edge-subsets of G and whose vector addition is the symmetric difference $\triangle$. An element C of $\mathcal{E}(G)$ is *generated* by a subset $\mathscr{C}$ of $\mathcal{E}(G)$ if there is $\mathscr{C}' \subseteq \mathscr{C}$ such that the symmetric difference of all the elements of $\mathscr{C}'$ is C. If all

elements of a subset $\mathscr{D}$ of $\mathcal{E}(G)$ are generated by a set $\mathscr{C}$, the set $\mathscr{C}$ *generates* $\mathscr{D}$. The *binary cycle space* $\mathcal{Z}(G)$ of a graph G is the subspace of $\mathcal{E}(G)$ that consists of all elements generated by the set of (the edge sets of) cycles in G.

Binary cycle spaces of local covers are already generated by the short cycles:

Lemma 3.2.5 ([DJKK22], Lemma 4.6)**.** *The binary cycle space of the r-local cover G_r of a (possibly infinite) graph G for $r \in \mathbb{N}$ is generated by its cycles of length $\leqslant r$.*

The short cycles in r-chordal graphs are generated by triangles. A *triangle* in a graph G is a cycle of length 3 (equivalently: clique of size 3) in G.

Lemma 3.2.6. *Let r and ℓ be integers with $3 \leqslant \ell \leqslant r$. Every cycle O of length ℓ in a (possibly infinite) r-chordal graph G is the symmetric difference of $\ell - 2$ triangles of G whose vertex sets are included in $V(O)$.*

Proof. Let O be a cycle in G of length ℓ. We proceed by induction on ℓ. If O has length three, then O is a triangle and so trivially the symmetric difference of $|O| - 2 = 1$ triangles. Now assume that O has length $\ell > 3$ (and at most r) and that the statement holds for all cycles of length less than ℓ. Since G is r-chordal, O has a chord e in G. Thus, $O + e$ contains two cycles O_1 and O_2 whose intersection is exactly e and whose union covers O. In particular, O_1 and O_2 are both shorter than O, the sum $|O_1| + |O_2|$ of their sizes is $|O| + 2$, and their symmetric difference $O_1 \triangle O_2$ is O. Now, we may apply the inductive hypothesis to obtain sets $\mathscr{C}_1$ of size $|O_1| - 2$ and $\mathscr{C}_2$ of size $|O_2| - 2$ triangles whose respective symmetric differences are O_1 and O_2, and whose vertex sets are included in the respective $V(O_i) \subseteq V(O)$. It follows that the symmetric difference of all triangles in $\mathscr{C} \coloneqq \mathscr{C}_1 \triangle \mathscr{C}_2 = (\mathscr{C}_1 \setminus \mathscr{C}_2) \cup (\mathscr{C}_2 \setminus \mathscr{C}_1)$ is $O_1 \triangle O_2 = O$ and

$$|\mathscr{C}| = |\mathscr{C}_1| + |\mathscr{C}_2| = (|O_1| - 2) + (|O_2| - 2) = |O_1| + |O_2| - 4 = |O| - 2 = \ell - 2,$$

where the first equality holds, since no triangle in $\mathscr{C}_1$ is in $\mathscr{C}_2$ and vice versa, as they are contained in their respective $V(O_i)$ but O_1 and O_2 intersect only in the edge e. Therefore, $\mathscr{C}$ is the desired set of triangles for O. □

If a graph G is wheel-free, then the converse of Lemma 3.2.6 also holds.

Theorem 7. *Let G be a (possibly infinite) wheel-free graph and let r be a positive integer. Then G is r-chordal if and only if all its cycles of length at most r are generated by triangles.*

Proof. If G is r-chordal, then Lemma 3.2.6 ensures that all its cycles of length at most r are generated by triangles. Conversely, assume that all cycles of G of length at most r are generated by triangles. Consider a cycle O in G of length at least 4 and at most r. Since the cycles in G of length at most r are generated by triangles, there is a finite set of triangles whose symmetric difference is O. Fix $\mathscr{C}$ to be a smallest such set. If all three vertices of some triangle in $\mathscr{C}$ are in O, then O has a chord, as desired. Suppose for a contradiction that this is not the case. Let T be a triangle in $\mathscr{C}$. By assumption, there is a vertex c of T that is not on O.

Let $e_1, \dots, e_q$ be the edges of $\bigcup_{C \in \mathscr{C}} C$ incident with c, with endpoints c and $v_1, \dots, v_q$, respectively. Let H be the subgraph of G with vertex set $\{v_1, \dots, v_q\}$ and edge set consisting of edges e with both endpoints in $\{v_1, \dots, v_q\}$ such that e is an edge of a triangle in $\mathscr{C}$ containing c. Since the symmetric difference of the triangles in $\mathscr{C}$ is exactly O and c is not in O, it follows that every edge $e_1, \dots, e_q$ disappears in the symmetric difference of $\mathscr{C}$. Therefore, each edge e_i is contained in at least two triangles of $\mathscr{C}$. Since c is an end of each edge e_i, it follows that every triangle containing e_i is contained in the neighborhood of c, and in particular uses an edge from H. Therefore, H has minimum degree at least two, and so H contains a cycle C. We may assume by reordering that C is the symmetric difference of the ℓ triangles C_i formed by the edges $e_i, e_{i+1}, v_i v_{i+1}$, where $e_{\ell+1} = e_1$ and $v_{\ell+1} = v_1$.

If $G[V(C)]$ contains an induced cycle of length at least four, then $G[V(C) \cup \{c\}]$ contains a wheel with hub c as an induced subgraph, a contradiction, so $G[V(C)]$ is chordal. Now, Lemma 3.2.6 gives a set $\mathscr{C}^*$ of $\ell - 2$ triangles in $G[V(C)]$ whose symmetric difference is C, and $\mathscr{C}' := (\mathscr{C} \setminus \{C_1, \dots, C_\ell\}) \triangle (\mathscr{C}^* \cup \{C\})$ contradicts the choice of $\mathscr{C}$, because $\mathscr{C}'$ is smaller than $\mathscr{C}$. □

We thank the user kabel abel on Mathematics Stack Exchange for their inspiring thoughts [kab24] which we developed further to prove Theorem 7. As a simple corollary of Theorem 7, we obtain the affirmative answer to the question asked by Jonas Kobler that was left unanswered in their discussion.

Corollary 3.2.7. *A (possibly infinite) graph G is chordal if and only if G is wheel-free and its binary cycle space is generated by triangles.* □

3.2.5 Proof of Theorem 6

Finally, we prove Theorem 6.

Proof of Theorem 6.

Claim 1. (ii) *implies* (i).
Proof. Since the r-local cover G_r is chordal by (ii), its $r/2$-balls are chordal by Lemma 3.2.2. Since p_r is $r/2$-ball-preserving, the $r/2$-balls in G are also chordal. So G is r-locally chordal. ■

Claim 2. (i) *implies* (iii).

Proof. Every cycle in G of length at most r is contained in the $r/2$-ball around each of its vertices and every wheel is contained in the $3/2$-ball around its hub. Since $r \geqslant 3$ and the $r/2$-balls of G are chordal by (i), neither a cycle of length at most r nor a wheel W_n with $n \geqslant 4$ are induced subgraphs of G. ■

Claim 3. (iii) *implies* (ii).

Proof. Since p_r is $r/2$-ball-preserving and G satisfies (iii), it follows that G_r is r-chordal and wheel-free. By Lemma 3.2.5, the binary cycle space of G_r is generated by the cycles of length at most r. Since G_r is r-chordal, by Lemma 3.2.6, every cycle of length at most r in G_r is generated by triangles. Therefore, every cycle in G_r is generated by triangles, so by Theorem 7, G_r is chordal. ■

This completes the proof. □

We remark that, if one wants to avoid using the r-local cover, then, analogous to Claim 3, one may show that (iii) implies (i) by using instead of Lemma 3.2.5 that the binary cycle space of $r/2$-balls is generated by its cycles of length at most r. The latter is immediately from the fact that the fundamental cycles of every depth-first search tree of each $r/2$-ball have length at most r.

3.3 Equivalence: Graph-decompositions & region representations

In this section, we introduce our major tools to describe the structure of locally chordal graphs, graph-decompositions and region intersection graphs. We begin with several definitions.

3.3.1 Graph-decompositions & region intersection graphs

Graph-decompositions [DJKK22] generalise the notion of tree-decompositions by allowing the bags V_h of the decompositions $(H, \mathcal{V})$ to be arranged along general graphs H instead of just trees:

Let G be a graph. A *graph-decomposition* of G is a pair $\mathcal{H} = (H, \mathcal{V})$ consisting of a graph H and a family $\mathcal{V}$ of vertex-subsets V_h of G indexed by the nodes h of H which satisfies

(H1) $G = \bigcup_{h \in H} G[V_h]$, and

(H2) $H[\{h \in V(H) \mid v \in V_h\}]$ is connected for every $v \in V(G)$.

We refer to H as the *decomposition graph*, and also say that $(H, \mathcal{V})$ is a *decomposition over* H. The vertex-subsets V_h are called the *bags* of $\mathcal{H}$, and the $\{h \in V(H) \mid v \in V_h\}$ are called the *co-bags* of $\mathcal{H}$. One may flesh out the bags and co-bags by choosing spanning subgraphs G_h of $G[V_h]$ such that $G = \bigcup_{h \in H} G_h$ and connected spanning subgraph H_v of $H[\{h \in V(H) \mid v \in V_h\}]$. Then the G_h are the *parts* and the H_v are the *co-parts* of the decomposition $\mathcal{H}$. In this chapter, we will always consider the induced subgraphs $G_h := G[V_h]$ of G from (H1) as the *parts*. Whenever we introduce a graph-decomposition $(H, \mathcal{V})$ without a specific choice of co-parts, we tacitly assume that we have chosen any family of co-parts H_v.

A *tree-decomposition* $(T, \mathcal{V})$ is precisely a graph-decomposition $(H, \mathcal{V})$ whose decomposition graph H is the tree T.

We remark that any given bag V_h of a graph-decomposition $(H, \mathcal{V})$ of G is precisely the set of vertices v of G whose corresponding co-part H_v has h as a node h. Thus, one may

equivalently axiomatise *graph-decompositions* as follows. A *graph-decomposition* of G over a graph H is a map that assigns to each vertex v of G a subgraph H_v of H such that

(H1V) H_v is nonempty for every vertex v of G,

(H1E) $H_u \cap H_v$ is nonempty for every edge uv of G, and

(H2) H_v is connected for every vertex v of G.

The *regions* of a graph are its (nonempty) connected subgraphs. In this language, a graph-decomposition over H assigns to each vertex v a region H_v of H such that $H_u \cap H_v \neq \emptyset$ if $uv \in E(G)$. If the map also satisfies the converse of this condition, i.e. $H_u \cap H_v \neq \emptyset$ if and only if $uv \in E(G)$, then we obtain a concept called *region representations*:

A *region representation $v \mapsto H_v$ of G over a graph H* is a map that assigns to each vertex v of G a region H_v of H such that uv is an edge of G if and only if H_u and H_v have nonempty intersection. If such a region representation of G over H exists, then G is a *region intersection graph over H*. For a vertex-subset X of G, we write H_X for the union $\bigcup_{x\in X} H_x$ of the region corresponding to the elements of X. Given a class $\mathcal{H}$ of graphs, a graph G is a *region intersection graph over $\mathcal{H}$* if G is a region intersection graph over some $H \in \mathcal{H}$.

We remark that a region representation of G restricts to a region representation of each of its induced subgraphs. Moreover, an induced subgraph G' of G is connected if and only if the union $H_{V(G')}$ of its corresponding regions is connected. A region representation $v \mapsto T_v$ of G over a tree T is better known as a *subtree representation of G over the tree T*. Similarly, if G has a subtree representation over a tree T, then G is a *subtree intersection graph over T*.

The following is now immediate from the definitions, as the converse of (H1E) holds if and only if every bag V_h is a clique. A graph-decomposition $(H, \mathcal{V})$ of G is *into cliques* if each bag V_h of $(H, \mathcal{V})$ is a clique of G.

Lemma 3.3.1. *Let G and H be (possibly infinite) graphs.*

(i) *If $v \mapsto H_v$ is a region representation of G over H, then $(H, \mathcal{V})$ with $V_h := \{v \in V(G) \mid h \in V(H_v)\}$ is a graph-decomposition of G into cliques, for which the family $(H_v)_{v\in G}$ is a suitable choice of co-parts.*

(ii) *If $(H, \mathcal{V})$ is a graph-decomposition of G into cliques, then $v \mapsto H_v$ is a region representation of G over H, where the H_v are the co-parts of $(H, \mathcal{V})$.* □

We remark that Lemma 3.3.1 yields a 1-to-1 correspondence between region representations of a graph and its graph-decompositions into cliques, which allows us to discuss the graph-decomposition *corresponding* to a given region representation and vice versa.

Let G be a graph, and let Γ be a subgroup of the group $\mathrm{Aut}(G)$ of automorphisms of G. A graph-decomposition $(H, \mathcal{V})$ of G is Γ-*canonical* if there exists an action of Γ on H which commutes with $h \mapsto V_h$, i.e. $\gamma(V_h) = V_{\gamma\cdot h}$. Equivalently, considering the graph-decomposition presented as $v \mapsto H_v$, the graph-decomposition $(H, \mathcal{V})$ is Γ-canonical if $v \mapsto H_v$ commutes

with the action of Γ on H, i.e. $\gamma(H_v) = H_{\gamma(v)}$. Moreover, $(H, \mathcal{V})$ is *canonical* if it is $\mathrm{Aut}(G)$-canonical. A region representation $v \mapsto H_v$ of G over a graph H is Γ-*canonical* or *canonical* if its corresponding graph-decomposition is Γ-canonical or canonical, respectively.

3.3.2 Clique graphs & graph-decompositions into all maximal cliques

Let G be a graph. We denote the set of all maximal cliques in G by $K(G)$. A graph-decomposition $(H, \mathcal{V})$ of G is *into the maximal cliques in* G if $h \mapsto V_h$ is a bijection from the nodes of H to the set $K(G)$ of all maximal cliques in G.

The corresponding concept for region representations is *clique graphs*. By $\mathbf{K}(G)$ we denote the intersection graph of the maximal cliques in G, i.e. the graph $\mathbf{K}(G)$ whose vertex set is $K(G)$ and where two maximal cliques K_1 and K_2 of G are adjacent in $\mathbf{K}(G)$ when $K_1 \cap K_2 \neq \emptyset$. A *clique graph* of G is a spanning subgraph H of $\mathbf{K}(G)$ such that $v \mapsto H_v \coloneqq H[K_G(v)]$ is a region representation of G over H. A *clique tree*[5] of G is a clique graph which is a tree. Equivalently, a clique tree of G is a spanning tree T of $\mathbf{K}(G)$ such that $v \mapsto T_v \coloneqq T[K_G(v)]$ is a subtree representation of G over T.

The following is immediate from the definitions.

Lemma 3.3.2. *Let G be a (possibly infinite) graph. Suppose H is a spanning subgraph of $\mathbf{K}(G)$, and set $H_v \coloneqq H[K_G(v)]$ for every vertex v of G. Then $K_G(v) = K(B_{3/2}(v))$ for every vertex v of G. Moreover, for every two vertices u, v of G, we have*

(i) $K_{B_{3/2}(v)}(u) = K_G(u) \cap K_G(v)$, *and*

(ii) $H_u \cap H_v = H[K_G(u) \cap K_G(v)] = H_v[K_{B_{3/2}(v)}(u)]$. □

In light of Lemma 3.3.2, for a graph G, a spanning subgraph H of $\mathbf{K}(G)$ is a clique graph if and only if $H[K_G(v)]$ is connected for all vertices v of G. For the curious reader, we remark that one may generalise $K_G(v) = K(B_{3/2}(v))$ to $\bigcup_{v \in N^{d-1}[v]} K_G(v) = K(B_{r/2}(v))$ for r-locally chordal graphs with $r = 2d + 1$ with d a positive integer.

A clique graph H of G is Γ-*canonical* for a subgroup Γ of $\mathrm{Aut}(G)$ or *canonical* if its corresponding region representation of G is Γ-canonical or canonical, respectively. The automorphisms γ of G induce an action on the set of all maximal cliques in G, and thus on $\mathbf{K}(G)$, by $K \mapsto \gamma(K)$. Note that a clique graph H is Γ-canonical if and only if H is Γ-invariant, i.e. $\gamma(H) = H$ for every $\gamma \in \Gamma$.

The following is immediate from the definitions and Lemma 3.3.1.

Lemma 3.3.3. *Let G be a (possibly infinite) graph.*

(i) *If H is a clique graph of G, then assigning to each node K of H (where $K \in K(G)$) the bag $V_K \coloneqq K$ yields a graph-decomposition of G into the maximal cliques in G.*

(ii) *If $(H, \mathcal{V})$ is a graph-decomposition of G into the maximal cliques of G, then the map $h \mapsto V_h$ embeds H onto a clique graph of G.* □

[5] Clique trees are also known as *junction trees*, see e.g. [WJ08].

We remark that Lemma 3.3.3 yields a 1-to-1 correspondence between clique graphs of a graph and its graph-decompositions $(H, \mathcal{V})$ into its maximal cliques (up to the images of H under $h \mapsto V_h$ in $\mathbf{K}(G)$), which allows us to dicuss the graph-decomposition (into its maximal cliques) *corresponding* to a given clique graph and vice versa.

3.3.3 Canonical graph-decompositions via coverings

We conclude this section by discussing the relationship between graph-decompositions, region representations, and local covers. The main result here is a construction from [DJKK22] that folds each given canonical tree-decompositions of the r-local cover G_r of a graph G to a graph-decompositions of the graph G.

Construction 3.3.4 ([DJKK22], Construction 3.8 and Lemma 3.9)**.** Let G be a (possibly infinite) connected graph and $r \in \mathbb{N}$. Let $(T, \hat{\mathcal{V}})$ be a Γ-canonical tree-decomposition of G_r, where $\Gamma := \Gamma(p_r)$ is the group of deck transformations of p_r. We define the graph-decomposition $(H, \mathcal{V})$ *obtained from* $(T, \hat{\mathcal{V}})$ *by folding via* p_r as follows. The decomposition graph H is the orbit-graph of T under the action of Γ, i.e. $H = T/\Gamma$. The bag V_h corresponding to a node h of H is $p_r(\hat{V}_t)$ for a (equivalently: every) t in the Γ-orbit h.[6] As co-parts H_v of $(H, \mathcal{V})$ we choose the projections $p_r(T_{\hat{v}})$ of the co-part $T_{\hat{v}}$ of $\mathcal{T}$ corresponding to any lift $\hat{v}$ of v to G_r, where we denote also the quotient map from T to the orbit-graph $H = T/\Gamma$ by p_r.

Consider the equivalent description of the tree-decomposition $(T, \hat{\mathcal{V}})$ as $v \mapsto T_v$. Denote the quotient map from T to the orbit-graph $H = T/\Gamma$ also by p_r. Then the equivalent description of the graph-decomposition $(H, \mathcal{V})$ is given by $v \mapsto H_v = p_r(T_{\hat{v}})$ for a (equivalently: every) $\hat{v} \in p_r^{-1}(v)$.

It turns out that certain nice properties of tree-decompostions of G_r transfer to the graph-decompositions of G obtained by folding via p_r. First, canonical tree-decompositions yield canonical graph-decompositions:

Lemma 3.3.5 ([DJKK22], Lemma 3.12)**.** *Let $\mathcal{H}$ be a graph-decomposition of a (possibly infinite) graph G obtained from a Γ-canonical tree-decomposition $\mathcal{T}$ of G_r by folding via p_r. Then $\mathcal{H}$ is canonical if $\mathcal{T}$ is canonical.*

Next, tree-decompositions of G_r into cliques yield graph-decompositions of G into cliques:

Lemma 3.3.6. *Let G be a (possibly infinite) graph and $r \geqslant 3$ an integer. Suppose that $\mathcal{T}$ is a $\Gamma(p_r)$-canonical tree-decomposition of G_r and let $\mathcal{H}$ be the graph-decomposition $\mathcal{H}$ obtained from $\mathcal{T}$ by folding via p_r. Then $\mathcal{T}$ is into cliques if and only if $\mathcal{H}$ is into cliques.*

Equivalently, the map $v \mapsto T_v$ corresponding to the tree-decomposition $\mathcal{T}$ is a subtree representation of G_r if and only if the map $v \mapsto H_v$ corresponding to the graph-decomposition $\mathcal{H}$ is a region representation of G.

[6] Formally, the parts of $(H, \mathcal{V})$ are not the induced subgraphs $G[V_h]$ but only the projections $p_r(G_r[\hat{V}_t])$. We will omit this detail, as it will not be relevant to this chapter, as soon as we consider only decompositions into cliques.

In the context of Lemma 3.3.6, we now refer to the map $v \mapsto H_v$ as the region representation of G *obtained from the subtree representation* $v \mapsto T_v$ *by folding via* p_r.

Proof. If the tree-decomposition $\mathcal{T}$ is into cliques, then Lemma 3.2.4 (i) yields the graph-decomposition $\mathcal{H}(G)$ is into cliques. The converse follows from Lemma 3.2.4 (ii). □

Now, Lemma 3.3.6 suggests a proof outline to obtain nice graph-decompositions and region representations of locally chordal graphs. By Theorem 6, a graph G is r-locally chordal if and only if its r-local cover G_r is chordal. Therefore, if we could obtain a Γ-canonical tree-decomposition of G_r into cliques, then the graph-decomposition of G obtained by folding via p_r is also into cliques and corresponds to a region representation of G. It turns out that we can obtain canonical tree-decompositions of locally finite chordal graphs into cliques:

Theorem 8. *Let G be a locally finite, connected chordal graph. Then G admits a canonical tree-decomposition into cliques (equivalently: G has a canonical subtree representation).*

If we ask only for $\Gamma(p_r)$-canonical tree-decompositions rather than canonical tree-decompositions, we can even obtain such tree-decompositions into *maximal* cliques:

Theorem 3.3.7. *Let G be a connected, locally finite graph and $p\colon \hat{G} \to G$ a normal covering. If $\hat{G}$ is chordal, then $\hat{G}$ admits a $\Gamma(p)$-canonical tree-decomposition into its maximal cliques (equivalently: $\hat{G}$ has a $\Gamma(p)$-canonical clique tree).*

The proofs of Theorem 8 and Theorem 3.3.7 are in Section 3.13. In Section 3.5, Theorem 8 and Theorem 3.3.7 are used to construct nice graph-decompositions of 3-locally-chordal graphs into cliques.

In general, it is not possible that a graph admits a tree-decomposition which is both canonical and into its maximal cliques.

Example 3.3.8. The stars $K_{1,t}$ with $t \geqslant 3$ leaves are chordal graphs which do not admit a canonical tree-decomposition into its maximal cliques (equivalently: not have a canonical clique tree).

Proof. Consider the chordal graph $K_{1,t}$, the star with t leaves, for $t \geqslant 3$. For every tree-decomposition $(T, \mathcal{V})$ of $K_{1,t}$ into the maximal cliques of $K_{1,t}$, the bags $\mathcal{V}$ are exactly the set of edges of $K_{1,t}$. For any pair of edges e, f of $K_{1,t}$, there is an automorphism of G mapping e to f, so if $(T, \mathcal{V})$ is canonical, then T must be vertex-transitive. But a finite vertex-transitive tree has at most two vertices, since every finite tree with at least three vertices contains a vertex of degree one and a vertex of degree greater than one. Therefore, as $t \geqslant 3$, the graph $K_{1,t}$ does not admit a canonical tree-decomposition into its maximal cliques. □

3.4 Overview: Region representations of (locally) chordal graphs

In this section, we further discuss why region intersection graphs (equivalently: graph-decompositions into cliques) are the right tools to describe the structure of locally chordal

graphs. First, we review the classic characterizations of chordal graphs as those graphs which have subtree representations (equivalently: admit tree-decompositions into cliques). Then, we explain how these characterizations generalize to the locally chordal case.

It is well-known[7] that subtree intersection graphs are chordal.

Proposition 3.4.1 (Folklore)**.** *Every subtree intersection graph, finite or infinite, is chordal. Equivalently, every graph that admits a tree-decomposition into cliques is chordal.*

For finite graphs, it is well-known that the converse holds as well, see e.g. Gavril [Gav74, Theorem 2 and 3]: the connected chordal graphs are precisely the subtree intersection graphs (equivalently: those graphs which admits a tree-decomposition into cliques), or even more strictly those graphs that have a clique tree (equivalently: admit a tree-decomposition into its maximal cliques). Hofer-Temmel and Lehner extended this result to locally finite graphs:

Theorem 3.4.2 (Hofer-Temmel & Lehner [HL18], Theorem 3)**.** *Let G be a locally finite, connected graph. Then G has a clique tree (equivalently: G admits tree-decomposition into its maximal cliques).*

We remark that one may have derived Theorem 3.4.2 already from Diestel's more general forbidden simplicial minors characterization of (possibly infinite) countable subtree intersection graphs [Die88, Theorem 1]. But, in contrast to Diestel, Hofer-Temmel and Lehner extended the techniques designed for finite graphs to locally finite graphs.

More generally, a region representation $v \mapsto H_v$ of a graph G over a graph H is *acyclic* if $H_{V(G)}$ is acyclic. A clique graph H of G is *acyclic* if its corresponding region representation is acyclic (equivalently: H is acyclic). Thus, the following is obtained from Proposition 3.4.1 and from applying Theorem 3.4.2 to the components of the graph G.

Theorem 3.4.3. *Let G be a locally finite (not necessarily connected) graph. Then the following are equivalent:*

(i) *G is chordal.*

(ii) *G has an acyclic clique graph.*

(iii) *G has an acyclic region representation.* □

As cycles are finite, a region representation $v \mapsto H_v$ of G over a graph H is acyclic if and only if only if for every finite vertex-subset X of G the union H_X is acyclic. More generally, let H be a graph and let $\mathbf{H}$ be a family of subgraphs of H. Then, $\mathbf{H}$ is *r-acyclic* for an integer $r \geqslant 0$ if every union of at most r elements in $\mathbf{H}$ is acyclic. A region representation $v \mapsto H_v$ of a graph G over H is *r-acyclic* if $(H_v \mid v \in V(G))$ is r-acyclic. A clique graph of G is *r-acyclic* if its corresponding region representation is r-acyclic.

In light of the notion *r-acyclic*, we may refer to acyclic region-representations, families of subgraphs or clique graphs as being *∞-acyclic*. Theorem 3.4.3 says that chordal graphs

[7] This statement is immediate from the first paragraph of the proof presented in [Gav74, Theorem 3].

are exactly the graphs which have ∞-acyclic region representations. On the other end of the spectrum, we observe the following:

Example 3.4.4. Every graph G has a 2-acyclic region representation $v \mapsto H_v$ over some graph H such that the bag $V_h = \{v \in V(G) \mid h \in V(H_v)\}$ is a clique of size at most 2 for every node h of H.

Proof. Let G be a graph and let H be the graph formed from G by subdividing each edge e once by a vertex v_e. Then assign each vertex v of G the star S_v in H formed by v together with v_e for all edges e incident to v in G. Now $v \mapsto H_v$ is a 2-acyclic region representation of G. It is immediate from the definition that $V_{v_e} = \{u, v\}$ for every edge $e = uv$ of G and $V_v = \{v\}$ for every vertex v of G. □

Given Theorem 3.4.3 and Example 3.4.4, it is natural to ask whether there is a nice characterization of graphs which have r-acyclic region representations for general values of r. The class of graphs with r-acyclic region representations is (inclusion-wise) monotone decreasing as r goes to infinity, eventually approaching the class of chordal graphs. It turns out that r-locally chordal graphs are exactly the right characterization here.

Theorem 9. *Let G be a locally finite graph and $r \geqslant 3$ an integer. Then the following are equivalent:*

(i) *G is r-locally chordal.*

(ii) *G has an r-acyclic clique graph.*

(iii) *G has an r-acyclic canonical region representation.*

(iv) *G has an r-acyclic region representation.* □

3.4.1 Proof of Theorem 9

The proof of Theorem 9 involves two steps. First, we prove that Theorem 9 holds for $r = 3$:

Theorem 10. *Let G be a locally finite graph. The following are equivalent:*

(i) *G is wheel-free (equivalently: 3-locally chordal).*

(ii) *G has a 3-acyclic clique graph.*

(iii) *G has a 3-acyclic canonical region representation.*

(iv) *G has a 3-acyclic region representation.*

Next, we prove that for $r > 3$, it turns out that 3-acyclic region representations are equivalent to r-acyclic region representations for r-locally chordal graphs:

Theorem 11. *Let $v \mapsto H_v$ be a 3-acyclic region representation of a (possibly infinite) graph G over a graph H, and let $r > 3$ be an integer. Then G is r-locally chordal if and only if $v \mapsto H_v$ is r-acyclic.*

The proof of Theorem 10 appears in Section 3.5 and the proof of Theorem 11 appears in Section 3.6. Before proceeding to the proofs, we remark that 3-locally-chordal graphs have been analyzed through a similar view in the past. Specifically, in the nineties, Gavril [Gav96] studied 3-locally-chordal graphs[8] and obtained a characterization similar to that in Theorem 10. Next, we discuss his result and how it relates to ours.

3.4.2 Comparison to Gavril's work

Gavril [Gav96, Theorem 5] characterized the graphs that are 3-locally chordal (equivalently: wheel-free) as precisely the graphs which have Helly, 2-acyclic region representations. To state his result, let us first introduce *Helly*. Let H be a graph and let $\mathbf{H}$ be a family of subgraphs of H. Then, $\mathbf{H}$ is *Helly* if every finite set of pairwise intersecting elements of $\mathbf{H}$ has nonempty intersection. A region representation $v \mapsto H_v$ of a graph G over H is *Helly* if $(H_v \mid v \in V(G))$ is Helly. A clique graph of G is *Helly* if its corresponding region representation is. A graph-decomposition $(H, \mathcal{V})$ is *Helly* if the family $(H_v \mid v \in V(G))$ of its co-parts is Helly.

Theorem 3.4.5 (Gavril [Gav96], Theorems 4 & 5)**.** *Let G be a finite graph. The following are equivalent:*

(i) *G is 3-locally chordal (equivalently: wheel-free).*

(ii) *G has a Helly, 2-acyclic clique graph.*

(iii) *G has a Helly, 2-acylic region representation.*

We show in this section that Theorem 3.4.5 is the same as Theorem 9 (without (iii)) for $r = 3$. In particular, we prove that a region representation is 3-acyclic if and only if it is Helly and 2-acyclic.

Lemma 3.4.6. *Let $\mathbf{H}$ be a family of regions of a (possibly infinite) graph H. Then $\mathbf{H}$ is 3-acylic if and only if $\mathbf{H}$ is 2-acyclic and Helly.*

We remark that, by Lemma 3.4.6, the region representations in Theorem 9 are Helly, since they are in particular 3-acyclic.

To prove Lemma 3.4.6 not only for finite but also infinite graphs H, we first extend two statements of Gavril on finite graphs H to infinite graphs H.

Lemma 3.4.7 (Gavril [Gav96], Lemma 1 and Theorem 2)**.** *Let $\mathbf{H}$ be a 2-acyclic family of regions of a (possibly infinite) graph H.*

(i) *If the trees in $\mathbf{H}$ have an element in common, then the union of $\mathbf{H}$ is a subtree of H.*

[8] which he called *neighborhood chordal*.

(ii) $\mathbf{H}$ *is Helly if and only if the intersection of every three pairwise intersecting elements of* $\mathbf{H}$ *is nonempty.*

Proof. Gavril [Gav96, Lemma 1 and Theorem 2] proved (i) and (ii) for finite graph H, respectively. In the following, we reduce the case that H is infinite to the case that H is finite.

Claim 1. (i) *holds for infinite* H*, as well.*

Proof. The union $\bigcup \mathbf{H}$ is connected, as the elements of $\mathbf{H}$ are connected and all share a node, say t. Suppose for a contradiction that $\bigcup \mathbf{H}$ contains a cycle C. Then there is a finite subfamily $\mathbf{H}^* = (H_i \mid i \in [n])$ of $\mathbf{H}$ such that $C \subseteq \bigcup \mathbf{H}^*$. For each H_i in $\mathbf{H}^*$, we may choose a finite connected subgraph H_i' of the connected graph H_i that contains t and the finite subgraph $C \cap H_i$ of H_i. Now, their union $H' \coloneqq \bigcup_{i\in[n]} H_i'$ is a finite graph. As $H_i' \subseteq H_i$ and $\mathbf{H}$ is 2-acyclic, the family $\mathbf{H}'$ consisting of the regions H_i' of H' is 2-acyclic. The choice of the H_i' ensures that t is a common node of the H_i'. Applying Gavril's (i) for finite graphs to the family $\mathbf{H}'$ of regions in the finite graph H' yields that $\bigcup \mathbf{H}'$ is a tree. This contradicts that the choice of the H_i' yields that $\bigcup \mathbf{H}'$ contains the cycle C. ■

Claim 2. (ii) *holds for infinite* H*, as well.*

Proof. Let $\mathbf{H}^* = (H_i \mid i \in [n])$ be a finite subfamily of pairwise intersecting elements of $\mathbf{H}$. For each two H_i, H_j in $\mathbf{H}^*$, fix a node $h_{i,j} = h_{j,i}$ in their intersection $H_i \cap H_j$. Now, for each H_i in $\mathbf{H}^*$, we may choose a finite connected subgraph H_i' of the connected graph H that contains $h_{i,j}$ for every H_j in $\mathbf{H}^*$, since $\mathbf{H}^*$ is finite and H_i is connected. Now, their union $H' \coloneqq \bigcup_{i\in[n]} H_i'$ is a finite graph. As $H_i' \subseteq H_i$ and $\mathbf{H}$ is 2-acyclic, the family $\mathbf{H}' \coloneqq (H_i' \mid i \in [n])$ is 2-acyclic. The choice of the H_i' ensures that each two elements H_i', H_j' of $\mathbf{H}'$ still intersect in $h_{i,j} = h_{j,i}$. Now, applying Gavril's (ii) for finite graphs to the family of regions in the finite graph H' yields that $\bigcap \mathbf{H}'$ is nonempty. Since $\bigcap \mathbf{H}' \subseteq \bigcap \mathbf{H}$, also $\bigcap \mathbf{H}$ is nonempty. ■

This completes the proof. □

Additionally, we need the following lemma to deduce that a Helly and 2-acyclic family is 3-acyclic.

Lemma 3.4.8. *Let* $\mathbf{H}$ *be a 2-acyclic family of regions of a (possibly infinite) graph* H*. If the union of three elements of* $\mathbf{H}$ *contains a cycle, then each two of these regions intersect.*

Proof. Let T_1, T_2, T_3 be three regions in $\mathbf{H}$ whose union contains a cycle C. Note that the T_i are trees, as $\mathbf{H}$ is 1-acyclic. Moreover, since $\mathbf{H}$ is 2-acyclic, the cycle C contains an edge f_i of T_i which is not contained in any other T_j with $j \neq i$. Consider the path P_i in C which contains f_i and such that the intersection of P_i with $\bigcup_{j\neq i} T_j$ is exactly the ends of P_i; in particular, $P_i \subseteq T_i$.

To ensure that T_i intersects both T_j with $j \neq i$, it suffices to show that the ends of P_i are not both in the same T_j for a $j \neq i$. So suppose for a contradiction that the ends of P_i are

contained in a common T_j for a $j \neq i$. Consider the path Q_i in T_j joining the two ends of P_i. Because f_i is not in T_j but in $P_i \subseteq T_i$, the paths Q_i and P_i are distinct. As they join the same vertices, $P_i \cup Q_i \subseteq T_i \cup T_j$ contains a cycle, which contradicts that $\mathbf{H}$ is 2-acyclic. □

Corollary 3.4.9. *Let $v \mapsto H_v$ be a region representation of a graph G over a (possibly infinite) graph H. Suppose that a vertex-subset X of G is of minimum size such that H_X contains a cycle. If X has size at most 3, then X is a clique of G.* □

The following well-known fact about trees will be frequently used in the rest of the chapter. We include a proof for completeness.

Lemma 3.4.10. *Let T and T' be two trees. Then $T \cup T'$ is a tree if and only if $T \cap T'$ is a tree.*

Proof. Assume that $T \cup T'$ is a tree. As $T \cap T' \subseteq T \cup T'$, it remains to show that $T \cap T'$ is connected. For every two given nodes s, t of $T \cap T'$, their unique s–t path P in $T \cup T'$ is the unique s–t path in T and also in T'. Hence, P is the unique s–t path in $T \cap T'$. Thus, $T \cap T'$ is a tree.

Conversely, assume that $T \cap T'$ is a tree. In particular, $T \cap T'$ is nonempty. Thus, $T \cup T'$ is connected, as T and T' are trees. It remains to show that $T \cup T'$ is acyclic. Suppose for a contradiction that $T \cup T'$ is not acyclic. As T, T' are trees, we may pick cycle C in $T \cup T'$ such that $P \coloneqq T \cap C$ and $Q \coloneqq T' \cap C$ are paths. But as the endvertices of P (equivalently: Q) are in $T \cap T'$, and thus there is a path R in the tree $T \cap T'$ joining them. Now R is either distinct from P or from Q. Hence $R \cup P \subseteq T$ or $R \cup Q \subseteq T'$ contains a cycle, which is a contradiction to T, T' being trees. □

We are now ready to prove Lemma 3.4.6.

Proof of Lemma 3.4.6. Let $\mathbf{H}$ be a family of regions of a graph H.

Claim 1. *If $\mathbf{H}$ is Helly and 2-acyclic, then $\mathbf{H}$ is 3-acyclic.*

Proof. Consider a three-element subset $\mathbf{H}'$ of $\mathbf{H}$. Suppose for a contradiction that their union contains a cycle. Note that $\mathbf{H}'$ is inclusion-minimal with this property, as $\mathbf{H}$ is 2-acyclic. Hence, Lemma 3.4.8 yields that they are pairwise intersecting. Thus, since $\mathbf{H}$ is Helly, the trees in $\mathbf{H}'$ have an element in common. By Lemma 3.4.7 (i), it follows that the union of $\mathbf{H}'$ is a tree, and thus acyclic. ■

Claim 2. *If $\mathbf{H}$ is 3-acyclic, then $\mathbf{H}$ is Helly and 2-acyclic.*

Proof. By Lemma 3.4.7 (ii), it suffices to consider three pairwise intersecting elements T_1, T_2, T_3 of $\mathbf{H}$ and show that their intersection is nonempty. Since $\mathbf{H}$ is in particular 2-acyclic, $T_1 \cup T_2$ is acyclic. So, since their intersection $S \coloneqq T_1 \cap T_2$ is nonempty, it is a tree by Lemma 3.4.10. As we are done otherwise, we assume that T_3 does not intersect S. Fix $x_i \in T_i \cap T_3$ for $i = 1, 2$ witnessing that T_i and T_3 intersect. In particular, $x_i \notin S$, and thus $x_i \notin T_{3-i}$ by assumption. Let P be the path from x_1 to x_2 in T_3, and let Q be the path from x_1 to x_2 in $T_1 \cup T_2$. In

particular, P and Q are distinct because $P \subseteq T_3$ avoids S and, since $x_i \in T_i \setminus T_{3-i}$ for $i = 1, 2$, the path $Q \subseteq T_1 \cup T_2$ meets $T_1 \cap T_2 = S$. Now, $P \cup Q \subseteq \bigcup_{i=1}^{3} T_i$ contains a cycle, contradicting that $\mathbf{H}$ is 3-acyclic. Therefore, $\mathbf{H}$ is Helly. ■

Therefore, being 3-acyclic is equivalent to being 2-acyclic and Helly. □

3.5 Region representations of wheel-free graphs

In this section, we prove Theorem 10, which we restate here.

Theorem 10. *Let G be a locally finite graph. The following are equivalent:*

(i) *G is wheel-free (equivalently: 3-locally chordal).*

(ii) *G has a 3-acyclic clique graph.*

(iii) *G has a 3-acyclic canonical region representation.*

(iv) *G has a 3-acyclic region representation.*

We remark that, in view of Lemma 3.4.6, Gavril [Gav96, Theorems 4 and 5] showed the equivalences of all the above but (iii) for finite graphs. Our proof of Theorem 10 will indeed construct the 3-acyclic canonical region representation in (iii) as the one corresponding to the graph-decomposition $\mathcal{H}_r(G)$ obtained from $\mathcal{T}(N(G_r))$ by folding via p_r. Moreover, our construction of the 3-acyclic clique graph in (ii) corresponds to a graph-decomposition into the maximal cliques in G obtained from $\mathcal{H}_r(G)$ by contracting some edges of its decomposition graph H and merging the parts corresponding to the ends of the respective edges.

3.5.1 Interplay: Local covering & cliques

Before we are ready to prove Theorem 10, we need two lemmas regarding the interaction between cliques and local coverings. These are tools to help us project the structure from the r-local cover G_r down to G.

Lemma 3.5.1. *Let G be a (possibly infinite) connected graph and $r \geqslant 3$ an integer. Let $\hat{v} \mapsto T_{\hat{v}}$ be a $\Gamma(p_r)$-canonical subtree representation of G_r over a tree T and let $v \mapsto H_v$ be the region representation of G over $H = T/\Gamma(p_r)$ obtained from $v \mapsto T_v$ by folding via p_r. Then, for every clique X of G and lift $\hat{X}$ of X to G_r, the map p_r restricts to an isomorphism from $T_{\hat{X}}$ to H_X.*

Proof. Fix a clique X of G and a lift $\hat{X}$ of X to G_r. Since $H_X = \bigcup_{\hat{x} \in \hat{X}} p_r(T_{\hat{x}})$, the map p_r restricts to a surjection from $T_{\hat{X}}$ to H_X. Next, we show that p_r restricts to an injection from $T_{\hat{X}}$ to H_X. First, we need the following observation. Let t be a node of T and define $V_t \coloneqq \{\hat{v} \in V(G_r) \mid t \in V(T_{\hat{v}})\}$. Since $\hat{v} \mapsto T_{\hat{v}}$ is a subtree representation, it follows that V_t is a clique of G_r. Therefore, for every $\hat{v} \in V_t$, it holds that $V_t \subseteq \bigcup K_{G_r}(\hat{v})$.

By the above observation, $\bigcup_{t \in V(T_{\hat{x}})} V_t \subseteq \bigcup K_{G_r}(\hat{x})$ for every $\hat{x} \in \hat{X}$, and thus $\bigcup_{t \in V(T_{\hat{X}})} V_t \subseteq \bigcup K_{G_r}(\hat{X})$. Similarly, for all $\gamma \in \Gamma(p_r)$, it holds that $\bigcup_{t \in V(T_{\gamma(\hat{X})})} V_t \subseteq \bigcup K_{G_r}(\gamma(\hat{X}))$. By

Lemma 3.2.4 (iii), $\bigcup K_{G_r}(\hat{X})$ is disjoint from $\bigcup K_{G_r}(\gamma(\hat{X}))$ for all $\gamma \in \Gamma(p_r)$ distinct from the identity. Therefore, $T_{\hat{X}}$ is disjoint from $T_{\gamma(\hat{X})}$ for all $\gamma \in \Gamma(p_r)$ distinct from the identity. We conclude that $p_r(T_{\hat{X}})$ is isomorphic to $T_{\hat{X}}$, and so p_r restricts to an isomorphism from $T_{\hat{X}}$ to H_X. □

Lemma 3.5.2. *Let G be a (possibly infinite) connected graph. Suppose that $\hat{v} \mapsto T_{\hat{v}}$ is a $\Gamma(p_3)$-canonical subtree representation of G_3 over T. Let $v \mapsto H_v$ be the region representation of G obtained from $\hat{v} \mapsto T_{\hat{v}}$ by folding via p_3. Then, $v \mapsto H_v$ is 3-acyclic.*

Proof. Suppose for a contradiction that there is a set $X \subseteq V(G)$ of size at most three such that H_X contains a cycle. By Corollary 3.4.9, X is a clique. Fix some lift $\hat{X}$ of X to G_3. Since p_3 restricts to an isomorphism from $T_{\hat{X}}$ to H_X by Lemma 3.5.1 and $T_{\hat{X}}$ is a tree, it follows that H_X is a tree, a contradiction. □

In the context of Lemma 3.5.2, we have $H = \bigcup_{v \in G} H_v$, and thus it follows from $v \mapsto H_v$ being 2-acyclic by Lemma 3.5.2 that the H_v are induced subgraphs of H. In particular, the co-parts H_v of the corresponding graph-decomposition $(H, \mathcal{V})$ from Construction 3.3.4 are precisely the induced subgraph from (H2).

The intersection graph $\mathbf{K}(G_r)$ of the maximal cliques of G_r covers the intersection graph $\mathbf{K}(G)$ of the maximal cliques of G via p_r:

Lemma 3.5.3. *Let G be a (possibly infinite) graph and $r \geqslant 3$ an integer. Then p_r induces a covering map from $\mathbf{K}(G_r)$ to $\mathbf{K}(G)$ via $\hat{X} \mapsto p_r(\hat{X})$ which restricts to an isomorphism from $\mathbf{K}(G_r)[K_{G_r}(\hat{X})]$ to $\mathbf{K}(G)[K_G(X)]$ for every clique X of G and each of its lifts $\hat{X}$ to G_r. In particular, the maximal cliques of G_r are precisely the lifts of the maximal cliques of G.*

In the context of Lemma 3.5.3, we denote the covering map from $\mathbf{K}(G_r)$ to $\mathbf{K}(G)$ induced by p_r again by p_r.

Proof. First, we show that $V(\mathbf{K}(G_r))$ is the set $K(G)$ of lifts of the maximal cliques of G. Let $\hat{X}$ be a maximal clique of G_r. Then Lemma 3.2.4 (i) ensures that $X \coloneqq p_r(\hat{X})$ is a clique of G. Let Y be a maximal clique of G containing X. By Lemma 3.2.4 (ii), there is a lift $\hat{Y}$ of Y to G_r. By potentially shifting $\hat{Y}$ with a deck transformation of p_r, we assume that $\hat{Y}$ meets $\hat{X}$. Since distinct lifts of the clique X are disjoint, $\hat{Y}$ must contain the whole lift $\hat{X}$ of X. In particular, since $\hat{X}$ is a maximal clique, $\hat{Y} = \hat{X}$, and thus $X = Y$ is a maximal clique of G. Similarly, every maximal clique of G_r projects to a maximal clique of G. This proves that the maximal cliques of G_r are precisely the lifts of the maximal cliques of G and that p_r induces a surjection from $V(\mathbf{K}(G_r))$ to $V(\mathbf{K}(G))$.

Lemma 3.2.4 (iii) ensures that, for every clique X and each of its lifts $\hat{X}$ to G_r, the map induced by p_r restricts to a bijection from $K_{G_r}(\hat{X})$ to $K_G(X)$. Moreover, for every two cliques $Y, Z \in K_G(X)$ and their respective lifts $\hat{Y}, \hat{Z} \in K_{G_r}(\hat{X})$, the map p_r restricts to a bijection from $\hat{Y} \cap \hat{Z}$ to $Y \cap Z$. Thus, p_r induces an isomorphism from $\mathbf{K}(G_r)[K_{G_r}(\hat{X})]$ to $\mathbf{K}(G)[K_G(X)]$. In particular, p_r induces a covering map from $\mathbf{K}(G_r)$ to $\mathbf{K}(G)$. □

3.5.2 Properties of 3-acyclic region representations

Next, we prove two lemmas about 3-acyclic region representations.

Lemma 3.5.4. *Let $(H, \mathcal{V})$ be a 3-acyclic graph-decomposition of a connected (possibly infinite) graph G. If K is a finite clique of G, then some bag V_h of $(H, \mathcal{V})$ contains K.*

In particular, let $v \mapsto H_v$ be a 3-acyclic region representation of a connected (possibly infinite) graph G over some graph H. If K is a finite clique of G, then $\bigcap_{v \in K} H_v$ is nonempty.

Proof. By Lemma 3.4.6, $(H, \mathcal{V})$ is Helly. Since K is a clique, (H1) yields that the co-parts H_v with $v \in K$ are pairwise intersecting. Therefore, by Helly, there is some node h in the intersection of all the co-parts H_v with $v \in K$, i.e. $V_h \supseteq K$.

The in-particular part follows immediately from the correspondence of region representations and graph-decompositions into cliques (Lemma 3.3.1). □

Even though one could derive from Theorem 3.4.5 and Lemma 3.4.6 that the wheels W_n with $n \geqslant 4$, and thus also the graphs that contain W_n as an induced subgraph, do not have a 3-acyclic region representation, we include a basic proof.

Lemma 3.5.5. *A wheel W_n with $n \geqslant 4$ does not have a 3-acyclic region representation. Moreover, no (possibly infinite) graph G which contains W_n with $n \geqslant 4$ as an induced subgraph has a 3-acyclic region representation.*

Proof. Assume W_n has a 3-acyclic region representation $v \mapsto H_v$ over H. Let W_n consist of cycle $O = x_1 x_2 \dots x_n x_1$ and hub x complete to O.

Set $H_i \coloneqq H_{x_i} \cap H_x$ for $i = 1, \dots, n$. Since $v \mapsto H_v$ is 3-acyclic, in particular 2-acyclic, and $H_{x_i} \cap H_x \neq \emptyset$, it follows that H_i is a tree for $i = 1, \dots, n$. Next, we claim that $H_i \cap H_j \neq \emptyset$ if and only if $x_i x_j \in E(G)$. If $H_i \cap H_j \neq \emptyset$, then $H_{x_i} \cap H_{x_j} \neq \emptyset$, so $x_i x_j \in E(G)$. Conversely, suppose $x_i x_j \in E(G)$. Then, $x_i x_j x$ is a triangle in W_n, so by Lemma 3.5.4, $H_{x_i} \cap H_{x_j} \cap H_x \neq \emptyset$. In particular, $(H_{x_i} \cap H_x) \cap (H_{x_j} \cap H_x) \neq \emptyset$, and thus $H_i \cap H_j \neq \emptyset$.

Finally, consider the mapping $x \mapsto H_x$ and $x_i \mapsto H_i$ for $i = 1, \dots, n$. Now, this is a subtree representation of W_n over the tree H_x, contradicting Proposition 3.4.1. The moreover-part follows immediately from the fact that every (3-acyclic) region representation of a graph G restricts to a (3-acyclic) region representation of the induced subgraph of G. □

3.5.3 Proof of Theorem 10

We are now ready to prove Theorem 10.

Proof of Theorem 10. Clearly, (ii) implies (iv) and (iii) implies (iv). By Lemma 3.5.5, (iv) implies (i).

Claim 1. (i) *implies* (ii).

Proof. By Theorem 6, the 3-local cover G_3 of G is chordal, so G_3 has a $\Gamma(p_3)$-canonical clique tree T by Theorem 3.3.7. By Lemma 3.5.3, p_3 induces a covering map from $\mathbf{K}(G_3)$ to $\mathbf{K}(G)$ via $\hat{X} \mapsto p_3(\hat{X})$ for each node $\hat{X}$ of $\mathbf{K}(G_3)$. Since T is a spanning subgraph of $\mathbf{K}(G_3)$, its projection $H \coloneqq p_3(T)$ is a spanning subgraph of $\mathbf{K}(G)$. We claim that this graph H is a clique graph of G. Since H is a spanning subgraph of $\mathbf{K}(G)$, it remains to show that $H[K_G(v)]$ is a tree for each $v \in V(G)$. Fix a vertex v of G and consider $H_v = H[K_G(v)]$. By Lemma 3.5.3, the covering map from $\mathbf{K}(G_3)$ to $\mathbf{K}(G)$ induced by p_3 restricts to an isomorphism from $T[K_{G_3}(\hat{v})]$ to H_v for each lift $\hat{v}$ of v to G_3. Since T is a clique tree of G_3, it follows that $T[K_{G_3}(\hat{v})]$ is a tree, so H_v is a tree. This proves that H is a clique graph of G. By Lemma 3.5.2, H is 3-acyclic. ■

Claim 2. (i) *implies* (iii).

Proof. By Theorem 6, the 3-local cover G_3 of G is chordal. By Theorem 8, there is a canonical subtree representation $\hat{v} \mapsto T_{\hat{v}}$ of G_3 over a tree T. By Lemma 3.3.6, we now consider $v \mapsto H_v$ the region representation of G over the orbit-graph $H = T/\Gamma(p_3)$ obtained from $\hat{v} \mapsto T_{\hat{v}}$ by folding via p_3. By Lemma 3.3.5 and the equivalence between graph-decompositions and region representations, since $\hat{v} \mapsto T_{\hat{v}}$ is canonical, it follows that $v \mapsto H_v$ is canonical. Finally, $v \mapsto H_v$ is 3-acyclic by Lemma 3.5.2. ■

This completes the proof. □

3.6 Characterizations via 3-acyclic region representations

In this section, we prove Theorem 11, which we restate here.

Theorem 11. *Let $v \mapsto H_v$ be a 3-acyclic region representation of a (possibly infinite) graph G over a graph H, and let $r > 3$ be an integer. Then G is r-locally chordal if and only if $v \mapsto H_v$ is r-acyclic.*

First, we prove some preparatory lemmas, which enable us to prove Theorem 3.6.1, the statement analogous to Theorem 11 for chordal graphs.

Theorem 3.6.1. *Let $v \mapsto H_v$ be a 3-acyclic region representation of a (possibly infinite) graph G over a graph H. Then G is chordal if and only if the region representation $v \mapsto H_v$ is acyclic.*

Then we deduce from Theorem 3.6.1 the following more general version of Theorem 11:

Theorem 3.6.2. *Let $v \mapsto H_v$ be a 3-acyclic region representation of a (possibly infinite) graph G over a graph H, and let X be a vertex-subset of G. Then $G[X]$ is an induced cycle of length at least 4 if and only if X is inclusion-minimal such that H_X contains a cycle.*

3.6.1 Interplay: Helly, acyclic & clique graphs

It is well-known[9] that families of subtrees of a tree are Helly. Thus, subtree representations and clique trees are Helly.

[9] A proof is given in [Gav74, Lemma 1].

Theorem 3.6.3 (Folklore)**.** *Every family of subtrees of a tree is Helly. In particular, subtree representations are Helly.*

We remark that we restricted the notion of an infinite family **H** of subgraphs of a graph being *Helly* to just its finite subsets due to the following.

Example 3.6.4. The final segments of a fixed ray form an infinite family of subtrees (equivalently: regions) of a tree whose elements pairwise intersect, but whose intersection is empty, contrary to what one would hope for by extending the Helly property also to all infinite subfamilies.

3.6.2 Proof of Theorem 11

Now we prove the main results of this section.

A pair $\{A, B\}$ of vertex-subsets of a graph G G is a *separation* of G if $A \cup B = V(G)$ and there is no edge between $A \setminus B$ and $B \setminus A$. Its *separator* is $A \cap B$, its *sides* are A and B and its *proper sides* $A \setminus B$ and $B \setminus A$.

Proof of Theorem 3.6.1. For the sake of this proof, we may assume $H = H_{V(G)}$ by disregarding all edges which are not in H_v for some vertex v of G. Then the components G' of G are in a 1-to-1 correspondence with the components of H via $G' \mapsto H_{V(G')}$. As a graph is chordal if and only if each of its components is chordal, we may assume for this proof that G, and thus H, is connected.

If H is acyclic, and thus a tree, then $v \mapsto H_v$ is a subtree representation of G over a tree H, and thus by Proposition 3.4.1, G is chordal. To prove the converse, let us assume that G is chordal. As G is connected and $v \mapsto H_v$ represents G, the graph H is connected, as well. It remains to show that H is acyclic. Suppose for a contradiction that C is a cycle in H. By choosing for each edge f of the cycle C a vertex x_f of G such that f is an edge of H_{x_f}, we find a finite set $X = \{x_f \mid f \in E(C)\}$ such that H_X contains a cycle. Moreover, as $v \mapsto H_v$ represents G, the subgraph $G[X]$ is connected. Thus, it suffices that H_Y is a tree for every finite vertex-subset Y of G such that $G[Y]$ is connected We proceed by induction on the number of maximal cliques in the finite chordal graph $G[Y]$.

The base case follows immediately from Lemma 3.5.4 together with Lemma 3.4.7 (i). So we may assume that $G[Y]$ is not complete, i.e. there are two non-adjacent vertices $u, v \in Y$. In particular, u, v are separated by $Y \setminus \{u, v\}$ in $G[Y]$. Pick an inclusion-minimal u–v separator S in $G[Y]$. Fix a separation $\{A, B\}$ of $G[Y]$ with separator $A \cap B = S$ such that $u \in A$ and $v \in B$. As every maximal clique K of $G[Y]$ satisfies either $K \subseteq A$ or $K \subseteq B$, both subgraphs $G[A], G[B]$ induced on the sides of $\{A, B\}$ have fewer maximal cliques than $G[Y]$. Since $G[Y]$ is chordal, S is a clique by Theorem 3.2.1. Thus, the subgraphs $G[A], G[B]$ are connected. As induced subgraphs of a chordal graph, $G[A], G[B]$ are also chordal. The induction assumption applied to A and B yields that both H_A and H_B are trees. Since $\{A, B\}$ is a separation and $v \mapsto H_v$ represents G, their intersection $H_A \cap H_B$ is precisely H_S, which is also a tree by Lemma 3.5.4 together with Lemma 3.4.7 (i). Thus, $H_Y = H_A \cup H_B$ is a tree by Lemma 3.4.10. □

Indeed, we prove the converse of Theorem 3.6.2 in slightly more detail.

Lemma 3.6.5. *Let $v \mapsto H_v$ be a 2-acyclic region representation of a (possibly infinite) graph G over a graph H. If X is inclusion-minimal vertex-subset of G such that H_X contains a cycle, then $G[X]$ is an induced cycle.*

Proof. Assume that X is inclusion-minimal so that H_X contains a cycle. Since $v \mapsto H_v$ is 2-acyclic, X has size at least 3. If X has size 3, then $G[X]$ is an induced cycle by Corollary 3.4.9. Thus, we assume that X has size at least 4. As X is inclusion-minimal with the property that H_X contains a cycle and X has size at least 4, the restriction of $v \mapsto H_v$ to X is a 3-acyclic region representation of $G[X]$. Therefore, Theorem 3.6.1 applied to the restriction of $v \mapsto H_v$ to X yields that $G[X]$ is not chordal, i.e. $G[X]$ contains an induced cycle $G[Y]$ of length at least 4 for some $Y \subseteq X$. Now, by Theorem 3.6.1 applied to the restriction of $v \mapsto H_v$ to $Y \subseteq X$ that H_Y contains a cycle, and thus the inclusion-minimal choice of X yields $Y = X$, as desired. □

Proof of Theorem 3.6.2. The backwards direction follows immediately from Lemma 3.6.5. For the forwards direction, assume that $G[X]$ is an induced cycle of length at least 4. In particular, $G[X]$ is connected but not chordal. Thus, Theorem 3.6.1 applied to $v \mapsto H_v$ restricted to X, which is a 3-acyclic region representation of $G[X]$, yields that H_X is not acyclic, i.e. H_X contains a cycle. □

Proof of Theorem 11. Since G admits a 3-acyclic region representation $v \mapsto H_v$, no wheel W_n is an induced subgraph of G for $n \geqslant 4$ by Lemma 3.5.5. Theorem 3.6.2 yields that G has no induced cycles of length at least 4 and at most r if and only if $v \mapsto H_v$ is r-acyclic. The former is equivalent to being r-locally chordal by Theorem 6. □

3.7 Graph-decompositions of locally chordal graphs

In view of our characterization of locally chordal graphs using region representations and the equivalence between region representations and graph-decompositions, we have already the tools to prove a characterization of locally chordal graphs in terms of their graph-decompositions. Specifically, say that a graph-decomposition $(H, \mathcal{V})$ of a graph G is *r-acyclic* if its collection $(H_v \mid v \in V(G))$ of regions of H is r-acyclic. Now, Theorem 9 and the equivalence between region representations and graph-decompositions given by Lemma 3.3.1 and Lemma 3.3.3 together show the following:

Theorem 3.7.1. *Let G be a locally finite graph and $r \geqslant 3$ an integer. The following are equivalent.*

(i) *G is r-locally chordal.*

(ii) *G admits an r-acyclic graph-decomposition into its maximal cliques.*

(iii) *G admits a canonical r-acyclic graph-decomposition into cliques.*

(iv) *G admits an r-acyclic graph-decomposition into cliques.*

However, in the context of graph-decompositions, we would like to say something structurally stronger than the characterization in Theorem 3.7.1. A well-known key property of a tree-decomposition $(T, \mathcal{V})$ of a graph G is that every edge f of the decomposition tree T *induces a separation s^f of G* whose separator is the adhesion set V_f corresponding to f, see e.g. [Die25, Lemma 12.3.1]. Formally:

(T2') For every edge $f = t_1t_2$ of T and for $i = 1, 2$, set $A_i^f := \bigcup_{t \in V(T_i)} V_t$ and T_i is the component of $T - f$ which contains t_i. Then $s^f := \{A_1^f, A_2^f\}$ is a separation of G and its separator $A_1^f \cap A_2^f$ is the adhesion set $V_f = V_{t_1} \cap V_{t_2}$ corresponding to f.

A *bond* of a graph is an inclusion-minimal cut of the graph. Note that the bonds of a tree are precisely the set consisting of single edges. Analogous to (T2'), one may prove along the lines of the proof of [DJKK22, Lemma 3.4] that every bond c of the decomposition graph H of a graph-decomposition $(H, \mathcal{V})$ of G induces a separation s^c of G:

(H2') For every bond $c = \{A_1, A_2\}$ of H and for $i = 1, 2$, set $A_i^c := \bigcup_{h \in A_i} V_h$. Then $s^c := \{A_1^c, A_2^c\}$ is a separation of G and its separator $A_1^c \cap A_2^c$ is precisely the union $\bigcup_{f \in E_H(A_1, A_2)} V_f$ of the adhesion sets in the cut.

In view of (H2'), a graph-decomposition represents the structure of the underlying graph. Indeed, one could show that an equivalent definition of *graph-decomposition* is that a pair $(H, \mathcal{V})$ is a graph-decomposition if it satisfies (H1) and (H2').

Given a graph-decomposition $(H, \mathcal{V})$ of a graph G, the edges of H do not satisfy a property similar to (T2') in general; if the decomposition graph is not a tree, there is no reason to hope that already single edges of H correspond to separations of G. However, they do sometimes correspond to a local analog of separations. A very recent paper [CJKK25] investigates the relationship between graph-decompositions and local separations. It turns out that, since the graph-decompositions in our context are very well-structured, we can also prove that they interact nicely with local separations. The goal of this section is to make this interaction precise. We will at the end of the section be able to prove with Theorem 12 that r-locally chordal graphs are precisely those graphs which admit graph-decompositions whose edges induce r-local separation in an obvious way.

3.7.1 Additional properties of graph-decompositions.

Let $(H, \mathcal{V})$ be a graph-decomposition of G. For a vertex-subset X of G, let H_X be the subgraph of H formed by those nodes h and those edges $e = h_0h_1$ whose corresponding bags V_h and adhesion sets $V_e = V_{h_0} \cap V_{h_1}$ meet X, i.e. $H_X = \bigcup_{x \in X} H_x$. We remark that our notation is consistent, in that $H_v = H_{\{v\}}$. In the following, we consider graph-decomposition $(H, \mathcal{V})$ which additionally satisfy

(H2T) H_v is a tree for every vertex v of G, and

(H3) $H_u \cap H_v = H[\{h \in V(H) \mid u, v \in V_h\}]$ is connected for every edge uv of G.

We remark that both (H2T) and (H3) hold for every tree-decomposition. (H3) also holds for a graph-decomposition obtained from a tree-decomposition of G_r by folding via p_r [DJKK22, Paragraph after Lemma 3.9]. It turns out that the properties (H2T) and (H3) are equivalent to being 2-acyclic.

Lemma 3.7.2. *A graph-decomposition $(H, \mathcal{V})$ of a (possibly infinite) graph G into cliques satisfies (H2T) and (H3) if and only if it is 2-acyclic.*

To prove Lemma 3.7.2, we need the following fact.

Lemma 3.7.3. *Let $\mathbf{T}$ be a set of trees. Then $\mathbf{T}$ is 2-acyclic if and only if the intersection $T \cap T'$ of every two $T, T' \in \mathbf{T}$ is empty or connected.*

Proof. Assume that $\mathbf{T}$ is 2-acyclic. Let T and T' be two trees of $\mathbf{T}$ that meet, so in particular, their union $T \cup T'$ is connected. Since $\mathbf{T}$ is 2-acyclic, it follows that $T \cup T'$ is a tree. By Lemma 3.4.10, $T \cap T'$ is a tree and thus connected.

Conversely, assume the intersection $T \cap T'$ of any two trees $T, T' \in \mathbf{T}$ is empty or connected. We claim that $\mathbf{T}$ is 2-acyclic. Fix $T, T' \in \mathbf{T}$. If $T \cap T'$ is empty, then their union $T \cup T'$ is acyclic because T and T' are trees. Thus, T and T' meet, and so by assumption, $T \cap T'$ is a tree, as it is acyclic as subgraph of the tree T. Now, by Lemma 3.4.10, $T \cup T'$ is a tree, and thus acyclic. □

Now we prove:

Proof of Lemma 3.7.2. (H2T) precisely says that the regions H_v are trees. Thus, as $(H, \mathcal{V})$ is into cliques, the desired equivalence follows immediately from Lemma 3.7.3. □

3.7.2 Interplay: Graph-decompositions & local separations

We now formally introduce what we mean by a local separation from [CJKK25]. For a vertex-subset X of G, we denote the set of edges of G with exactly one end in X by $\partial_G X$. A *pre-separation* of G is a set $\{E_0, X, E_1\}$ with $X \subseteq V(G)$ and disjoint E_0, E_1 whose union is $\partial_G X$. We remark that every separation $\{A, B\}$ of G *induces the pre-separation* $\mathrm{pre}(\{A, B\}) := \{\partial_G(A \setminus B), A \cap B, \partial_G(B \setminus A)\}$, see Figure 3.2a.

Let X be a vertex-subset of G. A walk which starts and ends in X but internally avoids X is an *X-walk*. An *r-local X-walk* in G is an X-walk that is contained in a cycle of length at most r in G or consists of a single edge with both ends in X. We remark that every reduced non-trivial r-local X-walk is either an X-path or a walk once around a cycle of length at most r that meets X precisely once. A pre-separation $\{E_0, X, E_1\}$ of G is an *r-local separation* of G if there is no r-local X-walk which traverses an edge of E_0 and an edge of E_1. See Figure 3.2 for examples. The pre-separations $\mathrm{pre}(s)$ induced by separations s of G are r-local separations. We also note the following useful fact about r-local separations.

(a) A pre-separation induced by a separation. (b) A local separation not induced by a separation.

Figure 3.2: Examples of local separations $\{E_0, X, E_1\}$. The partition classes of $\{E_0, E_1\}$ are depicted in red and blue, respectively. The set X is formed by the yellow vertices.

Lemma 3.7.4. *Let G be a (possibly infinite) graph, r a positive integer, and $\{E_0, X, E_1\}$ an r-local separation of G. Then, a cycle of length at most r meets each of E_0 and E_1 an even number of times.*

Proof. Let O be a cycle in G of length at most r. By symmetry, it suffices to show that O meets E_1 an even number of times. If O does not intersect X, then O does not meet E_1, as desired. Otherwise, O meets X. Then, the reduced non-trivial r-local X-walks in O partition the edge set of O. As each of them meets E_1 exactly zero or two times, O meets E_1 an even number of times. □

Graph-decompositions are related to pre-separations in the following way.

Lemma 3.7.5. *Let $(H, \mathcal{G})$ be graph-decomposition of a (possibly infinite) graph G which satisfies (H2T) and (H3). Fix an edge $f = h_0h_1$ of H. For every vertex $v \in V_f$, let H_0^v, H_1^v be the components of $H_v - f$ containing h_0 and h_1, respectively, and set $E_i^v := \{vw \in \partial_G V_f \mid \exists h \in V(H_i^v) \colon w \in V_h\}$ for $i = 0, 1$. Then $E_i := \bigcup_{v \in V_f} E_i^v$ forms a pre-separation $\{E_0, V_f, E_1\}$ of G.*

We think that satisfying (H2T) and (H3) is not the only way to guarantee that every edge of the decomposition graph induces a pre-separation which has the potential to be a local separation. However, in this chapter, our constructed graph-decompositions always satisfy (H2T) and (H3). Thus, in the context of Lemma 3.7.5, we say that f *trivially induces the pre-separation* $\{E_0, V_f, E_1\}$ of G. If $\{E_0, V_f, E_1\}$ is an r-local separation of G, we say that f *trivially induces an r-local separation*.

Proof. Fix $v \in V_f$. By (H1), every edge vw in $\partial_G V_f$ is contained in the union of the two sets E_0^v and E_1^v. It remains to show that E_0^v and E_1^v are disjoint. Since H_v is a tree by (H2T), the two components H_0^v and H_1^v are distinct. Consider an edge $vw \in \partial_G V_f$. Since $vw \in \partial_G V_f$, it follows that $w \notin V_f$, and thus that f is not in the tree H_w. Since $(H, \mathcal{V})$ satisfies (H3), we know that $H_w \cap H_v$ is connected, so H_w meets exactly one of H_1^v and H_0^v. As edge vw was chosen arbitrarily, this proves that E_0 and E_1 are disjoint, and thus that $\{E_0, V_f, E_1\}$ is a pre-separation of G. □

We are now ready to state the main result of this section. A graph-decomposition $(H, \mathcal{V})$ of G *trivally induces* r*-local separations* if $(H, \mathcal{V})$ satisfies (H2T) and (H3) and the pre-separation induced by each edge of H is an r-local separation of G.

Theorem 12. *Let G be a locally finite graph and $r \geqslant 3$ an integer. The following are equivalent:*

(i) *G is r-locally chordal.*

(ii) *G admits a graph-decomposition into its maximal cliques which trivially induces r-local separations.*

(iii) *G admits a canonical graph-decomposition into cliques which trivially induces r-local separations.*

(iv) *G admits a graph-decomposition into cliques which trivially induces r-local separations.*

We conclude this subsection with a more precise statement regarding the structure of pre-separations in graphs with graph-decompositions into cliques which we will need later:

Lemma 3.7.6. *Let $(H, \mathcal{V})$ be a graph-decomposition of a (possibly infinite) graph G into cliques which satisfies (H2T) and (H3). Suppose $X \subseteq V(G)$ is inclusion-minimal such that H_X contains a cycle. Then, $G[X]$ is an induced cycle, and there exists an edge f in a cycle C in H_X with induced pre-separation $\{E_0, V_f, E_1\}$ such that $G[X]$ meets V_f in a single vertex whose two incident edges in $G[X]$ are contained in distinct E_i.*

Proof. By Lemma 3.7.2, $(H, \mathcal{V})$ is 2-acyclic. Choose $X \subseteq V(G)$ inclusion-minimal such that H_X contains a cycle C. By Lemma 3.6.5 via Lemma 3.3.1, $G[X]$ is an induced cycle of G. Fix $v \in X$. Since X is inclusion-minimal containing C, there is an edge f of C such that v is the unique vertex of X whose region H_v contains f, i.e. $f \in E(C \cap H_v) \setminus E(H_{X-v})$, where $X - v = X \setminus \{v\}$. Consider the pre-separation $\{E_0, V_f, E_1\}$ of G induced by f. Since $f \notin H_x$ for any $x \in X - v$, it follows that $V_f \cap X = \{v\}$.

Let P be an inclusion-maximal connected subgraph of $(C \cap H_v) - E(H_{X-v})$ containing f. Since $P \subseteq H_v$ and $(H, \mathcal{V})$ is 2-acyclic, $P \neq C$, so P is a subpath of C. Let h_0 and h_1 be the ends of P. By the choice of P, there are $x_0, x_1 \neq v$ in X such that h_i is a node of H_{x_i} for $i = 0, 1$. Since $(H, \mathcal{V})$ is 2-acyclic, $x_0 \neq x_1$. Now, up to symmetry, $vx_0 \in E_0$ and $vx_1 \in E_1$. This completes the proof. □

3.7.3 Equivalence: Characterization (iii) and Characterization (iv)

Our main tool to prove Theorem 12 is a theorem that proves an equivalence between r-acyclic region representations and graph-decompositions satisfying the properties we want.

Theorem 3.7.7. *Let G be a (possibly infinite) graph and $r \geqslant 3$ an integer. A region representation $v \mapsto H_v$ of G over a graph H is r-acyclic if and only if its corresponding graph-decomposition $(H, \mathcal{V})$ of G into cliques trivially induces r-local separations.*

Proof. First, suppose that G has an r-acyclic region representation $v \mapsto H_v$ over H. Let $(H, \mathcal{V})$ be the graph decomposition into cliques induced by $v \mapsto H_v$. As $v \mapsto H_v$ is r-acyclic for $r \geqslant 2$, the induced graph-decomposition is 2-acyclic, so Lemma 3.7.2 ensures that $(H, \mathcal{V})$ satisfies (H2T) and (H3). It remains to show that every edge of H induces an r-local separation of G.

Let $f = h_0h_1$ be an edge of H. We claim that f induces an r-local separation of G. Let $\{E_0, V_f, E_1\}$ be the pre-separation of G induced by f. Suppose for a contradiction that there is an r-local V_f-walk $W = v_0e_1v_1 \dots e_\ell v_\ell$ in G whose first edge e_1 is in E_0 and whose last edge e_ℓ is in E_1. We may choose W shortest among these. Since $(H, \mathcal{G})$ is into cliques, V_f is a clique. Hence, depending on whether v_0 equals v_ℓ or not, $v_0v_1 \dots v_\ell$ or $v_0v_1 \dots v_\ell v_0$ is a cycle C of length at most r in G because W is r-local.

We claim C is an induced cycle. Indeed, whenever there is a chord $e = v_iv_j$ of C with $0 < i < j < \ell$, then the walk $v_0v_1 \dots v_iv_jv_{j+1} \dots v_\ell$ is shorter r-local V_f-walk, which contradicts the choice of W. By symmetry, it remains to consider a chord $e = v_0v_j$ with $j = 2, \dots, \ell - 1$. As W is an V_f-walk, $e \in \partial V_f$, and thus there is $j \in \{0, 1\}$ with $e \in E_j$. If $j = 1$, then $v_0v_jv_{j+1} \dots v_\ell$ contradicts the choice of W. If $j = 2$, then $v_0v_1 \dots v_jv_0$ contradicts the choice of W.

Set $X \coloneqq \{v_0, \dots, v_\ell\}$. Since W is r-local, the size of X is at most r. As C is an induced cycle of G, we may apply Theorem 3.6.2 to obtain that $H_X = \bigcup_{x \in X} H_x$ contains a cycle, which contradicts that $v \mapsto H_v$ is r-acyclic. Therefore, every edge of H induces an r-local separation of G.

Conversely, suppose G admits a graph-decomposition $(H, \mathcal{V})$ into cliques which trivially induces r-local separations. Let $v \mapsto H_v$ be the region representation of G over H induced by $(H, \mathcal{V})$. It remains to show that it is r-acyclic.

Suppose for a contradiction that $v \mapsto H_v$ is not r-acyclic. Then there is a set X of at most r vertices such that $H_X = \bigcup_{x \in X} H_x$ contains a cycle. We may choose X to be inclusion-minimal. By Lemma 3.7.6, there is an edge f in a cycle O of H_x with its induced pre-separation $\{E_0, V_f, E_1\}$ such that the induced cycle $C \coloneqq G[X]$ meets X in a single vertex y and its two incident edges e_0, e_1 on C satisfy $e_i \in E_i$ for $i = 1, 2$. Since X meets V_f only in y and $|X| \leqslant r$, a closed walk W once around C based at y is an r-local V_f-walk. Then the two edges e_0, e_1 incident to y on C are contained in the same E_i of the r-local separation $\{E_0, V_f, E_1\}$, which is a contradiction. □

3.7.4 Proof of Theorem 12

We can now prove the main result of this section.

Proof of Theorem 12. By Theorem 9, G is r-locally chordal if and only if G has an r-acyclic region representation $v \mapsto H_v$. By Theorem 3.7.7, the graph-decomposition $(H, \mathcal{V})$ of G into cliques corresponding to $v \mapsto H_v$ trivially induces r-local separations. By the definition, $(H, \mathcal{V})$ is canonical if and only if $v \mapsto H_v$ is canonical. By Lemma 3.3.3, $(H, \mathcal{V})$ is into the maximal cliques in G if and only if $v \mapsto H_v$ is a clique graph. By Theorem 9, each of these is also equivalent to G being r-locally chordal. This completes the proof. □

3.8 Characterization via graph-decompositions that locally derive tree-decompositions

Recall that each $r/2$-ball $B_G(v, r/2)$ of a locally chordal graph is chordal, and thus admits a tree-decompositon $(T^v, \mathcal{V}^v)$ into cliques (Theorem 3.4.2). In this section, we show that indeed r-acyclic graph-decompositions of an r-locally chordal graph form a common description of some choice of such tree-decompositions $(T^v, \mathcal{V}^v)$ of the $r/2$-balls into cliques.

Our first observation is that honest r-acyclic graph-decompositions always have high-girth model graphs, i.e. the models are locally trees. A graph-decomposition is *honest* if all its bags and all its adhesion sets are non-empty.

Lemma 3.8.1. *Let $r \geqslant 3$ be an integer, let G be a graph, and let $(H, \mathcal{V})$ be an r-acyclic graph-decomposition of G which is honest. Then H has girth greater than r.*

Proof. Let C be a cycle of H of length $\ell \geqslant 3$. Since $(H, \mathcal{V})$ is honest, we may choose for each edge f of C a vertex v_f of G in its adhesion set V_f. Now, $C \subseteq \bigcup_{f \in E(C)} v_f$. Since $(H, \mathcal{V})$ is r-acyclic, it follows that $\ell > r$. Therefore, H has no cycle of length at most r, i.e. H has girth greater than r. □

First, we formalize what the locally induced decompositions are: Let $(H, \mathcal{V})$ be a decomposition of a graph G. The decomposition which $(H, \mathcal{V})$ *induces r-locally in G* are the families $(H_{v,r}, \mathcal{V}^{v,r})$ where $H_{v,r} := \bigcup\{H_v \mid d_G(u, v) < r/2\}$ and $V_h^{v,r} := V_h \cap N^{r/2}[v]$ for each node h of $H_{v,r}$. In general, these need not be graph-decompositions according to our axioms.

We show that an r-acyclic decomposition $(H, \mathcal{V})$ induces locally tree-decompositions. Even the converse is true if we extend the notion *r-acyclic* to *connected-r-acyclic* by restricting the considered vertex-subsets in r-acyclic to those which are connected in G. A decomposition $(H, \mathcal{V})$ of G is *connected-r-acyclic* if the union $\bigcup_{x \in X} H_v$ is acyclic for every set $X \subseteq V(G)$ of at most r vertices such that $G[X]$ is connected.

Theorem 3.8.2. *Let $(H, \mathcal{V})$ be a decomposition of a graph G, and let $r \geqslant 1$ be an integer. Then $(H, \mathcal{V})$ is connected-r-acyclic if and only if the decompositions it induces r-locally are tree-decompositions.*

First, characterise those graph-decomposition whose locally induced decompositions are tree-decompositions:

Lemma 3.8.3. *Let $(H, \mathcal{V})$ be a decomposition of a graph G, and let $r \geqslant 1$ be an integer. Then the decompositions that are r-locally induced by $(H, \mathcal{V})$ are tree-decompositions if and only if*

- *r is odd and $H_{N^{r/2}[v]}$ is a tree, or*
- *r is even and $H_{N^{r/2-1}[v] \cup \{u\}}$ is a tree for every $u \in N^{r/2}(v)$.*

Proof. Let v be a vertex of G, set $B := B_{r/2}(v)$, and let $(T, \mathcal{B}) := (H^{v,r}, \mathcal{V}^{v,r})$ be the decomposition $(H, \mathcal{V})$ induces r-locally. We warn the reader that we use terms such as *part* and *co-part*

in the context of the pair $(T, \mathcal{B})$ before knowing that $(T, \mathcal{B})$ is a decomposition of B. Let us make the following observations about the definition of $(T, \mathcal{N})$, without assuming anything about $(H, \mathcal{V})$ or $(T, \mathcal{V})$.

Since (H1) holds for $(H, \mathcal{G})$, (H1) also holds for the pair $(T, \mathcal{B})$ with respect to B. It follows immediately from the definition of T and the bags B_t that

- if r is odd, then the co-part T_u of a vertex $u \in B_{r/2}(v)$ in $(T, \mathcal{B})$ is the co-part H_u of u in $(H, \mathcal{G})$, and
- if r is even, then the co-part T_u of a vertex $u \in B_{r/2}(v)$ in $(T, \mathcal{B})$ is either H_u if $u \in N^{r/2-1}[v]$ or $H_u \cap T$ if $u \in N^{r/2}(v)$.

Hence, since (H2) holds for $(H, \mathcal{V})$, to show (H2) for $(T, \mathcal{B})$ it suffices to prove that $H_u \cap T$ is connected if r is even and $u \in N^{r/2}(v)$. So let us assume that r is even; in particular, $T = H_{N^{r/2-1}[v]}$.

No matter which of the two assumptions about the decomposition $(H, \mathcal{G})$ we make, $T = H_{N^{r/2-1}[v]}$ is a tree and every H_u is a tree for every vertex u of G. Thus, if $u \in N^{r/2}(v)$, then the subgraph $H_u \cap T$ of the tree T is a tree if and only if it is connected. The latter is precisely (H2) of $(T, \mathcal{B})$. The former is equivalent to the property described in the statement of the lemma by Lemma 3.4.10. Hence, the desired statement holds. □

Lemma 3.8.4. *Let G be a graph, let $r \geqslant 1$ be an integer, and let U be a connected set of r vertices of G. Then some vertex $v \in U$ has distance at most $r/2$ from every $u \in U$ and at most one $u \in U$ has distance precisely $r/2$ from v. In particular, if r is odd, then some vertex $v \in U$ has distance less than $r/2$ from every $u \in U$.*

Proof. Let P be a longest path in the induced subgraph $G[U]$. Let v be one of the 'central' vertices of P: If $||P||$ is even, then let v be the vertex such that $||vP|| = ||Pv||$. Otherwise, $||P||$ is odd, and we let v be the vertex such that $||vP|| - 1 = ||Pv||$. In both cases, $\ell := ||vP|| \geqslant ||Pv||$.

Let T be a spanning tree of $G[U]$ containing P. If some inclusion-maximal path Q in T that ends in v is longer than ℓ, then at least one of QvP or PvQ is a path, which is longer than P. Since this contradicts the choice of P, we may assume that all leaves of T have distance at most ℓ from v in T. Hence, if $\ell < r/2$, then the tree T witnesses that every $u \in U$ has distance at most $\ell < r/2$ from v, as desired. Thus, we assume $||vP|| = \ell = r/2$; in particular, r is even, since ℓ is an integer. As $P \subseteq G[U]$, P has at most r vertices, so $||P|| \leqslant r - 1 = r/2 + (r/2 - 1)$. Hence, the choice of v yields $||Pv|| = (r/2 - 1)$, so $V(P) = U$. Thus, the path $P = T$ shows that the statement holds.

The in-particular part follows from the fact that distances are always integers. □

Proof of Theorem 3.8.2. Assume that the decompositions r-locally induced by $(H, \mathcal{G})$ are tree-decompositions. We claim that $(H, \mathcal{G})$ is connected-r-acyclic. Indeed, let U be a connected set of at most r vertices of G. Since U is connected in G, Lemma 3.8.4 yields some vertex $v \in U$ that has distance at most $r/2$ from every $u \in U$ but at most one of them has distance precisely $r/2$. Hence, H_U is acyclic by Lemma 3.8.3.

To prove the converse, assume that $(H, \mathcal{G})$ is connected-r-acyclic. By induction on $r \geqslant 1$, we show that the decompositions r-locally induced by $(H, \mathcal{G})$ are tree-decompositions by checking the equivalent properties described in Lemma 3.8.3. Let v be a vertex of G. The base case $r = 1$ follows immediately from $(H, \mathcal{V})$ being connected-r-acyclic, as $N^{1/2}[v] = \{v\}$ is connected. For the induction step, we assume that $r \geqslant 2$ and the decompositions $(r-1)$-locally induced by $(H, \mathcal{G})$ are tree-decompositions.

Set $d := \lfloor r/2 \rfloor$, and $d' := d - 1$. Set $N' := N^{d'}[v]$. By the induction assumption, $H_{N'}$ is a tree. If r is even, then we set $W := \emptyset$, $N^* := N'$, and $H^* := H_{N^*}$. If r is odd, then let W be any subset of $N^d(v)$ such that, for $N^* := N' \cup W$, the union $H^* := H_{N^*}$ is a tree. Let $u \in N^d(v) \setminus W$ be arbitrary.

To conclude the induction step, it suffices to show that $H^* \cup H_u$ is acyclic, and thus a tree by (H1). Suppose for a contradiction that $H^* \cup H_u$ contains a cycle, i.e. H_u contains an H^*-path Q with ends h_1, h_2. Since $u \in N^d(v)$, there is $u' \in N'$ so that uu' is an edge of G, so $H_u \cup H_{u'}$ is a tree because $(H, \mathcal{G})$ is connected-r-acyclic for $r \geqslant 2$. Hence, the H^*-path Q in H_u yields that $H_{u'} \subseteq H^*$ avoids h_i for some $i = 1, 2$. As $H_u \cup H_{u'}$ is a tree, there is a Q–$H_{u'}$ path Q' in H_u. If $H_{u'}$ contains the other h_j with $j = 3 - i$, then we choose Q' as the trivial path in h_j. The choice of Q' guarantees that some h_k with $k = 1, 2$, say h_1, is neither in $H_{u'}$ nor in Q'.

Let $v_1 \in N^*$ with $h_1 \in H_{v_1}$ and fix a shortest v–v_1 path P_1 in G. Fix a shortest v–u' path P in G. Since $u' \in N'$, it follows that $|P| \leqslant d' + 1 = d$.[10] Similarly, since $v_1 \in N^*$, it follows that $|P_1| \leqslant d + 1$ if r is odd and $|P_1| \leqslant d$ if r is even. As uu' is an edge of G, the set $U := V(P_1) \cup V(P) \cup \{u\}$ is connected in G and has size at most

$$|V(P_1) \cup V(P) \cup \{u\}| \leqslant |V(P_1)| + |V(P)| - |\{v\}| + |\{u\}| \leqslant r,$$

As $U \setminus \{u\} = V(P_1) \cup V(P)$ is connected in G, (H1) and (H2) togehter yield that $H_{U \setminus \{u\}}$ is connected. Let Q^* be the path in $Q \cup Q'$ that starts in h_1 follows Q until it meets Q' and then follows Q' until it meets $H_{U \setminus \{u\}}$. Note that this will happen at latest at the end of Q', which is in $H_{u'} \subseteq H_{U \setminus \{u\}}$. The choice of Q' guarantees that Q^* is non-trivial, as Q' does not contain h_1. Hence, Q^* is a $H_{U \setminus \{u\}}$-path in H_u, so Q^* closes to a cycle in H_U, which contradicts that $(H, \mathcal{V})$ is connected-r-acyclic. □

Theorem 13. *Let G be a graph, let $r \geqslant 1$ be an integer and let $(H, \mathcal{V})$ be a decomposition of G into cliques. Then $(H, \mathcal{V})$ is an r-acyclic decomposition into cliques if and only if the decompositions r-locally induced by $(H, \mathcal{V})$ are tree-decompositions.*

If r is odd, then the r-locally induced decompositions are into cliques if and only if $(H, \mathcal{V})$ is into cliques. Unfortunately, if r is even, even though the decompositions r-locally induced by an r-acyclic graph-decomposition $(H, \mathcal{V})$ into cliques are tree-decompositions, they might not be into cliques. But we may derive from $(H, \mathcal{V})$ tree-decompositions of its $r/2$-balls into cliques, as follows.

First, let us observe the following corollary of Lemma 3.2.3:

[10] Recall that the number $|P|$ of vertices of a path P of length ℓ is $\ell + 1$.

Corollary 3.8.5. *Let v be a vertex in an r-locally chordal graph G, and let $r \geqslant 2$ be an even integer. Then, for every vertex $u \in N^{r/2}(v)$, its neighbourhood $N(u) \cap N^{r/2-1}(v)$ in $B_{r/2}(v)$ forms a clique.* □

Let G be an r-locally chordal graph, and let $(H, \mathcal{V})$ be an r-acyclic decomposition of G for $r \geqslant 3$. Assume now that r is even. Fix any vertex v of G. Obtain $H^{v,r,*}$ from the disjoint union of $H^{v,r-1}$ and one single node h_u for each $u \in N^{r/2}(v)$ by adding a single edge $f_u = h_u h'_u$ where h'_u is a node of $H^{v,r-1}$ whose bag $V^{v,r-1}_{h'_u}$ contains the neighbourhood $N(u) \cap N^{r/2-1}(v)$, which exists by Lemma 3.5.4, as $N(u) \cap N^{r/2-1}(v)$ is a clique by Corollary 3.8.5. Consider any node h of $H^{v,r,*}$. If h is a node of $H^{v,r-1}$, then $V^{v,r,*}_h := V^{v,r-1}_h$. Otherwise, there exist a unique $u \in N^{r/2}(v)$ such that $h = h_u$. Then $V^{v,r,*}_h := \{u\} \cup N(u) \cap N^{r/2-1}(v)$. We refer to the decomposition $(H^{v,r,*}, \mathcal{V}^{v,r,*})$ as a decomposition *r-locally derived by* $(H, \mathcal{V})$. If r is odd, then we refer to the decomposition r-locally induced by $(H, \mathcal{V})$ to a decomposition *r-locally derived by* $(H, \mathcal{V})$.

Theorem 3.8.6. *Let $(H, \mathcal{V})$ be an r-acyclic decomposition of a graph G into cliques with $r \geqslant 3$. Then the decompositions r-locally derived by $(H, \mathcal{V})$ are tree-decompositions into cliques.*

Proof. The decomposition $(H, \mathcal{V})$ witnesses that G is r-locally chordal by Theorem 9. Fix a vertex v of G. If r is odd, then this follows from Theorem 13 and the definition of the decompositions r-locally induced by $(H, \mathcal{V})$, as $B_{r/2}(v)$ is an induced subgraph of G. Thus, we may assume that r is even. Since $r-1$ is odd, the decomposition $(H^{v,r-1}, \mathcal{V}^{v,r-1})$ $(r-1)$-locally induced by $(H, \mathcal{V})$ is a tree-decomposition of $B_{(r-1)/2}(v)$ into cliques. It is immediate from the above construction of the decomposition $(H^{v,r,*}, \mathcal{V}^{v,r,*})$ r-locally derived at v that it is a tree-decomposition of $B_{r/2}(v)$ into cliques. □

3.9 Characterization via minimal local separators

Recall that chordal graphs can also be characterized by the structure of their minimal separators.

Theorem 3.2.1 (Dirac [Dir61], Theorem 1)**.** *A (possibly infinite)*[11] *graph is chordal if and only if every minimal separator is a clique.*

How might Theorem 3.2.1 be generalized to locally chordal graphs? As we have seen in Theorem 12, the local analog of a separation is an *r-local separation*. In [CJKK25], we also introduced a local analog of (tight) separator. To define it here, we first need to introduce their local analog of components.

Let X be a vertex-subset of a graph G. The *r-local components at X* in G are the equivalence classes of the transitive closure of the relation on ∂X defined by letting

$$e \sim_r f \text{ at } X \text{ in } G :\Leftrightarrow e \text{ and } f \text{ are the first and last edge on the same } r\text{-local } X\text{-walk in } G, \text{ or } e = f.$$

[11] It is immediate from the proof presented in [BP93, Theorem 2.1] that this also holds for infinite graphs.

Now X is an *r-local separator* of G if there are at least two r-local components at X. An r-local component F at X is *tight at X* if every vertex in X is incident to some edge in F. Moreover, an r-local separator X of G is *tight* if there are at least two tight r-local components at X in G.

We remark that if one would set $r := \infty$, the ∞-local components at X precisely correspond to the *componental cuts at X*, i.e. the sets $E_G(C, X)$ consisting of the edges between X and a fixed component C of $G - X$. Moreover, the tight ∞-local components are precisely the componental cuts given by the components of $G - X$ that are tight at X. Now, one might hope that r-locally chordal graphs are precisely those graphs whose tight r-local separators are cliques. Unfortunately, not all tight r-local separators of an r-locally chordal graphs are cliques:

Example 3.9.1. The graph depicted in Figure 3.3 is r-locally chordal graph for $r = 3$ and has a tight r-local separator that is not a clique.

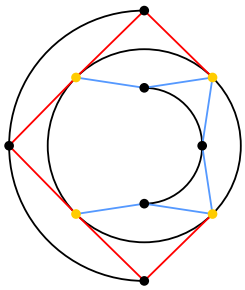

Figure 3.3: A tight r-local separator X, formed by the yellow vertices, in an r-locally chordal graph that is not a clique for $r = 3$. The two r-local components at X are indicated in red and blue, respectively.

Proof. Let G be the graph depicted in Figure 3.3. Each 3/2-ball of G is chordal, so G is 3-locally chordal. The red and blue edges, respectively, form tight 3-local components at the set X of yellow vertices, but X is not a clique in G. □

Recall that a set X is a tight separator of a graph G if and only if X is a *minimal (vertex) separator*, i.e. there are two vertices u, w such that X is an inclusion-minimal u–w separator. The example depicted in Figure 3.3 suggests that this equivalence does not hold for local separators. To put this on a formal basis, we define the following.

Let u, w be two vertices of a graph G of distance 2. Fix a vertex v of G with $uv, vw \in E(G)$. A vertex-subset X of G *r-locally separates u and w* in G with respect to v, or is an *r-local u–w separator* in G with respect to v, if X avoids u, w, contains v, and uv and vw are in distinct r-local components of X. When $r \geqslant 4$ or when $r \geqslant 3$ but X is a clique, we remark that this definition is independent of the choice of the vertex v. Note that $X := V(G) \setminus \{u, w\}$ always forms an r-local u–w-separator of G (with respect to v).

A vertex-subset X of G is a *minimal r-local (vertex) separator* in G if there are two vertices u, w of distance 2 in G such that X is an inclusion-minimal r-local u–w separator in G with

respect to some vertex v with $uv, vw \in E(G)$. Note that the tight r-local separator depicted in Figure 3.3 is indeed not a minimal r-local separator. With this definition, Theorem 3.2.1 extends to locally chordal graphs:

Theorem 14. *For an integer $r \geqslant 3$, a (possibly infinite) graph G is r-locally chordal if and only if every minimal r-local separator in G is a clique.*

Proof. We show the two directions separately.

Claim 1. *If every minimal r-local separator in a graph G is a clique, then G is r-locally chordal.*

Proof. By Theorem 6, it suffices to prove that G neither contains a wheel W_n with $n \geqslant 4$ nor a cycle of length at least four and at most r as an induced subgraph. Suppose for a contradiction that G contains a cycle O of length at least 4 but at most r (analogously: contains a wheel W_n with $n \geqslant 4$, rim O) as an induced subgraph. Let u, v, w be three consecutive vertices on O (analogously: let u, w be two non-adjacent vertices on the rim O and let v be the hub of wheel). In particular, u, w have distance 2 and $uv, vw \in E(G)$. Let X be an inclusion-minimal r-local u–w separator in G. Since uv and vw are in distinct r-local components at X, the short cycle O meets X in at least one vertex x other than v (analogously: each P_i of the two u–w paths P_1, P_2 in the rim O meets X in at least one vertex x_i). Thus, X is a minimal r-local separator in G, which contains two non-adjacent vertices v, x (analogously: x_1, x_2) of G. This is a contradiction, since X is a clique by assumption. ■

Claim 2. *If G is r-locally chordal, then every minimal r-local separator in G is a clique.*

Proof. By Theorem 3.7.1, the r-locally chordal graph G admits an r-acyclic graph-decomposition $(H, \mathcal{V})$ into cliques. Let X be a minimal r-local separator, i.e. there are three vertices u, v, w with $uv, vw \in E(G)$ such that X is an inclusion-minimal u–w separator in G with respect to v. Thus by definition, X avoids u, w, contains v, and the r-local components F_u, F_w at X which contain uv and vw, respectively, are distinct. Consider the sets N_u and N_w of neighbors of v in G outside X that are incident to edges in the r-local component F_u and F_w, respectively. Set $N' \coloneqq N_G(x) \setminus (X \cup N_u \cup N_w)$. Since F_u and F_w are distinct r-local components, N_u and N_w are disjoint and, as $r \geqslant 3$, the three sets N_u, N_w, N' are pairwise anticomplete. Hence, H_{N_u}, H_{N_w}, and $H_{N'}$ are pairwise disjoint because $(H, \mathcal{V})$ is into cliques. As u, w are neighbors of x in G, H_x meets H_u and H_w. Since H_x is connected, there is an H_u–H_w path P in H_x. As $H_{N_u} \supseteq H_u$ and $H_{N_w} \supseteq H_w$ are disjoint, there is a last edge on P leaving H_{N_u}; in particular, f has precisely one endnode in H_{N_u}. Hence, since H_{N_u}, H_{N_w}, and $H_{N'}$ are pairwise disjoint, f is contained in neither of them, so $V_f \subseteq N_G(x) \setminus (N_u \cup N_w \cup N') = X$ because $x \in V_f$ and $(H, \mathcal{V})$ is into cliques. By Theorem 3.7.7, the edge f trivially induces the pre-separation $\{E_0, V_f, E_1\}$ and $\{E_0, V_f, E_1\}$ is an r-local separation of G. The path P ensures that $uv \in E_0^v$ and $vw \in E_1^v$, where we refer to the definition of $\{E_0, V_f, E_1\}$ introduced in Lemma 3.7.5. Hence, $V_f \subseteq X$ r-locally separates u, w in G with respect to v. By the inclusion-wise minimality of X, we have $V_f = X$, so, X is a clique as desired. ■ This completes the proof. □

3.10 Computing clique graphs

In Theorem 9, we prove that G is r-locally chordal if and only if G has an r-acyclic clique graph. In this section, we provide a simple and efficient algorithm to compute r-acyclic clique graphs of finite r-locally chordal graphs. This algorithm is a generalization of a well-known algorithm to compute clique trees of chordal graphs.

The algorithm in particular finds a spanning subgraph of the intersection graph $\mathbf{K}(G)$ of the maximal cliques in G. For the rest of the chapter, we equip the edges of $\mathbf{K}(G)$ with weights according to the number of vertices in the intersection they represent; i.e. the *weight* of an edge K_1K_2 is $|K_1 \cap K_2|$. Let $\mathbf{K}$ be an edge-weighted graph. The *total weight* $w_{\mathbf{K}}(H)$ of a subgraph H of $\mathbf{K}$ is the sum of the weights of the edges in H. A *maximum weight spanning tree of* $\mathbf{K}$ is a spanning tree of $\mathbf{K}$ which has maximum total weight among all spanning trees of $\mathbf{K}$.

Maximum weight spanning trees of clique graphs are related to chordal graphs through the following theorem.

Theorem 3.10.1 ([Gav87], Theorem 1)**.** *Let G be a finite connected chordal graph. Clique trees of G are precisely the maximum weight spanning trees of $\mathbf{K}(G)$.*

Theorem 3.10.1 is especially useful since there is an efficient algorithm to compute maximum weight spanning trees:

Theorem 3.10.2 (Kruskal [Kru56])**.** *There is an algorithm, Algorithm 3.10.3, to compute a maximum weight spanning tree of every given finite connected edge-weighted graph G that runs in time $O(m \log n)$, where $m \coloneqq |E(G)|$ and $n \coloneqq |V(G)|$.*

Algorithm 3.10.3 (Kruskal [Kru56])**.** Let H be a finite connected edge-weighted input graph.

(1) Order the edges of H in order $e_1, \ldots, e_m$ of decreasing weight.

(2) Let T_0 be the edgeless graph with vertex set $V(H)$.

(3) For i from 1 to m, let $T_i \coloneqq T_{i-1} + e_i$ if $T_i + e_i$ is acyclic. Else, let $T_i \coloneqq T_{i-1}$

(4) After step m, set $T \coloneqq T_m$.

In view of the previous three results, computing a clique tree of a chordal graph is equivalent to computing a maximum weight spanning tree of $\mathbf{K}(G)$. In order to assess the running time of such an algorithm, we need to understand the size of $\mathbf{K}(G)$.

Proposition 3.10.4 (Folklore, e.g. [Gav72; Ros70])**.** *A chordal graph G on n vertices has at most n maximal cliques. Thus, $\mathbf{K}(G)$ has at most $\binom{n}{2} \leqslant n^2$ edges.*

All in all, ones derives the following from Theorems 3.10.1 and 3.10.2 and Proposition 3.10.4.

Theorem 3.10.5. *Given a finite connected chordal graph G on n vertices, computing $\mathbf{K}(G)$ and running Algorithm 3.10.3 on $\mathbf{K}(G)$ yields a clique tree of G in running time $O(n^2 \log n)$.*

In this section, we provide a parallel algorithm that, when G is a finite r-locally chordal (equivalently: wheel-free) graph, computes an r-acyclic clique graph of G.

Theorem 15. *There is a parallel algorithm, Algorithm 16, to compute an r-acyclic clique graph of any given finite r-locally chordal graph G that runs on $n = |V(G)|$ processors in time $O(\Delta^2 \log \Delta)$, where Δ is the maximum degree of G.*

We remark that Gavril's proof [Gav96, Theorem 5] of Theorem 10 (i)→(ii) is an algorithm with running time $O(n^3 \cdot \log n)$.[12]

By Theorem 11, it suffices to consider the case $r = 3$. Inspired by Gavril's work, our parallel algorithm works as follows.

Algorithm 16. *Let G be a 3-locally chordal (equivalently: wheel-free) input graph on $[n]$. Run on processor $v \in [n]$ the following:*

(1) Compute the intersection graph $\mathbf{K}(B_{3/2}(v))$ of the maximal cliques in $B_{3/2}(v)$.

(2) Order the edges of $\mathbf{K}(B_{3/2}(v))$ in order $e^v_1, \dots, e^v_{m(v)}$ of decreasing weight; decide ties of two edges by the lexicographic order on the union of their respective ends.

(3) Let H^v_0 be the edgeless graph with vertex set $V(\mathbf{K}(B_{3/2}(v)))$.

(4) For i from 1 to $m(v)$, let $H^v_i := H^v_{i-1} + e^v_i$ if $H^v_{i-1} + e^v_i$ is acyclic. Else, let $H^v_i := H^v_{i-1}$.

(5) After step $m(v)$, set $H^v := H^v_{m(v)}$.

After all n processor are finished, set $H := \bigcup_{v \in V(G)} H^v$.

3.10.1 Analysis of the algorithm

Algorithm 16 essentially works by running Kruskal's algorithm to find a clique tree of $B_{3/2}(v)$ for every $v \in V(G)$ simultaneously. In this section we prove that this algorithm indeed returns an r-acyclic clique graph when the input graph is r-locally chordal.

For this, by Theorem 11, we only need to ensure that the spanning subgraph H of $\mathbf{K}(G)$ returned by Algorithm 16 is a 3-acyclic clique graph. Indeed, it suffices to show that H is a 2-acyclic clique graph due to the following lemma.

Lemma 3.10.6. *Let G be a (possibly infinite) graph. Then every 2-acyclic clique graph H of G is Helly, and thus 3-acyclic.*

Proof. Let H be a 2-acyclic clique graph of G, and let $v \mapsto H_v := H[K_G(v)]$ be its corresponding region representation of G.

By Lemma 3.4.6, the 2-acyclic clique graph H is 3-acyclic, if H is Helly. To show that H is Helly, let X be a finite vertex-subset of G such that each two H_x, H_y with $x, y \in X$

[12] In fact, Gavril claims that his algorithm has running time $O(n^3)$. As far as we understand his argument, it yields $O(n^3 \cdot \log n)$.

intersect. In particular, X is a clique of G. Let K be a maximal clique of G containing X, which is thus a node of every H_x with $x \in X$. Hence, Lemma 3.4.7 (i) yields that H_X is a tree. Therefore, $x \mapsto H_x$ is a subtree representation of the complete graph $G[X]$ over the tree H_X. By Theorem 3.6.3, $\bigcap_{x \in X} H_x$ is nonempty. All in all, the 2-acyclic clique graph H is Helly, and thus 3-acyclic. □

We note that we cannot push this further to 1-acyclic due to the following example.

Example 3.10.7. Consider the wheel W_n with $n \geqslant 4$ whose hub is v and whose rim is the cycle $v_1 v_2 \dots v_n v_1$. Let G be the graph obtained from W_n by adding a vertex x and the edges xv_i for $i = 1, 2$. Now, G has a 1-acyclic clique graph but not a 2-acyclic clique graph. However, its induced subgraph W_n does not even have a 1-acyclic clique graph.

Proof.

Claim 1. *The wheel W_n with $n \geqslant 4$ has no 1-acyclic clique graph.*
Proof. The maximal cliques of W_n are $D_i := \{v, v_i, v_{i+1}\}$ for $i \in [n]$, where $v_{n+1} = v_1$, so $\mathbf{K}(W_n)$ is the cycle $D_1 D_2 \dots D_n D_1$. Now consider a clique graph H of W_n. Let $i \in [n]$. Consider D_0 to be D_n. Since $H_{v_i} = H[K_{W_n}(v_i)] = G[\{D_{i-1}, D_i\}]$ is connected, the edge $D_{i-1}D_i$ is in H. Thus, H is the whole cycle $\mathbf{K}(W_n)$. Since $H_v = H[K_{W_n}(v)] = H$, the clique graph H is not 1-acyclic. ■

Claim 2. *G has a 1-acyclic clique graph.*

Proof. The maximal cliques of G are $D_1, \dots, D_n$ and also $D := \{x, v_1, v_2\}$. Hence, the clique graph $\mathbf{K}(G)$ is obtained from $\mathbf{K}(W_n)$ by adding D and the edges DD_n, DD_1, D_2. Let H be the spanning cycle $DD_1 D_2 \dots D_n D$. We claim that H is a 1-acyclic clique graph of G. For $i \in [n] \setminus \{1, 2\}$, H_{v_i} consists of the single edge $D_{i-1}D_i$. H_{v_1} is the path $D_n D D_1$, H_{v_2} is the path DD_1D_2, and H_x is the graph on the single vertex D. Moreover, H_v is the path $D_1 D_2 \dots D_n$. All in all, G is a 1-acyclic clique graph. ■

To see that G has no 2-acyclic clique graph, one follows the argument presented in the proof of Claim 1 to show that a 1-acyclic clique graph H of G contains all but one edge of $\mathbf{K}(W_n)$. Thus, the union of H_v and either H_{v_1} or H_{v_2} forms a cycle, and thus H is not 2-acyclic. □

The next lemma yields that the spanning subgraph H of $\mathbf{K}(G)$ returned by Algorithm 16 is 2-acylic, if the parallel runs of the algorithm make the same decision about whether to contain an edge.

Lemma 3.10.8. *Let G be a (possibly infinite) graph and let H be a spanning subgraph of $\mathbf{K}(G)$. Then H is a 2-acyclic clique graph of G if and only if $H[K_G(v)]$ is a clique tree of $B_{3/2}(v)$ for all vertices v of G.*

Proof. Set $H_v := H[K_G(v)]$ for every vertex v of G. By Lemma 3.3.2, $H_u \cap H_v = H_v[K_{B_{3/2}(v)}(u)]$ for every two vertices u, v of G.

First, assume that H is a 2-acyclic clique graph. Fix a vertex v of G. Consider a vertex u of $B_{3/2}(v)$. It remains to show that $H_v[K_{B_{3/2}(v)}(u)]$ is connected. Since u is a vertex of $B_{3/2}(v)$, there is a maximal clique of G containing u and v. Thus, by Lemma 3.7.3, $H_u \cap H_v = H_v[K_{B_{3/2}(v)}(u)]$ is connected.

Conversely, assume that H_v is a clique tree of $B_{3/2}(v)$ for all vertices v of G. In particular, H_v is connected for all $v \in V(G)$. Therefore, the spanning subgraph H of $\mathbf{K}(G)$ is a clique graph. It remains to show that H is 2-acyclic. As the H_v are trees by assumption, it suffices to show that $H_v \cap H_u$ is empty or connected for every two vertices u, v of G by Lemma 3.7.3.

For this, let u, v be two vertices such that H_v and H_u meet. The assumption on H that $H[K_G(v)]$ is a clique tree of $B_{3/2}(v)$ yields that $H_v[K_{B_{3/2}(v)}(u)] = H_u \cap H_v$ is connected. All in all, H is a 2-acyclic clique graph of G, as desired. □

Now, we combine the above observations to prove Theorem 15.

Proof of Theorem 15. We claim that the spanning subgraph H outputted by Algorithm 16 is indeed an r-acyclic clique graph of G. Afterwards we discuss its running time.
Correctness. As $K_G(v) = K(B_{3/2}(v))$ by Lemma 3.3.2 and thus $\mathbf{K}(G)[K_G(v)] = \mathbf{K}(B_{3/2}(v))$, the graph $H = \bigcup_{v \in V(G)} H^v$ is a spanning subgraph of $\mathbf{K}(G) = \bigcup_{v \in V(G)} \mathbf{K}(G)[K_G(v)] = \bigcup_{v \in V(G)} \mathbf{K}(B_{3/2}(v))$. Note that H^v is a spanning tree of $H_v \coloneqq H[K_G(v)]$, as H^v is obtained by running Kruskal's algorithm to obtain a maximum weight spanning tree of $\mathbf{K}(B_{3/2}(v))$. In particular, H_v is connected. Therefore, H is a clique graph.

It remains to show that H is r-acyclic. As G is r-locally chordal, it suffices to show that H is 3-acyclic by Theorem 11. As, by Lemma 3.10.6, a clique graph is 2-acyclic if and only if it is 3-acyclic, Lemma 3.10.8 ensures that it remains to ensure that $H_v = H[K_G(v)]$ is a clique tree of $B_{3/2}(v)$. As H^v is obtained by running Kruskal's algorithm on $\mathbf{K}(B_{3/2}(v))$, H^v is a clique tree of $B_{3/2}(v)$ by Theorem 3.10.5. Thus, we need to show that the subgraph H_v is indeed H^v. This follows from the following.

Claim 1. *For every two vertices u, v of G, the graphs H^u and H^v agree on $K_G(u) \cap K_G(v)$.*

Proof. Order the edges of $\mathbf{K}(G)$ in order $e_1, \dots, e_m$ of decreasing weight where we decide ties of two edges by the lexicographic order on the union of their respective ends. Note that, for every vertex v of G, the order $e^v_1, \dots, e^v_{m(v)}$ is the order of the edges of $\mathbf{K}(B_{3/2}(v)) = \mathbf{K}(G)[K_G(v)]$ induced by $e_1, \dots, e_m$. For a vertex v of G and an index $i = 1, \dots, m$, let $I(i, v)$ be the maximum $j = 1, \dots, m(v)$ such that $e^v_1, \dots, e^v_j$ is a subsequence of $e_1, \dots, e_i$.

Let u and v be two vertices of G. If $K_G(u) \cap K_G(v)$ is empty, then the claim holds. Hence, we may assume that $K_G(u)$ and $K_G(v)$ meet.

We proceed by induction on $i = 0, 1, \dots, m$ to show that $H^u_{I(i,u)}$ and $H^v_{I(i,v)}$ agree on $K_G(u) \cap K_G(v)$. The base case $i = 0$ holds by definition of H^u_0 and H^v_0. Suppose that $H^u_{I(i,u)}$ and $H^v_{I(i,v)}$ agree on $K_G(u) \cap K_G(v)$ for some $i = 0, \dots, m-1$. As we are otherwise done, we assume that e_{i+1} is an edge of $\mathbf{K}(G)[K_G(u) \cap K_G(v)]$, i.e. $e_{i+1} = K_1K_2$ with $u, v \in K_j$ for $j = 1, 2$. By symmetry, it suffices to show that if $H^u_{I(i,u)} + e_{i+1}$ contains a cycle C, then $H^v_{I(i,v)} + e_{i+1}$ contains the cycle C, as well. Note that the K_1–K_2 path $P \coloneqq C - e_{i+1}$ is in

$H^u_{I(i,u)}$. If $P \subseteq \mathbf{K}(G)[K_G(u) \cap K_G(v)]$, then P is in $H^v_{I(i,v)}$ by induction assumption, and thus $C \subseteq H^v_{I(i,v)} + e_{i+1}$.

In the following we show that $P \subseteq \mathbf{K}(G)[K_G(u) \cap K_G(v)]$. Since $H^u_{I(i,u)}$ is by definition a subgraph of the tree H^u, the unique K_1–K_2 path in the spanning tree H^u of $\mathbf{K}(G)[K_G(u)]$ is P. Recall that H^u is even a clique tree of $B_{3/2}(u)$ by Theorem 3.10.5. Therefore,

$$H^u_v \coloneqq H^u[K_{B_{3/2}(u)}(v)] = H^u[K_G(u) \cap K_G(v)] \subseteq \mathbf{K}(G)[K_G(u) \cap K_G(v)]$$

is a subtree of H^u, which contains the nodes K_1 and K_2. In particular, the unique K_1–K_2 path in H^u_v agrees with the unique K_1–K_2 path P in H^u. So, $P \subseteq \mathbf{K}(G)[K_G(u) \cap K_G(v)]$, as desired. ■

Running time. Let v be a vertex of G. As $B_{3/2}(v)$ contains $d(v)+1$ many vertices and $d(v)$ is at most the maximum degree Δ, Theorem 3.10.5 yields that the algorithm on processor v has running time $O(\Delta^2 \log \Delta)$. □

3.10.2 Clique graphs as maximum weight spanning subgraphs

Let G be a graph. A spanning subgraph H of $\mathbf{K}(G)$ is *r-acyclic* if $(H_v \mid v \in V(G))$ for $H_v \coloneqq H[K_G(v)]$ is an r-acyclic family of subgraphs of H. A *maximum weight r-acyclic spanning subgraph of* $\mathbf{K}(G)$ is an r-acyclic spanning subgraph H of $\mathbf{K}(G)$ that has maximum total weight among all r-acyclic spanning subgraphs of $\mathbf{K}(G)$. It turns out that, just as clique trees are precisely maximum weight spanning trees of $\mathbf{K}(G)$ (Theorem 3.10.1), r-acyclic clique graphs are precisely the maximum weight r-acyclic spanning subgraphs of $\mathbf{K}(G)$. Specifically, we prove:

Theorem 17. *Let G be a finite r-locally chordal graph for an integer $r \geqslant 3$. Let H be a subgraph of $\mathbf{K}(G)$. Then the following are equivalent:*

(i) *H is an r-acyclic clique graph of G.*

(ii) *H is a maximum weight 2-acyclic spanning subgraph of $\mathbf{K}(G)$.*

(iii) *H is a 2-acyclic spanning subgraph of $\mathbf{K}(G)$ of weight $W(G) \coloneqq (\sum_{K \in K(G)} |K|) - |V(G)|$.*

First, we need the following lemma.

Lemma 3.10.9. *Let G be a finite graph. Suppose that H is a 1-acyclic spanning subgraph of $\mathbf{K}(G)$. Then the total weight of H is at most $W(G) = (\sum_{K \in K(G)} |K|) - |V(G)|$. Moreover, equality holds if and only if H is a clique graph of G.*

Proof. Recall that an acyclic graph T on n vertices has at most $n-1$ edges. Using this, the desired inequality follows from double counting the total weight $w(H)$:

$$\begin{aligned}\sum_{K_1K_2 \in E(H)} |K_1 \cap K_2| &= \sum_{v \in V(G)} |E(H[K_G(v)])| \\ &\leqslant \sum_{v \in V(G)} (|K_G(v)| - 1) \qquad = \left(\sum_{K \in K(G)} |K|\right) - |V(G)| = W(G).\end{aligned}$$

Recall that a spanning subgraph H of $\mathbf{K}(G)$ is a clique graph of G if and only if all $H[K_G(v)]$ are connected Moreover, it is well-known that an acyclic graph T on n vertices is connected if and only if T has $n-1$ edges (see e.g. [Die25, Corollary 1.5.2]). Thus, equality holds if and only if $H[K_G(v)]$ is connected for every vertex v of G. □

Proof of Theorem 17. First, consider an r-acyclic clique graph H of G. By Lemma 3.10.8, the induced subgraph $H[K_G(v)]$ are in particular trees for every vertex v of G. Thus, H has total weight $w(G)$ by Lemma 3.10.9.

Now, let H be a maximum weight 2-acyclic spanning H subgraph of $\mathbf{K}(G)$. As G is r-locally chordal, G has r-acyclic clique graphs by Theorem 9. Thus, the above shows that the weight of H is at least $w(G)$. As H is in particular 1-acyclic, Lemma 3.10.9 yields that the weight of H is at most $w(G)$, and thus precisely $w(G)$. Then the moreover-part of Lemma 3.10.9 yields that H is a clique graph. So, H is a 2-acylic clique graph, and thus H is 3-acyclic by Lemma 3.10.6. Since G is r-locally chordal, it follows from Theorem 11 that H is even r-acyclic.

All in all, this proved the equivalence of (i), (ii), and (iii). □

3.11 Application: Locally interval & circular arc graphs

In this section we present joint work with Sandra Albrechtsen, Nathan Bowler and Jana Nickel. We discuss an application of our characterization of locally chordal graphs to a subclass of finite chordal graphs, the interval graphs. A countable graph is an *interval graph* if it has a region representation $v \mapsto P_v$ over a (possibly infinite) path P.[13] We refer to a region representation over a path as an *interval representation*. Since paths are trees, interval graphs are a subclass of chordal graphs by Proposition 3.4.1. Their characterization by forbidden induced subgraphs is as follows:

Theorem 3.11.1 (Boland & Lekkeikerker [LB62], II)**.** *Let G be a countable[14] graph. Then the following are equivalent:*

(i) *G is interval.*

(ii) *G is chordal and contains none of the graphs depicted in Figure 3.4 as an induced subgraph.*

Analogously to locally chordal graphs, a graph G is *r-locally interval* for an integer $r \geqslant 0$ if every $r/2$-balls $B_G(v, r/2)$ with $v \in V(G)$ is interval. As a sanity check, we first show that every interval graph is indeed locally interval.

Lemma 3.11.2. *Every countable interval graph is r-locally interval for every integer $r \geqslant 0$.*

[13] Interval graphs are often equivalently defined as intersection graphs of intervals on a line.

[14] Even though Boland and Lekkeikerker prove Theorem 3.11.1 only for finite graphs, a standard compactness argument via e.g. König's infinity lemma (see [Die25, Lemma 8.1.2] for its statement) shows that a countable graph is interval if and only if all its finite induced subgraphs are interval.

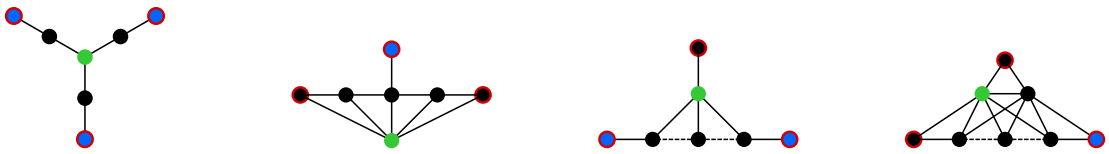

(a) Long claw (b) Whipping top (c) † (d) ‡

Figure 3.4: Chordal minimal forbidden induced subgraphs of interval graphs.

Proof. Let G be a locally finite interval graph, and $r \geqslant 0$ an integer. In particular, G is chordal, and thus G is r-locally chordal by Lemma 3.2.2. Suppose for a contradiction that an $r/2$-ball B in G around a vertex v contains a graph F depicted in Figure 3.4 as induced subgraph. If r is odd, then B is an induced subgraph of G, and thus the non-interval graph F is also an induced subgraph of G. Therefore, assume that r is even, and set $d := r/2$.

Since F is not an induced subgraph of G, some two vertices of F have to lie in $N^d(v)$. By Lemma 3.2.3, for every vertex $u \in N^d(v)$, its neighborhood $N_B(u)$ is a clique in B. Thus, only the red circled vertices in Figure 3.4 may be contained in $N^d(v)$. It is a simple check that every potential edge of G between some pair of red circled vertices yields an induced cycle of length at least 4 in G, which contradicts that G is chordal. □

What kind of graphs might be locally interval? Since every interval graph has a region representation over a path, it seems sensible to look for graphs that have a region representation over a graph that is "locally a path." Since paths have maximum degree two, no graph with a vertex of degree at least three can be "locally a path." Therefore, a good guess seems to be that locally interval graphs are those graphs which have region representations over a cycle. These form a well-known graph class, the *circular-arc* graphs.

A finite graph G is *circular-arc* if it has a region representation $v \mapsto C_v$ over a cycle C such that the C_v are paths (equivalently: $v \mapsto C_v$ is 1-acyclic). We refer to 1-acyclic region representations over a cycle as *circular-arc representations*. Analogously to region intersection graphs, we call a graph G an *r-acyclic* circular-arc graph if G has an r-acyclic circular-arc representation. Note that (∞-)acyclic circular-arc graphs are precisely interval graphs.

With these definitions, the following statement, conjectured by Paweł Rzążewski in private communication, seems morally true:

Conjecture 3.11.3 (Rzążewski 2024)**.** *Locally interval graphs are exactly circular-arc graphs.*

It turns out that Conjecture 3.11.3 is false for $r \leqslant 3$, but true for $r \geqslant 4$. Even in the case that $r = 3$, however, we will be able to use our tools to give a characterization of 3-locally interval graphs. First, we prove that for $r \geqslant 4$, the r-locally interval graphs are precisely the r-acyclic circular-arc graphs.

In the literature, 2-acyclic circular-arc representation are called *normal*. These types of circular-arc representations have often been considered when being additionally *Helly*. By Lemma 3.4.6, the 3-acyclic circular-arc representations are precisely the normal (equivalently:

2-acyclic) Helly circular-arc representations. Cao, Grippo and Safe characterized the 3-acyclic circular-arc graphs by forbidden induced subgraphs.

Theorem 3.11.4 (Cao, Grippo & Safe [CGS17], Theorem 1)**.** *Let G be a finite graph. Then the following are equivalent:*

(i) *G is 3-acyclic (equivalently: normal Helly) circular arc.*

(ii) *G is wheel-free (equivalently: 3-locally chordal) and contains none of the graphs depicted in Figures 3.4 and 3.5 as an induced subgraph.*

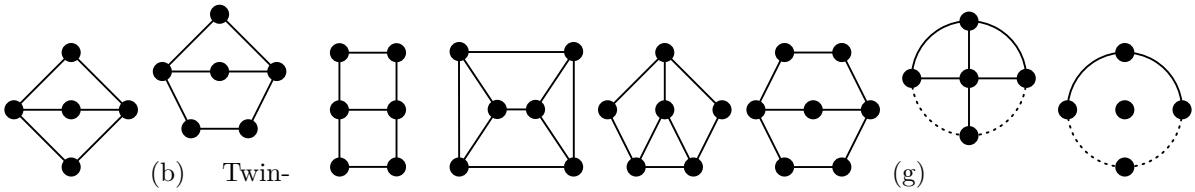

(a) $K_{2,3}$ (b) Twin-C^5 (c) Domino (d) $\bar{C^6}$ (e) FIS-1 (f) FIS-2 (g) Wheel W_n (h) C_n^*

Figure 3.5: Non-chordal minimal forbidden induced subgraphs of 3-acyclic circularc arc graphs, $n \geqslant 4$.

We can now prove the following:

Theorem 18. *Let G be a finite connected graph and $r \geqslant 4$ an integer. Then the following are equivalent:*

(i) *G is r-locally interval.*

(ii) *The r-local cover G_r of G is interval.*

(iii) *G is r-chordal, wheel-free, and contains none of the graphs depicted in Figure 3.4 as an induced subgraph. If $r = 4$, then G also contains none of the graphs depicted in Figure 3.6 as an induced subgraph.*

(iv) *G is r-acyclic circular arc.*

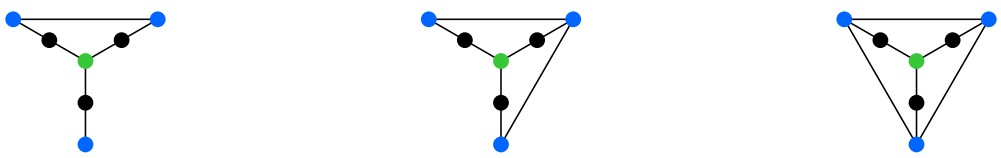

(a) Long Claw plus one edge (b) Long Claw plus two edges (c) Long Claw plus three edges

Figure 3.6: Additional forbidden induced subgraphs for r-acyclic circular arc graphs in the case $r = 4$.

Proof. We prove (ii)→(i)→(iii)→(ii), (iii)→(iv) and (iv)→(iii).

Claim 1. (ii) *implies* (i).

Proof. By (ii), the r-local cover G_r is interval, and thus Lemma 3.11.2 yields that G_r is r-locally interval. As p_r preserves the $r/2$-balls, it follows that G is r-locally interval, as desired. ■

Claim 2. *For every graph F depicted in Figure 3.4, there is a vertex v_F of F such that $F \subseteq B_F(v_F, 4/2)$. For every graph F depicted in Figure 3.6, there is a vertex v_F such that $F \cap B_F(v_F, 4/2)$ is the long claw.*

Proof. In both cases, let v_F be the green vertex in the respective figure. The blue vertices are the ones with distance 2 to v_F. ■

Claim 3. (i) *implies* (iii).

Proof. Since G is r-locally interval, it holds that G is r-locally chordal. Thus, by Theorem 6, G is wheel-free and r-chordal.

Suppose for a contradiction that G contains a graph from Figure 3.4 as an induced subgraph F. Then, F is contained in $B_F(v_F, 4/2) \subseteq B_G(v_F, 4/2)$ by Claim 2. Since F is an induced subgraph of G and $r \geqslant 4$, the graph F is an induced subgraph of the interval graph $B_G(v_F, r/2)$, which is a contradiction to Theorem 3.11.1. This completes the proof of the claim for $r \geqslant 5$.

Assume $r = 4$. Suppose for a contradiction that G contains a graph from Figure 3.6 as an induced subgraph F. Then, there is a vertex v_F of F such that $F \cap B_F(v_F, 4/2)$ is the long claw depicted in Figure 3.4. Since F is an induced subgraph of G and v_F is a vertex of G, we have $F \cap B_F(v_F, 4/2) = F \cap B_G(v_F, 4/2)$, and thus the long claw is an induced subgraph of $B_G(v_F, 4/2)$, which contradicts Theorem 3.11.1. This completes the proof of the claim for $r = 4$. ■

Claim 4. (iii) *implies* (ii).

Proof. By Theorem 3.11.1, it suffices to show that the r-local cover G_r has no graph F depicted in Figure 3.4 as an induced subgraph. Suppose for a contradiction that G_r contains a graph depicted in Figure 3.4 as an induced subgraph F. Then F is contained in $B_F(v_F, 4/2) \subseteq B_{G_r}(v_F, 4/2)$ by Claim 2.

First, assume that $r \geqslant 5$. Since F is an induced subgraph of G_r and $r \geqslant 5$, the graph F is an induced subgraph of the subgraph $B_{G_r}(v_F, 5/2)$ of G_r. As p_r preserves $r/2$-balls and $r \geqslant 5$, the graph F is also an induced subgraph of the induced subgraph $B_G(p_r(v_F), 5/2)$ of G, so F is an induced subgraph of G, which contradicts that F is not an induced subgraph of G by (iii).

Now, assume that $r = 4$. Since F is an induced subgraph of G_r and $r = 4$, the graph F is an induced subgraph of the subgraph $B_{G_r}(v_F, 4/2)$ of G_r. As p_r preserves $r/2$-balls and $r = 4$, the graph F is also an induced subgraph of $B_G(p_r(v_F), 4/2)$. Then only the blue vertices of F in Figure 3.4 have distance 2 to $p_r(v_F)$ in G. If F is not the long claw, then it is a simple check that every potential edge of G between some pair of the blue vertices yields an induced cycle of length 4 in G, which contradicts that G is in particular 4-chordal by (iii). Assume that F is the long claw. Then the addition of the potential edge of G between pairs of blue vertices

in Figure 3.4a gives one of the two graphs depicted in Figure 3.6, which then is an induced subgraph of G, contradicting (iii). ■

Claim 5. (iii) *implies* (iv).

Proof. First, we note that it suffices to show that G has none of the graphs depicted in Figure 3.5 as induced subgraphs. Suppose G forbids each graph in Figure 3.5. By (iii), G forbids each graph in Figure 3.4. Then G is a 3-acyclic circular-arc graph by Theorem 3.11.4. As G is r-locally chordal by (iii) and Theorem 6, it follows that G is in fact r-acyclic circular-arc by Theorem 11. Therefore, it remains to show that G forbids each graph in Figure 3.5 as induced subgraphs.

Every graph F depicted in Figure 3.5 but 3.5f contains an induced cycle of length 4, and thus is not an induced subgraph of G by (iii). It remains to show that the graph F depicted in Figure 3.5f is not an induced subgraph of G. Since F contains an induced cycle of length 5, it follows from (iii) that G excludes F as an induced subgraph when $r \geqslant 5$. If $r = 4$, then the induced subgraph F of G is precisely the graph depicted in Figure 3.6b, which is a contradiction to (iii). ■

Claim 6. (iv) *implies* (iii).

Proof. By Theorem 3.11.4 and $r \geqslant 4 > 3$, a r-acyclic circular arc graph G contains none of the graphs depicted in Figure 3.4 as induced subgraphs. This completes the proof for $r \geqslant 5$. For $r = 4$, it suffices to observe that none of the graphs in Figure 3.6 is circular arc. ■

This completes the proof. □

Now, we revisit the case that $r \leqslant 3$. For $r \leqslant 2$, Conjecture 3.11.3 is trivially wrong: every graph is 0-, 1-, and 2-locally interval, since the relevant balls are stars and stars are interval graphs, but there are graphs, e.g. the long claw depicted in Figure 3.4a, which are not circular arc.

For $r = 3$, Conjecture 3.11.3 still fails for the same reason: the long claw is still 3-locally interval but not a circular-arc graph. However, every 3-acyclic circular-arc graph is 3-locally interval by definition. One may characterize 3-locally interval graphs by forbidden induced subgraphs, as follows.

Proposition 3.11.5. *Let G be a finite graph. Then the following are equivalent:*

(i) *G is 3-locally interval.*

(ii) *G is wheel-free (equivalently: 3-locally chordal) and contains none of the graphs depicted in Figure 3.4 with one additional vertex that is complete to the whole graph as an induced subgraph.*

Proof. Every 3-locally interval graph is 3-locally chordal by Proposition 3.4.1. Conversely, let G be a 3-locally chordal graph such that G contains none of the graphs depicted in Figure 3.4

with an added apex vertex as induced subgraph. Consider a vertex v of G and the chordal 3/2-ball $B \coloneqq B_{3/2}(v)$ around v. By Theorem 3.11.1, it suffices to show that B contains none of the graphs depicted in Figure 3.4 as an induced subgraph. Suppose for a contradiction that B contains a graph from Figure 3.4 as an induced subgraph B'. Since v is complete to all other vertices in B but no vertex in the graphs depicted in Figure 3.4 is complete to all other vertices in that graph, v is not a vertex of B'. Thus, the subgraph induced on the vertex set of B' together with v is one of the induced subgraphs forbidden in G. □

3.12 Future work

In this section, we detail several directions for future work about locally chordal graphs.

3.12.1 Bounding the number of maximal cliques

One property of finite chordal graphs is that they have at most n maximal cliques (Proposition 3.10.4), where n is the number of their vertices. We can thus easily deduce that every r-locally chordal graph G for $r \geqslant 3$ has at most n^2 maximal cliques by observing that $B_{3/2}(v)$ is chordal for every $v \in V(G)$ and that every maximal clique of G lives in $B_{3/2}(v)$ for some v. However, we suspect that a better bound should hold for the number of maximal cliques in r-locally chordal graphs, where the bound depends on r and approaches n as r tends to infinity. Specifically, we conjecture the following:

Conjecture 3.12.1. *For every $\epsilon > 0$, there exists an integer $r \geqslant 0$ such that the number of maximal cliques in finite r-locally chordal graphs G is $O(|V(G)|^{1+\varepsilon})$.*

An extremal case of r-locally chordal graphs are those whose $r/2$-balls are tree, i.e. precisely the graphs of girth $> r$. Then Conjecture 3.12.1 holds, as the number of maximal cliques coincides with the number of edges in connected graphs G of girth $> r$, which is at most $n^{1+\frac{2}{r-2}}$ if the graph G has n vertices and at least $2n$ edges [Bol78, §III].

3.12.2 Algorithmic applications

One benefit of structural decompositions like tree-decompositions is that they enable the design of algorithms for graphs with "good" decompositions. Indeed, tree-decompositions play a large role in algorithmic graph theory. In Algorithm 16, we give an efficient algorithm to compute r-acyclic graph-decompositions into cliques (equivalently: r-acyclic clique graphs) of r-locally chordal graphs, which are structural descriptions that generalize tree-decomposition into cliques (equivalently: clique trees). Therefore, a very interesting follow-up question is whether these representations can be used in the design of efficient algorithms for locally chordal graphs.

One of the most successful algorithmic applications of chordal graphs is to the MAXIMUM INDEPENDENT SET problem. The MAXIMUM INDEPENDENT SET problem asks, given a graph G (and possibly a weight function $w : V(G) \to \mathbb{N}$ of the vertices), to return a maximum size (or maximum weight) independent set of G. MAXIMUM INDEPENDENT SET is a classic NP-hard

problem in general [Kar72], and is even hard to approximate in general [Hås99]. However, there is a linear-time algorithm to solve Maximum independent set in chordal graphs [Gav72]. One can also solve MIS in chordal graphs using standard dynamic programming on a tree-decomposition into cliques, which does not require that the chordal graphs have bounded treewidth, because independent sets intersect cliques, and thus bags of tree-decompositions, in at most one vertex. The nice interplay between chordal graphs and the Maximum Independent Set problem has inspired several other approaches to and perspectives on the Maximum Independent Set problem, for example the *PMC method* [FV10; LVV14] and the width-parameter *tree-independence number* [DMŠ24].

Since chordal graphs and Maximum Independent Set are such a fruitful pair, the Maximum Independent Set problem is a promising potential algorithmic application of locally chordal graphs. Our decompositions likewise have the property that they are into cliques, so a solution to Maximum independent set intersects each bag in at most one vertex. Since r-locally chordal graphs include, in particular, all graphs of girth greater than r, the class of graphs of high girth witnesses that Maximum Independent Set is NP-hard even in r-locally chordal graphs [Pol74]. However, we hope for an affirmative answer to the following.

Question 3.12.2. Is there a good approximation algorithm for Maximum Independent Set in r-locally chordal graphs?

There are decent approximation algorithms for Maximum Independent Set in high-girth graphs, e.g. [Mur92]. We hope that these algorithms can be modified to give approximation algorithms in locally chordal graphs given their r-acyclic graph-decompositions $(H, \mathcal{V})$ into cliques, by using the high-girth model graph H (on which the high-girth algorithms can be leveraged) and the fact that all bags are cliques (which intersect independent sets in at most one vertex).

3.12.3 Stronger Helly properties

Let H be a graph and let $\mathbf{H}$ be a family of subgraphs of H. The family $\mathbf{H}$ has the *Helly-Erdős-Pósa property* if, for every $k \geqslant 2$, every finite subfamily $\mathbf{H}'$ of $\mathbf{H}$ contains either k pairwise disjoint sets or there is a vertex-subset of H of size at most $k-1$ that meets every subgraph of $\mathbf{H}'$. Observe that when $k = 2$ this corresponds to the *Helly property* defined in this chapter, so the Helly-Erdős-Pósa property can be viewed as a generalization of the Helly property.

It is known that if T is a tree and $\mathbf{T}$ is a family of subtrees of T, then $\mathbf{T}$ satisfies the Helly-Erdős-Pósa property, see [GL70]. Therefore, if G is chordal, then the families corresponding to acyclic region representations of G (i.e. if G is connected, the subtree representations of G) satisfy the Helly-Erdős-Pósa property.

For r-locally chordal graphs, we proved that the families corresponding to their r-acyclic region representations satisfy the Helly property. A natural extension of this result is thus to ask whether locally chordal graphs, like chordal graphs, also satisfy the Helly-Erdős-Pósa property:

Question 3.12.3. Let $r \geqslant 3$ be an integer and G an r-locally chordal graph. For an r-acyclic region representation $v \mapsto H_v$ of G, does the family $(H_v \mid v \in V(G))$ satisfy the Helly-Erdős-Pósa property?

We thank Piotr Micek for suggesting Question 3.12.3 in private communication.

3.12.4 Generalizing treewidth

Following the correspondence of the tree-decompositions of a graph and its chordal completions, a final intriguing direction for further study is to exploit a similiar correspondence between r-acyclic graph-decompositions and and r-locally chordal completions to generalize the notion of treewidth. Inspired by our work on locally chordal graphs, we introduce together with Abrishami and Diestel a new (parameterized by r, as in r-locally chordal) width parameter called *r-acyclic width*. They will soon share the definition of r-acyclic width, as well as several observations and ideas for further study, in a preprint [ADK].

3.13 Canonical subtree representations of chordal graphs

The purpose of this section is to prove Theorem 8 and Theorem 3.3.7. We warn the reader that the techniques used in this section heavily rely on modern separation theory and are very different to the rest of this chapter. Let us begin by reviewing relevant background information.

3.13.1 Background: Separations

Let us recall some definitions regarding separations. Let G be a graph. A separation $\{A, B\}$ of G is *proper* if both $A \setminus B$ and $B \setminus A$ are non-empty. A separation $\{A, B\}$ of G is *tight* if each proper side contains the vertex set of a tight component of $G - (A \cap B)$. We remark that a separation $\{A, B\}$ is tight if and only if its separator $A \cap B$ is an inclusion-minimal a–b separator for some $a \in A \setminus B$ and $b \in B \setminus A$. The *order* of $\{A, B\}$ is the size $k := |A \cap B|$ of its separator. We also say that $\{A, B\}$ is a *k-separation* of G. Note that the automorphisms φ of G naturally act on the separations $s = \{A, B\}$ of G by mapping them to $\varphi(s) := \{\varphi(A), \varphi(B)\}$. For a subgroup Γ of $\mathrm{Aut}(G)$, a set S of separations of G is *Γ-invariant* if $\varphi(s) \in S$ for every $s \in S$ and every $\varphi \in \Gamma$.

3.13.2 Background: Interplay of tree-decompositions and separations

We refer to (A, B) and (B, A) as the *orientations* of the separation $\{A, B\}$ or just as *oriented separations* of G. Two separations $\{A, B\}$ and $\{C, D\}$ are *nested* if there are orientations of $\{A, B\}$ and $\{C, D\}$, say (A, B) and (C, D), such that $A \subseteq C$ and $B \supseteq D$. A set of separations is *nested* if its elements are pairwise nested.

Let $(T, \mathcal{V})$ be a tree-decomposition of a graph G. Recall that, by (T2'), every edge f of the decomposition tree T induces a separation s^f of G. A tree-decomposition is *regular* if all its induced separations are proper.

Lemma 3.13.1 ([DJKK22], Lemma 3.7)**.** *Let $(T, \mathcal{V})$ be a tree-decomposition of a graph G. If $(T, \mathcal{V})$ is regular, then for every automorphism γ of G there exists at most one automorphism φ of T such that $\gamma(V_t) = V_{\varphi(t)}$. In particular, if additionally $(T, \mathcal{V})$ is Γ-canonical for a subgroup Γ of $\mathrm{Aut}(G)$, then there is a unique action of Γ on T which commutes with $t \mapsto V_t$.*

Let G be a graph. A tree-decomposition $(T, \mathcal{V})$ of G is *point-finite* if T_v is finite for every vertex v of G. A set S of separations of G is *point-finite* if every vertex of G appears in the separator of at most finitely many separations in S. Given a nested set N of proper separations of a graph, let $\mathcal{T}(N)$ be defined as in [CJKK25, Construction 2.3].

Lemma 3.13.2 ([EKT22], Lemma 2.7)**.** *Let N be a nested set of proper separations of a connected graph G. If N is point-finite, then there is a tree decomposition $\mathcal{T}(N)$ of G such that $f \mapsto s^f$ is a bijection from the edges of T to N. In particular, $\mathcal{T}(N)$ is regular and point-finite. Moreover, if N is $\Gamma(G)$-invariant, then $\mathcal{T}(N)$ is Γ-canonical.*

3.13.3 Canonical tree-decompositions into cliques

Given two vertex-subsets X, Y of a graph G, a separation $\{A, B\}$ of G *distinguishes* X and Y if, up to renaming sides, $X \subseteq A$, X meets $A \setminus B$, $Y \subseteq B$ and Y meets $B \setminus A$. If $\{A, B\}$ has minimum order among the separations of G that separate X and Y, then $\{A, B\}$ distinguishes X and Y *efficiently*. A set S of separations of G *efficiently distinguishes* a set $\mathcal{X}$ of vertex-subsets of G if every two elements of $\mathcal{X}$ are efficiently distinguished by some $s \in S$.

Lemma 3.13.3. *Let X and Y be two distinct maximal cliques in a locally finite, chordal graph G. Then X and Y are distinguished by a separation of order less than the sizes of X and Y.*

Proof. If X lies in another component C of G than Y, then $\{C, V(G) \setminus C\}$ is as desired. So we assume otherwise; in particular, we may assume without loss of generality that G is connected. By Theorem 3.4.2, the chordal graph G admits a tree-decomposition $(T, \mathcal{V})$ into its maximal cliques. Consider the bags V_s, V_t with $V_s = X$ and $V_t = Y$. Let $f = ss'$ be the first edge of the unique s–t path in T. Since $(T, \mathcal{V})$ is into the maximal cliques of G, $V_f = V_s \cap V_{s'}$ is a proper subset of $X = V_s$ and $Y = V_t$. Thus, the separation s^f of G corresponding to f distinguishes X and Y and has order less than $|X|$ and $|Y|$. □

For a given locally finite, connected graph G, let $N(G)$ be defined as in [CJKK25, Construction 3.6].

Lemma 3.13.4. *Let G be a locally finite, connected chordal graph. Then $N(G)$ is an $\mathrm{Aut}(G)$-invariant nested set of tight separations of G that efficiently distinguishes every two distinct maximal cliques of G.*

Proof. The following proof uses the notion of 'bottleneck in G' from [CJKK25, § 3.1] as black box, in that it is only relevant to the proof that bottlenecks are sets of tight separations of G and, by Lemma 3.13.3, the set $\beta(K, K')$ of separations which efficiently distinguish some two distinct maximal cliques K, K' in G is a so-called bottleneck in G [CJKK25, Example 3.4].

Then [CJKK25, Theorem 3.7] yields that the nested set $N(G)$ of separations of G meets in particular every $\beta(K, K')$. Hence, $N(G)$ efficiently distinguishes the set $K(G)$ of all maximal cliques. The construction of $N(G)$ [CJKK25, Construction 3.6] is invariant under isomorphisms of the underlying graph and thus $N(G)$ is $\mathrm{Aut}(G)$-invariant. Moreover, all elements of $N(G)$ are tight, as they are contained in some so-called bottleneck, which by definition only contains tight separations. □

For experts, we remark that one may alternatively obtain a nested set like $N(G)$ by [EKT22, Theorem 1.3], as every maximal clique in a chordal graph yields a 'robust regular profile' and by invoking Lemma 3.13.3 one may prove that each two of them are 'distinguishable'. Unfortunately, the nested set from [DJKK22, Definition 5.2] does not work in general, as the 'tangles' induced by cliques K have in general only order $2/3 \cdot |K|$, which as a consequence may make the two tangles induced by two distinct maximal cliques indistinguishable.

Lemma 3.13.5. *Let G be a locally finite, connected chordal graph. Assume that N is a nested set of tight separations of G that efficiently distinguishes every two distinct maximal cliques in G. Then all separators of separations in N are cliques and N is point-finite.*

Proof. Since all separations in N are tight, so are their separators, which are hence cliques by Theorem 3.2.1. Since G is locally finite, every vertex of G is contained in only finitely many cliques, and since G is connected, there are only finitely many separations with the same finite separator. Thus, every vertex of G is contained in only finitely many separators of separations in N, so N is point-finite. □

Lemma 3.13.6. *[HL18, Lemma 14] If X is a tight separator of a (possibly infinite) chordal graph G, then there are two distinct vertices $u, v \in V(G) \setminus X$ such that X is a u–v separator and both u and v are complete to X.*

Lemma 3.13.7. *Let G be a connected, locally finite, chordal graph. Assume that N is a nested set of tight separations of G that efficiently distinguishes every two distinct maximal cliques. Then $\mathcal{T}(N)$ is a regular, point-finite tree-decomposition of G into cliques. Moreover, $\mathcal{T}(N)$ is canonical if N is $\mathrm{Aut}(G)$-invariant.*

Proof. By Lemma 3.13.5, N is point-finite and the separators of separations in N are cliques. As all separations in N are tight, they are also proper. So Lemma 3.13.2 ensures that $\mathcal{T} := \mathcal{T}(N)$ is a tree-decomposition $(T, \mathcal{V})$ such that the map $f \mapsto s^f$ is a bijection from $E(T)$ to N; in particular, $N = \{s^f \mid f \in E(T)\}$, and thus all adhesion sets of $(T, \mathcal{V})$ are cliques of G. The latter ensures that every part $G_t = G[V_t]$ of $\mathcal{T}$ is connected, since G is connected. Note that, by Lemma 3.13.2, $\mathcal{T}$ is regular and point-finite, and also $\mathcal{T}$ is canonical, if N is $\mathrm{Aut}(G)$-invariant.

We claim that $\mathcal{T}$ is into cliques. Suppose for a contradiction that some V_t is not a clique of G, i.e. there are two non-adjacent vertices $u, v \in V_t$. In particular, $V_t \setminus \{u, v\}$ separates u and v in the part $G_t = G[V_t]$ of $(T, \mathcal{V})$. Now let X be an inclusion-minimal u–v separator in G_t. Note that X is nonempty, as the part G_t of $\mathcal{T}$ is connected.

As G is chordal, so is its induced subgraph $G[V_t] = G_t$. Thus, Theorem 3.2.1 implies that X is a clique, and Lemma 3.13.6 ensures that there are two vertices $u', v' \in V_t \setminus X$ such that X is a u'–v' separator in G_t and both u' and v' are complete to X. Let K_1 and K_2 be two maximal cliques of G that contain $X \cup \{u'\}$ and $X \cup \{v'\}$, respectively. As the adhesion sets of $(T, \mathcal{V})$ are cliques, X is not only a u'–v' separator in the part G_t but also in G. In particular, X separates the cliques $K_1 \ni u'$ and $K_2 \ni v'$ in G. Note that every K_1–K_2 separator in G must contain their intersection $K_1 \cap K_2$, which contains at least X. Hence, the only K_1–K_2 separator of minimum size in G is indeed X. Since $N = \{s^f \mid f \in E(T)\}$ efficiently distinguishes every two distinct maximal cliques such as K_1 and K_2 by assumption, some separation s^f corresponding to an edge f of T separates K_1 and K_2 efficiently. Then its separator is X following the above analysis; in particular, s^f separates $u' \in K_1 \setminus X$ and $v' \in K_2 \setminus X$. But u', v' are in the common bag V_t and thus cannot be separated by a separation s^f corresponding to an edge f of T by the definition of s^f, which is a contradiction. □

Theorem 3.13.8. *Let G be a locally finite, connected chordal graph. Then $\mathcal{T}(N(G))$ is a point-finite, regular, canonical tree-decomposition of G into cliques.*

Proof. By Lemma 3.13.4, $N(G)$ is an $\mathrm{Aut}(G)$-invariant, nested set N of tight separations of G that efficiently distinguishes every two distinct maximal cliques of G. Then Lemma 3.13.7 yields that $\mathcal{T}(N(G))$ is a canonical tree-decomposition of G into cliques, which is regular and point-finite. □

Proof of Theorem 8. This follows immediately from Theorem 3.13.8 by choosing the tree-decomposition $\mathcal{T}$ to be $\mathcal{T}(N(G))$. □

3.13.4 Tree-decompositions of coverings into all maximal cliques

Next we prove Theorem 3.3.7.

In view of Example 3.3.8, it is unfortunately not the case that every (locally) finite chordal graph admits a canonical tree-decomposition into its maximal cliques. For our purposes, however, it is enough that the tree-decomposition be $\Gamma(p)$-*canonical*, and this turns out to be true. We need one more lemma before we are ready for the proof.

Recall that a covering $p\colon \hat{G} \to G$ is *normal* if the group $\Gamma(p)$ of deck transformations of p acts transitively on the fibres $p^{-1}(v)$ of p. Indeed, the desired result holds for all normal coverings of a graph, since the normal coverings yet satisfy the following weak version of Lemma 3.2.4 (iii).

Lemma 3.13.9. *Let G be a connected, locally finite graph and $p\colon \hat{G} \to G$ a normal covering, e.g. the r-local covering. Then for every clique K, the intersection $K \cap \gamma(K)$ is empty whenever $\gamma \in \Gamma(p) \setminus \{\mathrm{id}_{\hat{G}}\}$.*

Proof. As G is simple, it has no loops. Thus, the fibres of p are independent sets in $\hat{G}$. In particular, every two distinct vertices in a clique K of $\hat{G}$ are in distinct fibres of p. Thus, the

desired statement follows, as $\Gamma(p)$ always acts freely and, since P is normal, also transitively on the fibres of p. □

Now we can prove Theorem 3.3.7.

Proof of Theorem 3.3.7. Set $\Gamma \coloneqq \Gamma(p)$. Theorem 3.13.8 yields a tree-decomposition of G that is a point-finite, regular, Γ-canonical and also into cliques. From now on let $(T, \mathcal{V})$ be any tree-decomposition of $\hat{G}$ with these properties. We claim that if there is no edge $f = st$ of T such that $V_s \subseteq V_t$, then $(T, \mathcal{V})$ is into the maximal cliques of G, i.e.

(i) for every maximal clique K of $\hat{G}$, there is a unique node $t(K)$ of T with $K = V_{t(K)}$, and

(ii) every V_t is a maximal clique.

(i): First, we show that $t(K)$ exists. By Lemma 3.3.1, we may consider the subtree representation $v \mapsto T_v$ of G corresponding to $(T, \mathcal{V})$. Thus, $\bigcap_{v \in K} T_v$ is non-empty, since the T_v with $v \in K$ are pairwise intersecting because K is a clique and $v \mapsto T_v$ is Helly by Theorem 3.6.3. Let t be some node in the intersection $\bigcap_{v \in K} T_v$, i.e. $K \subseteq V_t$. As $(T, \mathcal{V})$ is into cliques and K is a maximal clique, V_t is indeed equal to K.

Secondly, we show that $t(K)$ is unique. Suppose for a contradiction that there are two distinct nodes t, s of T with $V_t, V_s = K$. Then (T2) ensures that all vertices t' on the t–s path in T satisfy $X \subseteq V_{t'}$. As $(T, \mathcal{V})$ is into cliques and X is a maximal clique, we have $X = V_{t'}$. Thus, every edge $f = t's'$ on the t–s path in T contradicts the assumption, as $V_{t'} = X = V_{s'}$.

(ii): Suppose for a contradiction that V_t is not a maximal clique. Let K be a maximal clique that contains V_t. Let $f = ts$ be the first edge on the t–$t(K)$ path in T. Then (T2) ensures that $V_t \subseteq V_s$, which contradicts our assumption.

Since $(T, \mathcal{V})$ is Γ-canonical and regular, the deck transformations of p act by Lemma 3.13.1 uniquely on T such that this action commutes with $t \mapsto V_t$. We claim that T is countable. Indeed, since $(T, \mathcal{V})$ is point-finite and the connected graph G is locally finite, T is the union $\bigcup_{v \in V(G)} T_v$ of countably many finite graphs, and thus countable. Hence, $E(T)$ countably many Γ-orbits $M_1, M_2, \dots$. Set $I_0 \coloneqq \emptyset$. We iteratively construct an inclusion-wise increasing sequence of $I_i \subseteq [i]$ such that the tree-decomposition $(T^i, \mathcal{V}^i)$ of $\hat{G}$ obtained by contracting $M^i \coloneqq \bigcup_{j \in I_i} M_j$ in T, i.e. set $T^i \coloneqq T/M^i$, set $V^i_t \coloneqq \bigcup_{s \in C} V_s$ if t is the contraction vertex given by some component C of $(V(T), M^i)$, and else $V^i_t \coloneqq V_t$, is into cliques, and we have $V^i_s \not\subseteq V^i_t$ and $V^i_s \not\supseteq V^i_t$ for every edge $f = st \in \bigcup_{j \in [i] \setminus I_i} M_j$. It is immediate from the definition that the tree-decompositions $(T^i, \mathcal{V}^i)$ are point-finite, regular, Γ-canonical, and the Γ-orbits of $E(T^i)$ are the M_j with $j \in [i] \setminus I_i$ and the $M_i, M_{i+1}, \dots$.

Assume that I_{i-1} is defined for some integer $i \geqslant 1$. Let $f_i = st$ be an edge in M_i. If $V_s \not\subseteq V_t$ and $V_s \not\supseteq V_t$, then $I_i \coloneqq I_{i-1}$ is as desired. Otherwise, we claim that $I_i \coloneqq I_{i-1} \cup \{i\}$ is as desired. Observe that the Γ-orbit M_i of f_i in $E(T^{i-1})$ is a matching in T^{i-1} by Lemma 3.13.9. The tree-decomposition $(T^i, \mathcal{V}^i)$ of $\hat{G}$ is (canonically isomorphic to) the one obtained from $(T^{i-1}, \mathcal{V}^{i-1})$ by contracting M_i in T^{i-1}. As only independent edges $f = st$ with, say, $V^{i-1}_s \subseteq V^{i-1}_t$ have been contracted, $(T^i, \mathcal{V}^i)$ is still into cliques.

Set $I \coloneqq \bigcup_{i\in\mathbb{N}} I_i$, and let $(T', \mathcal{V}')$ be the tree-decomposition of $\hat{G}$ obtained from $(T, \mathcal{V})$ by contracting $\bigcup_{j\in I} M_j$ in T. Again, $(T', \mathcal{V}')$ is by definition point-finite, regular, Γ-canonical, and we have $V'_s \not\subseteq V'_t$ and $V'_s \not\supseteq V'_t$ for every edge $f = st \in E(T')$. We note that every bag V'_t is a clique of G, as otherwise there would have been a bag $V^i_{t'} \subseteq V'_t$ which was not a clique as well, which contradicts the recursive construction. □

Corollary 3.13.10. *Let G be a connected, locally finite graph and $p\colon \hat{G} \to G$ a normal covering, e.g. the r-local covering. If $\hat{G}$ is chordal, then its group $\Gamma(p)$ of deck transformations is free.*

Proof. Set $\Gamma \coloneqq \Gamma(p)$. By Theorem 8, $\hat{G}$ admits a Γ-canonical tree-decomposition $(T, (V_t)_{t\in T})$ into cliques. Lemma 3.13.9 ensures that Γ acts freely on the decomposition tree T. Hence, it follows from [Löh17, Theorem 4.2.1] that Γ is free. □

3.14 Comparison: Canonical graph-decomposition via coverings

We conclude by highlighting the relationship between the graph-decompositions obtained in this chapter and the graph-decompositions developed and studied in [CJKK25], which is a slight adaptation of the one from [DJKK22, Theorem 1], as discussed in [CJKK25]. We warn the reader that this section assumes some familiarity with [CJKK25].

We start by reviewing the graph-decompositions from [CJKK25]. Let G be a locally finite, connected graph and r a positive integer. For $k \in \mathbb{N}$, we denote by $N^{\leqslant k}(G)$ the subset of $N(G)$ of all separations of order at most k. Then $\mathcal{T}^{\leqslant k} \coloneqq \mathcal{T}(N^{\leqslant k}(G_r))$ is a tree-decomposition by [CJKK25, Theorem 3.12]. Let $\mathcal{H}_r^{\leqslant k} \coloneqq \mathcal{H}_r^{\leqslant k}(G)$ be the graph-decomposition of G obtained from $\mathcal{T}(N^{\leqslant k}(G_r))$ by folding via p_r. In [CJKK25, Theorem 1], they discuss the computability of these graph-decompositions $\mathcal{H}_r^{\leqslant k}$ for finite graphs G.

Note that, for $k \leqslant k'$, one may obtain $\mathcal{H}_r^{\leqslant k}$ from $\mathcal{H}_r^{\leqslant k'}$ by contracting all edges hh' of the decomposition graph $H_r^{\leqslant k'}$ of $\mathcal{H}_r^{\leqslant k'}$ which are the orbit of an edge st of the decomposition tree $T^{\leqslant k'}$ of $\mathcal{T}^{\leqslant k'}$ that induce a separation of G of order $> k$ and merge the bags assigned to h and h'. Thus, we say that $\mathcal{H}_r^{\leqslant k'}$ *refines* $\mathcal{H}_r^{\leqslant k}$. If also $\mathcal{T}(G_r) \coloneqq \mathcal{T}(N(G_r))$ is a tree-decomposition then the graph-decomposition $\mathcal{H}_r(G)$ obtained from $\mathcal{T}(G_r)$ by folding via p_r is their *limit*.

Theorem 19. *Let G be a connected, locally finite graph and $r \geqslant 3$ an integer. The following are equivalent:*

(i) *G is r-locally chordal.*

(ii) *The graph-decomposition $\mathcal{H}_r(G)$ obtained from the canonical tree-decomposition $\mathcal{T}(G_r)$ by folding via p_r is well-defined and into cliques.*

Proof. Assume that G is r-locally chordal. Then its r-local cover G_r is chordal by Theorem 6. Thus, $\mathcal{T}(N(G))$ is into cliques by Theorem 3.13.8. Therefore, Lemma 3.3.6 ensures that $\mathcal{H}_r(G)$ is into cliques.

Conversely, assume that $\mathcal{H}_r(G)$ is well-defined and into cliques. Then Lemma 3.3.6 yields that $\mathcal{T}$ is into cliques. So, the r-local cover G_r of G is chordal by Proposition 3.4.1. Therefore, Theorem 6 ensures that G is r-locally chordal. □

PART II

Coarse graph theory

CHAPTER 4

A structural duality for path-decompositions into parts of small radius

This chapter is based on [ADEFJKW23], which is joint work with Sandra Albrechtsen, Reinhard Diestel, Ann-Kathrin Elm, Eva Fluck, Raphael W. Jacobs, and Paul Wollan. The here presented version has a completely rewritten introduction, which takes into account several developments occurring close or parallel to the publication of [ADEFJKW23].

4.1 Introduction

4.1.1 Fat minors and quasi-isometries

Following Gromov's ideas on coarse geometry [Gro93] into the realm of graphs, Georgakopoulos and Papasoglu [GP23] 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. For example, all locally finite Cayley graphs of a finitely generated group are quasi-isometric to each other (see e.g. [Löh17, Proposition 5.2.5]).

As their favorite problem of metric graph-theoretic flavour, Georgakopoulos and Papasoglu proposed the following conjecture, whose qualitative converse is immediate:

Conjecture 4.1.1 ([GP23], Conjecture 1.1)**.** *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 $X \in \mathcal{X}$.*

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 4.1.4 for the formal definition).

Georgakopoulos and Papasoglu [GP23] 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 [CDEHV08] yields the case $\mathcal{X} = \{K_{2,3}\}$ (quasi-isometric to an outerplanar graph). Fujiwara and Papasoglu [FP23] solved the $\mathcal{X} = \{K_4^-\}$-case (quasi-isometric to a cactus); see also [AJKW24] for a simpler proof. Moreover, Albrechtsen, Jacobs, Knappe and Wollan [AJKW24] solved the case $\mathcal{X} = \{K^4\}$. In contrast to these positive results, Davies, Hickingbotham, Illingworth and McCarty [JRFR24] showed that Conjecture 4.1.1 is false in general.

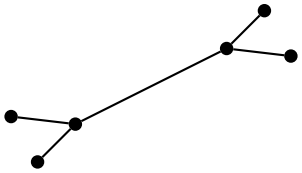

Figure 4.1: The 'wrench'/'spanner' graph W

In this chapter, we establish Conjecture 4.1.1 for three further small cases but in a stronger form:

Theorem 20. *Conjecture 4.1.1 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 4.1. In all these cases, Conjecture 4.1.1 even holds true when replacing 'fat minors' by 'quasi-geodesic topological minors'.*

We refer the reader to Theorems 21 to 23 for the respective functions f.

In Theorem 20, 'quasi-geodesic topological minors' are a metric version of topological minors, whose model in G is 'quasi-geodesically' embedded (see Section 4.1.4 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 4.3.14).

Our work on Theorem 20 was independent of Conjecture 4.1.1 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 4.1.1, through [GP23]. Our approach is not based on quasi-isometries, but on the notion of graph-decompositions, which we describe in what follows.

4.1.2 Graph-decompositions

Graph-decompositions [DJKK22] 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 [DJKK22] as well as the study of local separations [CJKK25] and of locally chordal graphs [AJKK25].

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 $\mathcal{H}$,

[1] Georgakopoulos and Papasoglu independently prove the more general case $\mathcal{X} = \{K_{1,m}\}$ in a second version of their paper [GP23].

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) [DK05; RS86b].[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 [DG07] and *tree-breadth* in [DK14a]. Similar to the classical tree-width notion [RS86a], several computationally hard problems such as the computation of the metric dimension of a graph [BFGR17] can be efficiently solved on graphs of small radial tree-width (for a summary, see [DG07, Section 1] or [Lei17]).

4.1.3 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'.

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 [BS24, p. 4.1]. Here, we extend their observation to arbitrary graphs H.

Proposition 4.1.2. *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.*

[2] By Diestel and Kühn [DK05, 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 [RS86b].

(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 4.3.3 for the detailed statement on the functions g and h_1, h_2.

The correspondence given by Proposition 4.1.2 paves the way for a new proof method towards Conjecture 4.1.1: 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 [AJKW24].

Proposition 4.1.2 in particular implies that we reached our initial goal from the beginning of Section 4.1.2 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.

4.1.4 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 4.1.1, 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 4.3.14). Hence, as fat minors form an obstruction to small radial width (Lemma 4.3.15), so do quasi-geodesic topological minors (Proposition 4.3.16).

Our main result, Theorem 20, 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 21 to 23, from which we then deduce Theorem 20.

Theorem 21 (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 $18k+1$.*

[3]Note that in metric spaces this property is often called $(c, 0)$-quasi-geodesic. In [BR22] it is called c-multiplicative.

Moreover, P is $(1, 18k + 2)$-quasi-isometric to G.

Theorem 22 (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 23 (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 20 given Theorems 21 to 23. 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 4.3.14 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 20 from Theorems 21 to 23 by applying the respective 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 4.3.8 yields the desired $f(K)$-quasi-isometry from G to H. □

4.1.5 Sketch of the proofs

Let us give a brief overview of the proof techniques for Theorems 21 to 23. For the proof of Theorem 21 (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 21 immediately generalizes to Theorem 22 (radial cycle-width). The proof of Theorem 23 (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.

[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 23 in Section 4.6 for details).

4.1.6 An open conjecture about quasi-isometries to trees

For the special case of $\mathcal{X} = \{K^3\}$, Georgakopoulos and Papasoglu answered Conjecture 4.1.1 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 [BS24] characterised the graphs quasi-isometric to a forest using a similar kind of obstruction. Our results Theorems 21 to 23 yield another natural guess for such an obstruction – long quasi-geodesic cycles:

Conjecture 4.1.3. *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 4.1.2 and [BS24, p. 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 $f(k)$. In Section 4.8 we show that the proof technique described in Section 4.1.5 for Theorems 21 and 22 does not work for Conjecture 4.1.3.

In Section 4.7, we show that the strengthening of Conjecture 4.1.1 for quasi-geodesic topological minors analogous to Conjecture 4.1.3 fails earlier than Conjecture 4.1.1 itself.

4.1.7 Structure of this chapter

Section 4.2 collects some basic definitions. In Section 4.3 we introduce the notion of graph-decompositions and their radial width and spread, and show how they are related to quasi-isometries, proving Proposition 4.1.2. We also prove in Section 4.3.4 that quasi-geodesic topological minors yield fat minors and are an obstruction to small radial width. The following three Sections 4.4–4.6 contain the proofs of Theorems 21 to 23, respectively. We finally discuss further problems resulting from our work in Sections 4.7 and 4.8.

4.2 Preliminaries

We mainly follow the notations from [Die25]. We recall that for two sets X, Y of vertices of a graph G an *X–Y path* in G is a path in G whose only vertex in X is its first vertex and whose only vertex in Y is its last vertex. Moreover, for a subgraph H of G an *H-path* in G is a non-trivial path in G which meets H precisely in its endvertices. For the remainder of the section, we briefly present the definitions we will require going forward. All graphs in this chapter are finite.

Let G be a graph. We write $\mathrm{d}_G(v, u)$ for the distance of the two vertices v and u in G. For two sets U and U' of vertices of G, we write $\mathrm{d}_G(U, U')$ for the minimum distance of two elements of U and U', respectively. If one of U or U' is just a singleton, then we may omit the braces. Further, if X is a subgraph of G, then we abbreviate $d_G(U, V(X))$ as $d_G(U, X)$ for notational simplicity.

Given a set X of vertices of G, the *r-ball (in G) around X*, denoted as $B_G(X, r)$, is the set of all vertices in G of distance at most r from X in G. If $X = \{v\}$ for some $v \in V(G)$, then we omit the braces, writing $B_G(v, r)$ for the r-ball (in G) around v.

Further, the *radius* $\mathrm{rad}(G)$ of G is the smallest number $k \in \mathbb{N}$ such that there exists some vertex $v \in G$ with $\mathrm{d}_G(v, w) \leqslant k$ for every vertex $w \in G$. Note that G has radius at most k if and only if there is some vertex $v \in G$ with $V(G) = B_G(v, k)$. If G is not connected, then we define its radius to be ∞. Additionally, if $U \subseteq V(G)$, then the *radius of U in G* is the smallest number $k \in \mathbb{N}$ such that there exists some vertex $v \in G$ with $U \subseteq B_G(v, k)$.

4.3 Graph-decompositions and quasi-isometries

In this section we first introduce the concept of graph-decompositions and their radial width and spread and collect some properties of these notions. We then study the connection between graph-decompositions of small radial width and spread and quasi-isometries to their decomposition graph, proving Proposition 4.1.2. Finally, we define quasi-geodesic topological minors and show that they yield fat minors, which indeed form obstruction to small radial width.

4.3.1 Graph-decompositions

Let us recall the notion of graph-decompositions:

Definition 4.3.1 (Graph-decomposition [DJKK22])**.** Let G and H be graphs and let $\mathcal{V} = (V_h)_{h \in H}$ be a family of sets V_h of vertices of G.[5] We call $(H, \mathcal{V})$ an *H-decomposition* of G, a *decomposition of G modelled on H*, or just a *graph-decomposition*, if

(H1) $\bigcup_{h \in H} G[V_h] = G$, and

(H2) for every vertex $v \in G$, the graph of $H_v := H[\{h \in H \mid v \in V_h\}]$ is connected.

The sets V_h are called the *bags* of this graph-decomposition, their induced subgraphs $G[V_h]$ are its *parts*, and the graph H is its *decomposition graph*. Whenever a graph-decomposition is introduced as $(H, \mathcal{V})$, we tacitly assume $\mathcal{V} = (V_h)_{h \in H}$. For a class $\mathcal{H}$ of graphs, we call $(H, \mathcal{V})$ an *$\mathcal{H}$-decomposition* if it is an H-decomposition with $H \in \mathcal{H}$.

A graph-decomposition $(H, (V_h)_{h \in H})$ of a graph G is *honest* if all its bags V_h are non-empty and $V_h \cap V_{h'} \neq \emptyset$ for every edge $hh' \in H$.

If G is connected and every bag of $(H, \mathcal{V})$ is non-empty, then it follows from Definition 4.3.1 that H has to be connected as well. More generally, we have the following lemma:

Lemma 4.3.2. *Let $(H, \mathcal{V})$ be a graph decomposition of a graph G. Then for any connected subgraph X of G, the graph $H_X := H[\{h \in H \mid X \cap V_h \neq \emptyset\}]$ is connected.*

[5] In [DJKK22] a graph-decomposition is, more generally, defined as a pair $(H, (G_h)_{h \in H})$ of a graph H and a family of (not necessarily induced) subgraphs G_h of G indexed by the nodes of H, instead of a family of vertex sets V_h as in this chapter.

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 X is connected, $H_X = \bigcup_{x \in X} H_x$ is connected as well. □

The next lemma shows that we can obtain a new graph-decomposition by enlarging all bags of a given graph-decomposition simultaneously in that we replace each of its bags by the r-ball around it for some globally fixed $r \in \mathbb{N}$. A version of Lemma 4.3.3 for tree-decompositions was first proven in [DJKK22, Lemma A.1].

Lemma 4.3.3. *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.*

Proof. For notational simplicity, we adopt the following conventions for the course of this proof. Given a subgraph X of G, we write $H_X := H[\{h \in H : V_h \cap V(X) \neq \emptyset\}]$ and $H'_X := H[\{h \in H : V'_h \cap V(X) \neq \emptyset]$.

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 $\mathrm{d}_G(v, w) \leqslant r$, and we fix a shortest v–w path P in G. Then every vertex $p \in P$ satisfies $\mathrm{d}_G(v, p) \leqslant r$ 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 Lemma 4.3.2, there exists a h–$V(H_v)$ path in H_P and hence in H'_v, as desired. □

4.3.2 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.

Definition 4.3.4 (Radial width)**.** Let $(H, \mathcal{V})$ be a graph-decomposition of a graph G modelled on a graph H. The *(inner-)radial width* of $(H, \mathcal{V})$ is

$$\mathrm{radw}(H, \mathcal{V}) := \max_{h \in V(H)} \mathrm{rad}(G[V_h]).$$

Note that if $G[V_h]$ is disconnected for some $h \in H$, then $\mathrm{radw}(H, \mathcal{V}) = \infty$. The *outer-radial width* of $(H, \mathcal{V})$ is $\max_{h \in V(H)} \mathrm{rad}_G(V_h)$. We remark that the outer-radial width is at most the (inner-)radial width.

Given a non-empty class $\mathcal{H}$ of graphs, the *(inner-)radial $\mathcal{H}$-width* of G is

$$\mathrm{radw}_{\mathcal{H}}(G) := \min \{\mathrm{radw}(H, \mathcal{V}) \mid (H, \mathcal{V}) \text{ is a graph-decomposition of } G \text{ with } H \in \mathcal{H}\}.$$

Note that the (inner-)radial $\mathcal{H}$-width will always be at most $\mathrm{rad}(G)$ for a connected graph G if $\mathcal{H}$ contains at least one non-empty graph. The outer-radial $\mathcal{H}$-width is defined analogously.

For classes $\mathcal{H}$ of graphs that we frequently use in this chapter, we name the radial $\mathcal{H}$-width as in Table 4.1.

Table 4.1: 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 (cf. Section 4.7).

$\mathcal{H}$ **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

Under the name 'tree breadth', the concept 'outer-radial tree-width' has previously been studied [DK14a; Lei17] with applications to tree-spanners and routing problems.

Definition 4.3.5 (Radial spread)**.** 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

$$\operatorname{rads}(H, \mathcal{V}) := \max_{v \in V(G)} \operatorname{rad}(H_v).$$

The *outer-radial spread* of $(H, \mathcal{V})$ is $\max_{v \in V(G)} \operatorname{rad}_H(H_v)$. Note that the outer-radial spread is at most the (inner-)radial spread.

We remark that the 'outer' versions of radial width and spread are only used in Lemmas 4.3.6, 4.3.9 and 4.3.10. Hence, we will usually omit the 'inner' when talking about (inner)-radial width and spread.

Berger and Seymour [BS24, p. 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. The next lemma generalises this to arbitrary decomposition graphs.

Lemma 4.3.6. *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. By Lemma 4.3.3, $(H, \mathcal{V}')$ 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 any vertex $x' \in V'_h$ a vertex $x \in V_h$ with $\mathrm{d}_{G[V'_h]}(x', x) = \mathrm{d}_G(x', x) \leqslant k$. Thus, we have $\mathrm{d}_{G[V'_h]}(z_h, x) \leqslant \mathrm{d}_{G[V'_h]}(z_h, x') + \mathrm{d}_{G[V'_h]}(x', x) \leqslant k + k = 2k$. 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 Lemma 4.3.2, 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 $2rk$. □

4.3.3 Interplay with 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 4.1.2. For this, let us recall the definition of quasi-isometries in the case of graphs.

Definition 4.3.7 (Quasi-isometry)**.** 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) $d_H(h,h') \leqslant m \cdot d_G(\varphi(h), \varphi(h')) + a$ for every $h, h' \in V(H)$,

(Q2) $d_G(\varphi(h), \varphi(h')) \leqslant M \cdot d_H(h,h') + A$ for every $h, h' \in V(H)$, and

(Q3) for every vertex $v \in G$, there exists a node $h \in H$ with $d_G(v, \varphi(h)) \leqslant r$.

Usually (cf. [GP23, Section 2.1]), quasi-isometries are denoted with only two parameters $L \geqslant 1$ and $C \geqslant 0$: an (L, C)-*quasi-isometry* is precisely an (L, LC, L, C, C)-quasi-isometry. Here, we use a more detailed set of parameters that allows us to describe in a more precise way how graph-decompositions and quasi-isometries interact.

The following is a well-known fact.

Lemma 4.3.8. *If a graph H is (L, C)-quasi-isometric to a graph G, then G is $(L, 3LC)$-quasi-isometric to H.*

We split the proof of Proposition 4.1.2 into the following two lemmas.

Lemma 4.3.9. *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 (Q3) by (H1).

For (Q2) 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 $h_0h_1 \dots h_\ell$ 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, $v_0P_0v_1 \dots v_\ell P_\ell v_{\ell+1}$ 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), 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 $v_0v_1 \dots v_\ell$ 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 $h_0 := h$ and $h_{\ell+1} := h'$. 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, $h_0P_0h_1 \dots h_\ell P_\ell h_{\ell+1}$ is an h–h' walk in H of length at most $2r_1(\ell + 1)$, as desired. □

Lemma 4.3.10. *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 (Q3), 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, $\mathrm{d}_G(\varphi(h_0), \varphi(h_1))$ is at most $2r' + 1$. By (Q1), 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, $\mathrm{d}_G(\varphi(h_0'), \varphi(h_1'))$ is at most $2r'$. So by (Q1), there is an h_0'–h_1' path P' in H of length at most $2mr' + a$. Let W be the walk $P_0 P' P_1$ 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 (Q2), we obtain

$$\mathrm{d}_G(v, \varphi(h)) \leqslant \mathrm{d}_G(v, \varphi(h')) + \mathrm{d}_G(\varphi(h), \varphi(h')) \leqslant r' + M(mr' + \lceil (m+a)/2 \rceil) + A.$$

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 4.1.2. Use Lemma 4.3.9 for (i) and Lemma 4.3.10 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 [BS24, p. 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 4.3.11. 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 4.3.11 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.

4.3.4 Quasi-geodesic topological minors

In this section, we study quasi-geodesic topological minors as obstructions to small radial width. We show in Lemma 4.3.14 that quasi-geodesic topological minors are also an instance of the more general 'fat minors' used in Conjecture 4.1.1. Fat minors are obstructions to small radial width (Lemma 4.3.15), which then implies that quasi-geodesic topological minors are also such obstructions (Proposition 4.3.16).

Let us recall the definition of $(\geqslant k)$-subdivisions T_kX and c-quasi-geodesity from the introduction.

Definition 4.3.12 (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 4.3.13 (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 $\mathrm{d}_G(u,v) \leqslant c \cdot \mathrm{d}_X(u,v)$. 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 4.1.1. For this, let us first recall the definition of 'fat minors'. 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 $x \neq x_0, x_1$.[6] The V_x are its *branch sets* and the E_e are its *branch paths*. Then X is a *minor* of G if G contains a model of X.

A model $(\mathcal{V}, \mathcal{E})$ of X in G is *K-fat* for $K \in \mathbb{N}$ if $\mathrm{d}_G(Y,Z) \geqslant K$ 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 if G contains a K-fat model of X. We remark that the 0-fat minors of G are precisely its minors.

Lemma 4.3.14. *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

[6] Note that we here deviate from the definition of model from [Die25]. However, 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 [Die25].

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 $T_k X$, this construction directly yields $\mathrm{d}_{X_G}(Y,Z) \geqslant k/2 - 1$ 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 $\mathrm{d}_G(Y,Z) \geqslant \mathrm{d}_{X_G}(Y,Z)/c \geqslant \frac{k-2}{2c}$; so X_G is indeed a K-fat model of X in G. □

Lemma 4.3.15. *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 Lemma 4.3.2.

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 Lemma 4.3.2, 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 4.3.16. *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.* □

4.4 Radial path-width

In this section we prove Theorem 21, which we restate here for convenience.

Theorem 21. *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 21.

Theorem 4.4.1. *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 4.4.1 indeed implies Theorem 21. 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 4.2).

Lemma 4.4.2. *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 4.3.10, 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 4.3.10 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 $k := \lfloor (j-i)/2 \rfloor$. □

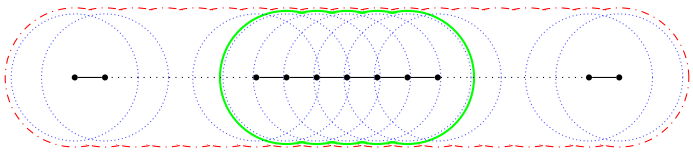

Figure 4.2: A decomposition modelled on the black path using the blue balls as given by Lemma 4.4.2 for $r = 2$, $c = 1$. The bag corresponding to the centre vertex is green.

Proof of Theorem 21 given Theorem 4.4.1. 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 4.4.1, either P' has length at most $18k + 2$ or every vertex of G has 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 4.4.2 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. □

The remainder of this section is devoted to the proof of Theorem 4.4.1. 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 4.4.3, 4.4.4 and 4.4.6 below. They are stated in a slightly more general form than needed for the proof of Theorem 4.4.1, which enables us to use them later also in the proofs of Theorems 22 and 23.

The first lemma, Lemma 4.4.3 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 4.4.4, 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 4.3). The third lemma, Lemma 4.4.6, 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 4.4.3. *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 $\mathrm{d}_P(u, x) = \mathrm{d}_G(u, x) \leqslant \mathrm{d}_G(u, w)$ as $w \in X$. We then have $\mathrm{d}_X(x, w) \leqslant c \cdot \mathrm{d}_G(x, w) \leqslant c \cdot (\mathrm{d}_G(x, u) + \mathrm{d}_G(u, w))$, 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

$$\begin{aligned}\mathrm{d}_{X\cup P}(u, w) &\leqslant \mathrm{d}_{X\cup P}(u, x) + \mathrm{d}_{X\cup P}(x, w) = \mathrm{d}_P(u, x) + \mathrm{d}_X(x, w)\\ &\leqslant \mathrm{d}_G(u, x) + c \cdot (\mathrm{d}_G(x, u) + \mathrm{d}_G(u, w)) \leqslant (2c + 1) \cdot \mathrm{d}_G(u, w),\end{aligned}$$

which shows the claim. □

Lemma 4.4.4. *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 4.3 for a sketch of the situation in Lemma 4.4.4.

The proof of Lemma 4.4.4 builds on the study of how cycles interact with a given separation of the graph. More formally, we have the following preparatory lemma.

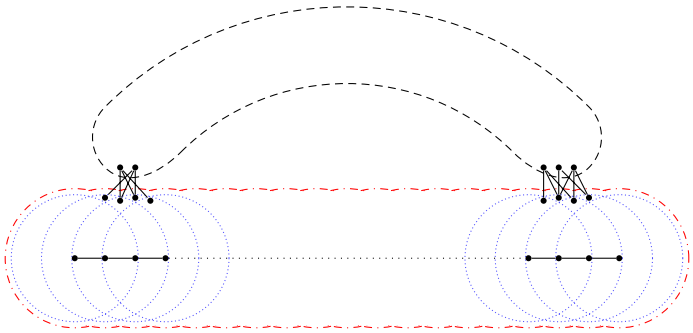

Figure 4.3: The setting of Lemma 4.4.4.

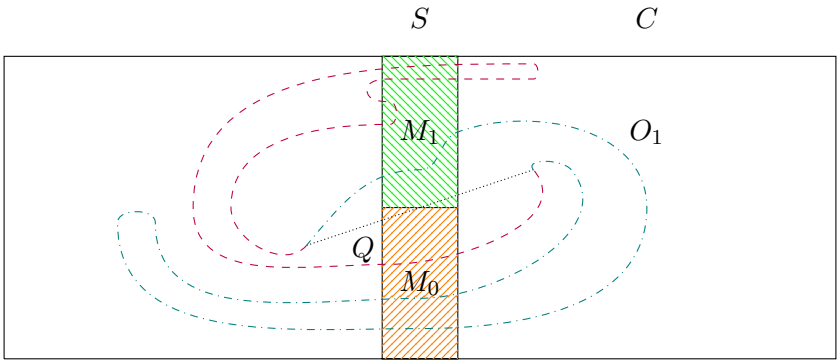

Figure 4.4: The setting of Lemma 4.4.5. Here, O_2 is the cycle containing Q (black, dotted line) and the purple, dashed part of O_1, and O_3 is the cycle containing Q and the teal, dash dotted part of O_1.

Lemma 4.4.5. *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 4.4 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 4.4.4. 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

$$n - m_0 - m_1 \leqslant \mathrm{d}_P(p_{j_0}, p_{j_1}) = \mathrm{d}_G(p_{j_0}, p_{j_1}) \leqslant \mathrm{d}_G(p_{j_0}, u) + \mathrm{d}_G(u, v) + \mathrm{d}_G(v, p_{j_1}) = d_G(u, v) + 2r. \tag{4.1}$$

So $\mathrm{d}_G(u, v) \geqslant n - m_0 - m_1 - 2r = m > 0$ 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 (4.1), 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 4.4.5 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 4.4.6. *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 $\mathrm{d}_G(v, P) \leqslant 2r + \ell$ 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 $n = ||P|| \geqslant ||Q|| = \mathrm{d}_G(v, u) = \mathrm{d}_G(v, q) + \mathrm{d}_G(q, u)$, where the last equality follows from Q being geodesic in G. Additionally, since $p_n \in U$, we have $\mathrm{d}_G(p_j, p_n) \leqslant \mathrm{d}_G(p_j, u) \leqslant \mathrm{d}_G(p_j, q) + \mathrm{d}_G(q, u) = r + \mathrm{d}_G(q, u)$. These two inequalities combine to

$$\mathrm{d}_G(v, q) \leqslant n - \mathrm{d}_G(q, u) \leqslant n - (\mathrm{d}_G(p_j, p_n) - r).$$

All in all, we get

$$\mathrm{d}_G(v, P) \leqslant \mathrm{d}_G(v, p_j) \leqslant \mathrm{d}_G(v, q) + r \leqslant n - (\mathrm{d}_G(p_j, p_n) - r) + r = n - (n - j) + 2r \leqslant 2r + \ell,$$

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 4.4.1.

Proof of Theorem 4.4.1. 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 4.4.3, $P \cup Q$ 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 for a contradiction that there is such a component C. Applying Lemma 4.4.4 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 4.4.6 with $U = \{p_n\}$. Otherwise, $N_G(C) \subseteq \bigcup_{i=n-3k}^{n} B_i$, and the claim follows by applying Lemma 4.4.6 with $U = \{p_0\}$. ■

By Claims 1 to 3, every vertex of G has distance at most $9k$ from P, as desired. □

4.5 Radial cycle-width

In this section we build on our result on radial path-width, Theorem 21, and use a small adaptation of its proof to address radial cycle-width by proving Theorem 22. More precisely, Theorem 22 follows from the lemmas that we have already shown in Section 4.4. Let us restate Theorem 22 here for convenience.

Theorem 22. *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 22.

Theorem 4.5.1. *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 4.5.1 implies Theorem 22.

Proof of Theorem 22 given Theorem 4.5.1. 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 4.5.1 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 4.4.2 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 4.5.1.

Proof of Theorem 4.5.1. Let G be a connected graph that contains no $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph. Applying Theorem 4.4.1 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 $\mathrm{d}_Q(u,v) = \mathrm{d}_C(u,v)$ for every two vertices $u, v \in Q$. Thus, Q is geodesic in G, since C is geodesic in G. By Lemma 4.4.3, $P \cup Q$ then is a 3-quasi-geodesic subgraph of G, as desired. □

4.6 Radial star-width

In this section we prove Theorem 23, which we restate here for convenience.

Theorem 23. *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 some* $C_k \in \mathbb{N}$ *such that some subdivided star is* $(1, C_k)$*-quasi-isometric to* G*.*

Recall that W is the 'wrench'/'spanner' graph depicted in Figure 4.1. 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 23 for details).

Before we start with the proof of Theorem 23, we first show an auxiliary lemma. Recall that in the proof of Theorem 21 we used Lemma 4.4.4 as a tool to show that a graph contains a geodesic $T_k K^3$. Similarly, we will use the following lemma in the proof of Theorem 23 to show that a graph contains a 3-quasi-geodesic $T_{3k} W$.

Lemma 4.6.1. *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* $\mathrm{d}_G(q, r) \geqslant 4 \max\{\mathrm{d}_G(u, P), \mathrm{d}_G(v, P)\}$*;*

(ii) $r \in Q$ *and* $\mathrm{d}_G(q, r) \geqslant 4 \max\{\mathrm{d}_G(v, Q), \mathrm{d}_G(p_0, q), \mathrm{d}_G(p_n, q)\}$*.*

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 4.4.3, 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 $\mathrm{d}_X(q, r) = \mathrm{d}_G(q, r) \leqslant \mathrm{d}_G(q, x) + \mathrm{d}_G(x, y) + \mathrm{d}_G(y, r)$ and thus

$$\begin{aligned}\mathrm{d}_G(x, y) &\geqslant \mathrm{d}_G(q, r) - \mathrm{d}_G(x, q) - \mathrm{d}_G(y, r) \geqslant 4 \max\{\mathrm{d}_G(u, P), \mathrm{d}_G(v, P)\} - \mathrm{d}_G(u, P) - \mathrm{d}_G(v, P) \\ &\geqslant 2 \max\{\mathrm{d}_G(u, P), \mathrm{d}_G(v, P)\} \geqslant 2 \max\{\mathrm{d}_G(x, q), \mathrm{d}_G(y, r)\}.\end{aligned}$$

Therefore, making again use of $\mathrm{d}_G(q, r) \leqslant \mathrm{d}_G(q, x) + \mathrm{d}_G(x, y) + \mathrm{d}_G(y, r)$, we obtain

$$\begin{aligned}\mathrm{d}_X(x, y) &= \mathrm{d}_X(x, q) + \mathrm{d}_X(q, r) + \mathrm{d}_X(r, y) = \mathrm{d}_G(x, q) + \mathrm{d}_G(q, r) + \mathrm{d}_G(r, y) \\ &\leqslant 2\,\mathrm{d}_G(x, q) + 2\,\mathrm{d}_G(y, r) + \mathrm{d}_G(x, y) \leqslant 3\,\mathrm{d}_G(x, y).\end{aligned}$$

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

$$\mathrm{d}_G(x, y) \geqslant \mathrm{d}_G(q, r) - \mathrm{d}_G(x, q) - \mathrm{d}_G(y, r) \geqslant 2 \max\{\mathrm{d}_G(v, Q), \mathrm{d}_G(p_0, q), \mathrm{d}_G(p_n, q)\},$$

and thus

$$\mathrm{d}_X(x, y) = \mathrm{d}_G(x, q) + \mathrm{d}_G(q, r) + \mathrm{d}_G(r, y) \leqslant 2\,\mathrm{d}_G(x, q) + 2\,\mathrm{d}_G(y, r) + \mathrm{d}_G(x, y) \leqslant 3\,\mathrm{d}_G(x, y). \qquad \square$$

Now we prove Theorem 23. We remark that we did not optimise the bounds on the radial width and radial spread.

Proof of Theorem 23. 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 4.3.8 and 4.3.9, there exists an (L,C)-quasi-isometry from G to S such that L and C depend on k only. Applying [NSS25, p. 1.6] (to $L, C, 2$)[7] then yields a constant C'_k such that G is $(1, C'_k)$-quasi-isometric to some subdivided star S'. By Lemma 4.3.8, 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 $0 \leqslant i \leqslant n$ 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 4.4.4 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_kK^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 4.4.6 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 4.4.2 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

[7] Note that we use here that subdivided stars admit path-decompositions into bags of size 3 and thus have path-width 2.

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 4.6.1 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 $G - G_P$, either $N_G(C)$ is a subset of $\bigcup_{i=0}^{3k} B_i$ or a subset of $\bigcup_{i=s-12k-3}^{s+12k+3} B_i =: D_s$ or a subset of $\bigcup_{i=n-3k}^{n} B_i$.*

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 4.4.4 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 4.4.4 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 4.4.6 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 4.4.6 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 4.4.2, the V'_i are the bags of a P-decomposition 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 $\mathscr{C}$ be the set of components that contain a vertex whose distance from P is at least $24k + 5$, and let $\mathscr{C}'$ be the set of all other components of $G - V''$. We now add all vertices in $V'' \setminus V'$ and all components from $\mathscr{C}'$ to the bag V'_s that is indexed by p_s. More precisely, we define

$$V_s := V'_s \cup (V'' \setminus V') \cup \{V(C) \colon C \in \mathscr{C}'\}.$$

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 $G[V_s]$.

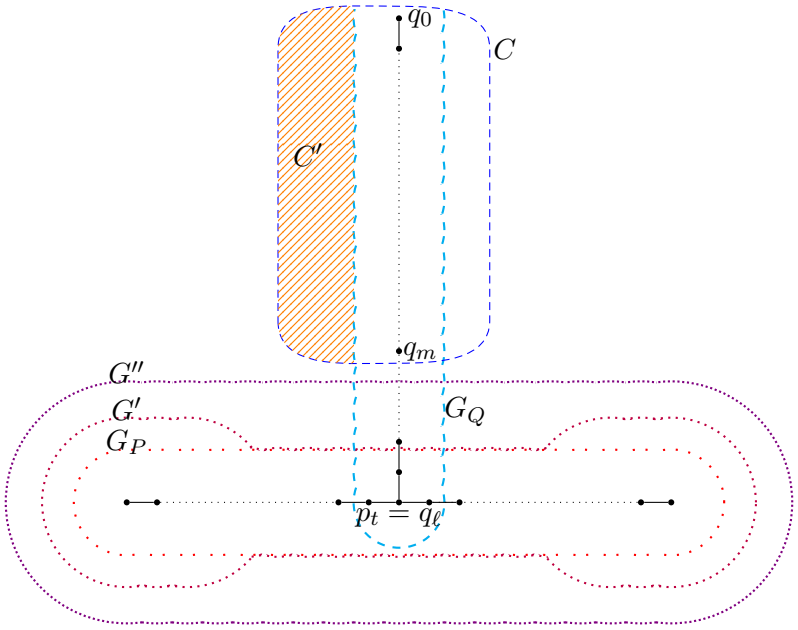

Figure 4.5: The setting of the proof of Claim 6.

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. ■

We now show that all components of $G - V'''$, i.e. those in $\mathscr{C}$, are path-like.

Claim 6. *Let C be a component in $\mathscr{C}$. 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 4.5 to follow the proof more easily. Let q_0 be a vertex of C of maximal distance from P. Since $C \in \mathscr{C}$ 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 4.6.1 (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 4.4.4 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 4.4.6 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 4.4.2, 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

$$\mathrm{d}_G(q_0,q_j)+\mathrm{d}_G(q_j,v)+\mathrm{d}_G(v,w)+\mathrm{d}_G(w,P) \leqslant (m-3k-2)+3k+1+(18k+4) = m+18k+3 < \ell$$

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 $\mathscr{C}$ are precisely the components of $G-V'''$. For every component $C \in \mathscr{C}$, 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. □

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 23 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.

4.7 Where quasi-geodesic topological minors are not the right obstruction

Davies, Hickingbotham, Illingworth and McCarty showed that Conjecture 4.1.1 is false in general. Hence, the strengthening of Conjecture 4.1.1 to quasi-geodesic topological minors also has a negative answer. However, this strengthening fails earlier than Conjecture 4.1.1, and in this section we exhibit sets $\mathcal{X}$ exemplifying this. To make this precise, let us give some background on minor-closed classes of graphs.

A class $\mathcal{G}$ of graphs is *minor-closed* if $G' \in \mathcal{G}$ whenever G' is a minor of some $G \in \mathcal{G}$. Further, a class $\mathcal{G}$ of graphs is *defined by a finite list* $\mathcal{X} = \{X_1, \dots, X_\ell\}$ *of forbidden (topological) minors* if a graph G is contained in $\mathcal{G}$ precisely if no X_i is a (topological) minor of G.

The Graph Minor Theorem by Robertson and Seymour characterises minor-closed graph classes.

Theorem 4.7.1 ([RS04])**.** *Every minor-closed class of graphs is defined by a finite list of forbidden minors.*

Robertson and Seymour showed in [RS90b, (2.1)] that every minor-closed graph property that is defined by a finite list of forbidden minors is also defined by a still finite, but usually significantly larger list of forbidden topological minors. This then yields the following corollary of Theorem 4.7.1:

Corollary 4.7.2. *Every minor-closed class of graphs is defined by a finite list of forbidden topological minors.* □

Consider a minor-closed class of graphs defined by a finite list $\{M_1, \dots, M_k\}$ of forbidden minors (Theorem 4.7.1) and by a finite list $\{X_1, \dots, X_\ell\}$ of forbidden topological minors (Corollary 4.7.2). How are the forbidden minors M_i and the forbidden topological minors X_j related to each other? Robertson and Seymour described their structure in [RS90b, Section 2] by 'splitting' a vertex. Equivalently, one may reformulate it as follows: the graphs $X_1, \dots, X_\ell$ are the models of the $M_1, \dots, M_k$ such that all the induced subgraphs on the branch sets are trees and if the branch set has size at least two, then every vertex in the branch set has degree at least 3 in G. In particular, each X_i *arises* from some M_j in that way.

With this at hand, the following question indeed generalises the strengthening of Conjecture 4.1.1 for quasi-geodesic topological minors:

Question 4.7.3. Let $\mathcal{M}$ be a finite set of finite graphs, and let $\mathcal{X}$ be the finite set of graphs which defines the class of graphs with not M minor with $M \in \mathcal{M}$ via forbidden topological minors; see Corollary 4.7.2. Does there exists a function $f' : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that every graph with no $(\geqslant ck)$-subdivision of some $X \in \mathcal{X}$ as a c-quasi-geodesic subgraph for some $c \in N$ is $f'(k)$-quasi-isometric to a graph with no topological X minor for any $X \in \mathcal{X}$?

As stated before, Question 4.7.3 has a negative answer even when $\mathcal{X}$ is the corresponding list of forbidden topological minors to very simple sets $\mathcal{M}$ such as $\{K_{3,3}, K^5\}$, which characterises the planar graphs [Wag37]. The following proposition provides a sufficient criterion for this, via the detour of graph-decompositions as established in Proposition 4.1.2.

Proposition 4.7.4. *Let $\mathcal{M}$ be a finite set of finite graphs, let $\mathcal{H}$ be the class of graphs without M minor for $M \in \mathcal{M}$, and let $\mathcal{X}$ be the finite list of forbidden topological minors characterizing $\mathcal{H}$ (which exists by Corollary 4.7.2). Let m be the smallest number of vertices of degree $\geqslant 3$ among all $M \in \mathcal{M}$. Suppose that for each $M \in \mathcal{M}$, if M has precisely m vertices of degree at least 3, then it also has a vertex of degree 4.*

Then there exists, for every given $R \in \mathbb{N}$, a graph G which has radial $\mathcal{H}$-width at least R and contains no $(\geqslant 1)$-subdivision of X_i for all $1 \leqslant i \leqslant \ell$.

Proof. Let $\mathcal{M} = \{M_1, \dots, M_k\}$. Without loss of generality, we may assume that M_1 has precisely m vertices of degree at least 3, has minimal maximum degree among all such M_j, and is minimal among these M_j with respect to the number of vertices of degree $\Delta(M_1)$. Let G' be the graph obtained from M_1 by subdividing each edge at least $4R$ times. We then let G be the graph that arises from G' by 'splitting' a vertex v of degree $\Delta(M_1)$ into two adjacent vertices, one of which has degree exactly 3. More formally, G is the model of G' that differs from G' only in that the branch set to v consists of the endvertices u_e and v_e of an edge $e \in G'$ such that u_e has degree 3 and thus v_e has degree $d_{G'}(v) - 1$. Note that G has precisely one more vertex of degree at least 3 than M_1, maximum degree at most $\Delta(M_1)$, and one vertex of degree $\Delta(M_1)$ less than M_1. We now show that G is as desired.

Suppose for a contradiction that G has an $\mathcal{H}$-decomposition of radial width less than R. From it, we then obtain an $\mathcal{H}$-decomposition of G' by adding v to every bag that contains at least one of v_e and u_e and deleting v_e and u_e from all bags. By the construction of G from G', this $\mathcal{H}$-decomposition of G' has again radial width less than R. But G' has radial $\mathcal{H}$-width at least R by Proposition 4.3.16, a contradiction.

It remains to show that G contains no $(\geqslant 1)$-subdivision of any $X \in \mathcal{X}$. Suppose for a contradiction that G contains a subgraph X_G that is a $(\geqslant 1)$-subdivision of some $X \in \mathcal{X}$. This X arises from some $M \in \mathcal{M}$. We now distinguish two cases depending on whether M has precisely m vertices of degree at least 3 or not.

If M has more than m vertices of degree at least 3, then M and hence X each have at least $m + 1$ vertices of degree at least 3. Thus, all vertices of degree at least 3 in G are branch vertices of X_G, and in particular, X_G contains both u_e and v_e as branch vertices. Since $d_G(u_e) = 3$, this yields $e \in X_G$. So X_G is not a $(\geqslant 1)$-subdivision of X, contradicting our assumption.

So suppose that M has precisely m vertices of degree at least 3. With the same argument as above, u_e and v_e cannot both be branch vertices of X_G of degree at least 3. Thus, X_G has at most m branch vertices of degree at least 3. This implies $X = M$, as X arises from M. Now we cannot have $\Delta(M) > \Delta(M_1)$, as $\Delta(G) \leqslant \Delta(M)$ and the subdivision X_G of $X = M$ is a subgraph of G. So by our minimal choice of M_1, M has maximum degree $\Delta(M_1)$ and contains at least as many vertices of degree $\Delta(M_1)$ as M_1. But G has one vertex of degree $\Delta(M_1)$ less than M_1, and this contradicts that the subdivision X_G of $X = M$ is a subgraph of G. This completes the proof. □

4.8 Counterexamples to versions in the star and tree case

A possible strengthening of Theorem 23 and Conjecture 4.1.3 is suggested by the proofs of the path- and cycle-case (see Theorems 4.4.1 and 4.5.1). The path- and cycle-decompositions which we exhibit there have an additional property: their decomposition graph H is itself a C-quasi-geodesic (in fact geodesic) subgraph of G such that every vertex of H is contained in its corresponding bag; here, the constant C is independent of the constant c in 'c-quasi-geodesic ($\geqslant ck$)-subdivision' in Theorem 23 and Conjecture 4.1.3. The star-decomposition constructed in the proof of Theorem 23, however, does not have this extra property, and in fact there is not always such a strengthened star-decomposition as the following example shows (see Example 4.8.3 for more details).

Example 4.8.1. Let $d, k \geqslant 1$ be integers. Let G be a subdivision of the graph W depicted in Figure 4.1 such that the edge between the two vertices of degree 3 is subdivided less than k times and every other edge is subdivided at least d times. Then G has radial star-width at most $\lceil k/2 \rceil$. But if a subdivided star S is a subgraph of G, then any S-decomposition of G with the extra property that every vertex of S is contained in its corresponding part has radial width at least $\lceil (d+1)/2 \rceil$.

Still, the proof of Theorem 23 can be adapted (with a slight worsening of the bounds) to obtain the following weaker property for star-decompositions: For a large, but bounded integer C, the decomposition star S can be obtained from a tree T that is a C-quasi-geodesic subgraph of G by contracting a subtree of bounded radius to a single vertex s. This vertex s is then the centre of the star S, and the set of vertices of T that were contracted to s is a subset of the bag of the decomposition that corresponds to s.

An analogue of Example 4.8.1 also holds in the more general case of tree-decompositions: there is no $C \in \mathbb{N}$ such that if a graph G has low radial tree-width, then there exists a tree-decomposition of G of small radial width whose decomposition tree is a C-quasi-geodesic subgraph of G. We prove this with our rather involved Example 4.8.4, which is related to constructions in [BS24; KLMPW03]. The existence of such an example suggests that any further work towards an answer to Conjecture 4.1.3 will have to build on new methods for constructing graph-decompositions of low radial width.

4.8.1 Counterexamples to stronger structural dualities

Here we present further details for the examples mentioned above. We first give examples of graphs of large radial path-width which contain neither a long c-quasi-geodesic cycle nor a c-quasi-geodesic ($\geqslant 1$)-subdivision of the claw for any $1 \leqslant c < 2$.

Example 4.8.2. Let k be an integer.

(i) Assume $k \geqslant 1$. By completing the neighbourhood of the centre vertex of the k-subdivided claw to a K^3, we obtain a graph G which contains a 2-quasi-geodesic k-subdivided claw but G contains neither a c-quasi-geodesic cycle of length $\geqslant 4$ nor a c-quasi-geodesic claw for every c with $1 \leqslant c < 2$.

(ii) ([BS24]) Assume $k \geqslant 2$. The $k \times k$ grid has radial tree-width at least $(k-2)/2$ by Proposition 4.3.16, since it contains a cycle of length $4(k-1)$ as 2-quasi-geodesic subgraph; but every c-quasi-geodesic cycle with $1 \leqslant c < 2$ has length 4.

Our second example, a rephrasing of Example 4.8.1, demonstrates that, given a graph G of small radial star-width, we cannot always find a star-decomposition whose decomposition star is a c-quasi-geodesic subgraph of G such that every vertex of this subgraph belongs to its corresponding bag.

Example 4.8.3. Let $k, d \geqslant 0$ be integers. Let G be a subdivision of the graph W depicted in Figure 4.1 such that the edge of W between the two vertices u, v of degree 3 is subdivided at most k-times and every other edge is subdivided at least d-times. Then G has a star-decomposition $(S, (V_t)_{t \in S})$ of radial width at most $\lceil (k+1)/2 \rceil$, but there is no star-decomposition $(S', (V'_t)_{t \in S})$ of radial width at most $\lfloor d/2 \rfloor$ such that the star S' is a subgraph of G.

Proof. First, we show that there is a star-decomposition $(S, (V_t)_{t \in S})$ of low radial width: Let S be the subdivided star obtained from G by contracting the unique u–v path P in G to a vertex s. Now let $V_s := V(P)$, and $V_t := \{t\} \cup N_G(t)$ for every other node $t \neq s$ of S. Then $(S, (V_t)_{t \in S})$ is the desired star-decomposition of G.

Secondly, we show that every star-decomposition $(S', (V'_t)_{t \in S})$ of G where S' is a subgraph of G has radial width at least $\lfloor d/2 \rfloor$: Since S' is a subgraph of G which is a star, it is either a path or a star centred at u or v. Hence, one of the four vertices of G of degree 1 has distance $d+1$ to S'. □

We now present an example that in general, even for graphs with low radial tree-width, we cannot hope to obtain tree-decompositions of low radial width where the decomposition tree is a c-quasi-geodesic subtree (with bounded c). The trees of wheels form a class of graphs that all have radial tree-width 1. But by Lemma 4.8.6, for all $c, n \in \mathbb{N}$ there is a tree of wheels G such that every H-decomposition of G, with a c-quasi-geodesic subtree H of G with $h \in V_h$ for all nodes $h \in H$, must necessarily have radial width at least n.

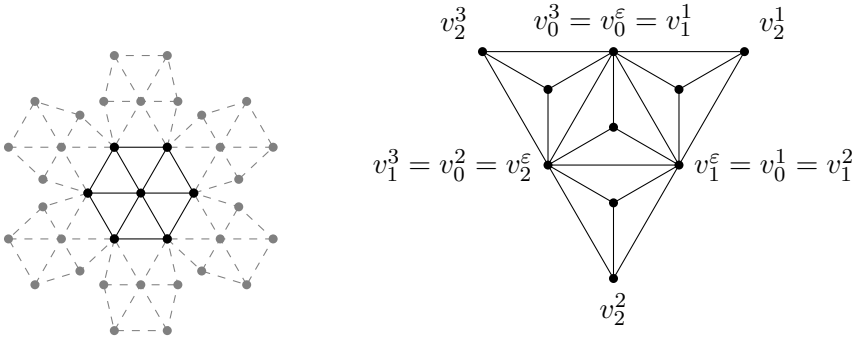

Figure 4.6: Two trees of wheels. The left graph is a tree of wheels of depth 1 for $n = 6$, and the edges that do not belong to the central wheel are dashed. The right graph is a tree of wheels of depth 1 for $n = 3$, and the vertices that are not hubs are labelled.

Example 4.8.4 (Tree of wheels)**.** Let $n \geqslant 3$ be an integer. Recall that a *wheel* W_n is the graph obtained from a cycle C^n and a single vertex K^1 by adding edges from all vertices in the cycle to K^1. We will denote the cycle, the *rim* of the wheel, by $v_0v_1 \cdots v_n$ where $v_0 = v_n$, and the edge v_iv_{i+1} by e_{i+1}. The vertex of K^1 is the *hub* of the wheel.

In the following construction, all wheels are copies of W_n. We start with a wheel in *depth* 0, which we refer to as the *central wheel*. In the first step, we add for every edge on the rim of the central wheel another wheel. More precisely, we take, for each edge e_i on the rim of the central wheel, a new wheel and identify the vertex v_1 of the new wheel with v_{i-1} of the central wheel, identify v_0 of the new wheel with v_i of the central wheel and identify the edge e_1 of the new wheel with e_i of the central wheel. We consider the new wheels to be in *depth* 1.

We now repeat this step: Say we have already added wheels in depth $j-1$ but not wheels in depth j. At each edge e_i with $i \neq 1$ of one of the wheels in depth $j-1$ we add another wheel. We consider the new wheel to be in *depth* j. More precisely, let e_i with $i \neq 1$ be an edge of a wheel W in depth $j-1$. We then add a new wheel for e_i, identify v_1 of the new wheel with v_i of W, identify v_0 of the new wheel with v_{i+1} of W and identify e_1 of the new wheel with e_i.

The graph which we obtained in step d (that is, we added the wheels in depth d but no wheels in depth $d+1$) is the *tree of wheels of depth* d.

Two trees of wheels of depth 1 are depicted in Figure 4.6. Note that for every tree of wheels, there is a tree-decomposition where every bag is the vertex set of a wheel. So every tree of wheels has radial tree-width 1.

For the following lemmas we need to establish some notation regarding trees of wheels. Let G be a tree of wheels of depth r. Every wheel in depth j was added at an edge e_i of a wheel in depth $j-1$, and this latter wheel was again added at an edge $e_{i'}$ of a wheel in depth $j-2$ to the central wheel and so on. So every wheel is uniquely identified by the indices of the edges to which previous wheels were added. More precisely, we associate the central wheel with the empty sequence ε. When a wheel is added at edge e_i of a wheel that is associated with the sequence $i_1i_2 \cdots i_j$, then the new wheel is associated with the sequence $i_1i_2 \cdots i_ji$. For a sequence s that is associated with a wheel, we also denote the edge e_i of that wheel by e_i^s

and the vertex v_i of that wheel by v_i^s.

In [KLMPW03, Section 6] a graph is described that is very close to the tree of wheels on $3+1$ vertices, only without the hubs. There it is also proven that for a given c, such a graph of sufficiently large radius does not have a spanning tree that is c-quasi-geodesic.

Lemma 4.8.5. *Let*[8] $n \geqslant 4$, $r, c \in \mathbb{N}$ *and let* G *be a tree of wheels on* $n+1$ *vertices of depth* d. *If* G *contains a* c*-quasi-geodesic tree that contains all vertices on rims of wheels of* G, *then* $c \geqslant 2(d+1)$.

Proof. We need some additional notation. Let $s \neq \varepsilon$ be a sequence associated with a wheel. Then G splits into two connected graphs G_1^s and G_2^s which only meet in $G[\{v_0^s, v_1^s\}]$. We choose G_1^s such that it contains the hub of the central wheel.

Assume that G contains a c-quasi-geodesic tree T that contains all vertices on rims of wheels of G. We will now recursively construct a sequence s such that the wheel associated with s contains two vertices whose distance in T is large. In the first step let $I \subseteq \{1, \dots, n\}$ consist of those indices i such that $v_i = v_1^i$ and $v_{i+1} = v_0^i$ are contained in the same component of $T \cap G_2^i$. As T is acyclic and the G_2^i are edge-disjoint, there is $i \in \{1, \dots, n\}$ with $i \notin I$. In particular, e_0^i is not contained in T, and v_0^i and v_1^i are contained in the same component of $T \cap G_1^i$, since T is connected and all vertices on rims of wheels are in T by assumption. Thus, the unique v_0^i–v_1^i path in T is contained in $T \cap G_1^i$ and has at least two edges.

Now we repeat the following step recursively. Assume that $j \leqslant d-1$ and that a sequence $s = i_1 i_2 \dots i_j$ associated with a wheel has been constructed such that v_0^s and v_1^s are not contained in the same component of $T \cap G_2^s$ and have distance at least $2j$ in T. Then it directly follows that there also is some $i \neq 1$ such that $v_i^s = v_1^{si}$ and $v_{i+1}^s = v_0^{si}$ are not contained in the same component of $T \cap G_2^s$. As G_2^{si} is a subgraph of G_2^s, the vertices v_0^{si} and v_1^{si} are not contained in the same component of $T \cap G_2^{si}$. Furthermore, the unique v_0^{si}–v_1^{si} path in T contains v_0^s and v_1^s and thus, since $n \geqslant 4$, v_0^{si} and v_1^{si} have distance at least $2(j+1)$ in T.

So assume we obtained a sequence $s = i_1 i_2 \dots i_r$ such that v_0^s and v_1^s are not contained in the same component of $T \cap G_2^s$ and have distance at least $2d$ in T. As above, there is some integer $i \neq 1$ such that v_i^s and v_{i+1}^s are not contained in the same component of $T \cap G_2^s$, and thus they have distance at least $2(d+1)$ in T while they are adjacent in G, implying $c \geqslant 2(d+1)$. □

Lemma 4.8.6. *Let*[9] $n \geqslant 6$, $r \in \mathbb{N}$ *and let* G *be a tree of wheels on* $n+1$ *vertices of depth* d. *If there is a tree* T *in* G *that is* c*-quasi-geodesic for some* $c \in \mathbb{N}$ *such that for some* $q \in \mathbb{N}$ *all vertices of* G *have distance at most* q *from* T, *then* $c \geqslant (2/3) \cdot (d - \lceil q/2 \rceil + 1)$.

Proof. Denote $\lceil q/2 \rceil$ by p. As otherwise the conclusion is void, we may assume that $d > p-1$.

For every sequence $s \neq \varepsilon$ of length at most $d-(p-1)$ associated with a wheel, the tree T contains v_0^s or v_1^s. Indeed, let t be a sequence of integers 4 of length $p-1$. Since $n \geqslant 6$, v_3^s and v_4^s have distance 2 from v_0^s and v_1^s in G. Applying this fact iteratively yields that v_3^{st} has distance $2p \geqslant q$ from v_0^s and v_1^s in G and even larger distance from all other vertices of G_1^s.

[8] The proof can be adapted to $n = 3$, but then we get a different bound on c.

[9] The proof can be adapted to $n \geqslant 3$, but then we get a different bound on c.

So T contains a vertex of G_2^s. In particular, T contains a vertex in $G_2^{s'} \subseteq G_1^s$ by applying the above argument on $s' = i$ for some i distinct to the first integer in the sequence of s. As T is connected, this implies that T contains v_0^s or v_1^s.

If it exists, let $s \neq \varepsilon$ be a sequence of length at most $d-(p-1)-(c+1)$ associated with a wheel and let t be a sequence of integers 2 of length $c+1$. Assume for a contradiction that T does not contain v_1^s. Since $v_1^s = v_1^{s2} = \cdots = v_1^{st}$ and the length of st is at most $d-(p-1)$, it follows from the above that T contains both v_0^s and v_0^{st}. And the path in T from v_0^s to v_0^{st} contains all $v_0^{st'}$ where t' is an initial segment of t. As the $v_0^{st'}$ are non-adjacent, the distance of v_0^s and v_0^{st} in T is at least twice the length of t which in turn is larger than c. This is a contradiction to the fact that v_0^s and v_0^{st} have distance 2 in G and that T is c-quasi-geodesic. So T contains v_1^s, and by letting t be a sequence of integers n of length $c+1$ the same argument yields that T contains v_0^s. Hence, T contains all vertices on rims of wheels of G associated to sequences of length at most $d-(p-1)-(c+1)$. So Lemma 4.8.5 yields $c \geqslant 2(d-(p-1)-(c+1)+1)$ and thus $c \geqslant (2/3)\cdot(d-\lceil q/2\rceil+1)$.

So we are left with the case that there is no sequence $s \neq \varepsilon$ of length at most $d-(p-1)-(c+1)$ associated with a wheel, that is, $d-p-c \leqslant 0$. This is equivalent to $c \geqslant d-p$, so either $d-p \geqslant (2/3)(d-p+1)$ and the lemma holds, or $d-p < (2/3)(d-p+1)$. In the latter case $d-p<2$, so $d-p+1 \leqslant 2$. Hence in this remaining case either $c \geqslant 4/3$ and the lemma holds, or $c=1$ and the lemma holds by Lemma 4.8.5. □

CHAPTER 5

A characterisation of graphs quasi-isometric to K_4-minor-free graphs

This chapter is based on [AJKW24], which is joint work with Sandra Albrechtsen, Raphael W. Jacobs and Paul Wollan.

5.1 Introduction

5.1.1 Quasi-isometry and fat minors

Gromov's [Gro93] 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 5.2.3 for the definition).

Following Gromov's idea into the realm of graphs, Georgakopoulos and Papasoglu [GP23] 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 [GP23, 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 5.2.2 for the definition).

Conjecture 24 ([GP23], Conjecture 1.1)**.** *Let X be a finite graph. Then 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 24 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 24 holds for any (not even necessarily finite) graph X.

Recently, Conjecture 24 has been disproved by Davies, Hickingbotham, Illingworth and McCarty for general graphs X [JRFR24, Theorem 1]. However, Conjecture 24 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 [Man05] (see [GP23] or [BS24] for a graph-theoretic proof). The case $X = K_{2,3}$ was proved by Chepoi, Dragan, Newman, Rabinovich,

and Vaxes [CDNRV12], which characterises graphs quasi-isometric to an outerplanar graph. Further, Fujiwara and Papasoglu [FP23] 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 [GP23]. Additionally, together with Diestel, Elm and Fluck, we showed in [ADEFJKW23] 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.

5.1.2 Our results

The main contribution of this chapter is the resolution of Conjecture 24 for $X = K_4$:

Theorem 25. *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 25, 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 24 for $X = K_4^-$ given by Fujiwara and Papasoglu [FP23] 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 24 from a different angle: Graph-decompositions [DJKK22] 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 [DJKK22] as well as the study of local separations [CJKK25] and of locally chordal graphs [AJKK25]. In [ADEFJKW23, § 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 Sections 5.2.4 and 5.2.5 for details). To prove Theorem 25, 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 25.

Our techniques additionally yield a new proof of Conjecture 24 for $X = K_4^-$, which is significantly shorter than the original one by Fujiwara and Papasoglu [FP23]:

Theorem 26. *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 26. 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 25 and 26 for arbitrary (finite or infinite) graphs[2]. Thus, Theorems 25 and 26 transfer verbatim to length spaces, a generalisation of geodesic metric spaces [GP23, Observation 2.1].

5.1.3 Applications

Bonamy, Davies, Esperet and Wesolek brought the following consequences of Theorem 25 to our attention [BDEW24].

Gromov [Gro93] introduced *asymptotic dimension* of metric spaces in the context of geometric group theory; see [BD08] 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 [BBEGLPS23, 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 $K \in \mathbb{N}$ [BBEGLPS23, Question 5]. It now follows from Theorem 25 that this is true for $X = K_4$.

Theorem 5.1.1. *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 [BBEGLPS23, Theorem 1.2]. Let $K \in \mathbb{N}$. By Theorem 25, every graph with no K-fat K_4 minor is $f(K)$-quasi-isometric to a graph in $\mathcal{H}$ where f is the function from Theorem 25. Since asymptotic dimension is preserved under quasi-isometry [BD08, Proposition 22], the class of graphs with no K-fat K_4 minor thus has asymptotic dimension at most 1, too. □

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 [Bur65] 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 [PT23]. In the following, we will refer to the induced subgraphs of graphs in the Burling sequence as *Burling graphs*.

[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 5.2.1 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 [FP23, 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 5.5.8.

[3] For the various equivalent definitions of 'asymptotic dimension', we refer the reader to [BBEGLPS23, § 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 $\leqslant f(r)$ [BBEGLPS23, Proposition 1.17].

From Theorem 5.1.1 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 5.1.2. *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 5.1.1, 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 [PT24, 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. □

5.1.4 An open conjecture

Even though Conjecture 24 is disproved for general graphs X [JRFR24], we would like to draw the reader's attention to the following special case[4] of Conjecture 24, which is still open. It is a coarse variant of Kuratowski's theorem and was first conjectured by Georgakopoulos and Papasoglu [GP23]:

Conjecture 27. *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.*

5.1.5 How this chapter is organised

In Section 5.2 we recall some important definitions and facts about the interplay between quasi-isometries and graph-decompositions of small radial width. Section 5.3 then contains some preparatory lemmas, which simplify the proofs of Theorems 25 and 26. Finally, we prove Theorem 26 in Section 5.4 and Theorem 25 in Section 5.5.

5.2 Preliminaries

We warn the reader that graphs in this chapter may be infinite, unless specified otherwise. Our notions mainly follow [Die25]. In particular, if $P = p_0 \dots p_n$ is a path in a graph G, then we denote by p_iPp_j for $i, j \in [n]$ the subpath $p_ip_{i+1} \dots p_j$ of P. Further, we denote by $\bar{P}$ the path $p_n \dots p_0$, and by $\mathring{P}$ the subpath p_1Pp_{n-1} of P. We recall that 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 C of G, a C*-path* is a non-trivial path which meets C precisely in its endvertices. A path *through* a set U of vertices of G is a path in G whose internal vertices are contained in U.

[4] Note that one may replace the graph X in Conjecture 24 with any finite set of finite graphs.

We denote by $\mathscr{C}(G)$ 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 $N_G(U) \cap V(C)$.

5.2.1 Distance, radius and balls

Let G be a graph. We write $\mathrm{d}_G(v, u)$ for the distance of the two vertices v and u in G. For two sets U and U' of vertices of G, we write $\mathrm{d}_G(U, U')$ for the minimum distance of two elements of U and U', respectively. If one of U or U' is just a singleton, then we may omit the braces, writing $\mathrm{d}_G(v, U') := \mathrm{d}_G(\{v\}, U')$ for $v \in V(G)$.

Given a set U of vertices of G, the *ball (in G) around U of radius $r \in \mathbb{N}$*, denoted as $B_G(U, r)$, 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 $\mathrm{d}_G(w, v) \leqslant k$ 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 $\mathrm{d}_G(U, V(Y))$, $\mathrm{rad}_G(V(Y))$, $B_G(V(Y), r)$, and $G[V(Y), r]$ as $\mathrm{d}_G(U, Y)$, $\mathrm{rad}_G(Y)$, $B_G(Y, r)$, and $G[Y, r]$, respectively.

5.2.2 Fat minors

In this chapter we deviate from the definition of minor in [Die25], 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 $x \neq x_0, x_1$.[5] 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 $\mathrm{d}_G(Y, Z) \geqslant K$ 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 if G contains a (K-fat) model of X. We remark that the 0-fat minors of G are precisely its minors. By $\mathrm{Forb}_{\preccurlyeq}(X)$, we denote the class of all graphs with no X minor.

[5] 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 [Die25].

5.2.3 Quasi-isometries

For $M \in \mathbb{R}_{\geqslant 1}$ and $A \in \mathbb{R}_{\geqslant 0}$, an (M, A)-*quasi-isometry* from a graph H to a graph G is a map $\varphi\colon V(H) \to V(G)$ such that

(Q1) $M^{-1} \cdot \mathrm{d}_H(h, h') - A \leqslant \mathrm{d}_G(\varphi(h), \varphi(h')) \leqslant M \cdot \mathrm{d}_H(h, h') + A$ for every $h, h' \in V(H)$, and

(Q2) for every vertex v of G, there exists a node h of H with $\mathrm{d}_G(v, \varphi(h)) \leqslant A$.

We say that H is (M, A)-*quasi-isometric* to G if there exists a (M, A)-quasi-isometry from H to G. The following is a well-known fact.

Lemma 5.2.1. *If a graph H is (M, A)-quasi-isometric to a graph G, then G is $(M, 3AM)$-quasi-isometric to H.* □

5.2.4 Graph-decompositions

Let G and H be graphs and let $\mathcal{V} = (V_h)_{h \in H}$ be a family of sets V_h of vertices of G. We call $(H, \mathcal{V})$ an *H-decomposition* of G, a *decomposition* of G *modelled on* H, or just a *graph-decomposition* [DJKK22], 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.

The sets V_h are called the *bags* of this graph-decomposition, their induced subgraphs $G[V_h]$ are its *parts*, and the graph H is its *decomposition graph*. Whenever a graph-decomposition is introduced as $(H, \mathcal{V})$, we tacitly assume $\mathcal{V} = (V_h)_{h \in H}$. A graph-decomposition $(H, \mathcal{V})$ is *honest*, if all its bags V_h are non-empty and for every edge h_0h_1 of H the bags V_{h_0} and V_{h_1} intersect non-emptily. By [ADEFJKW23, Lemma 2.3], 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.

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.

Let G be a graph and let $\mathcal{H} = (H, \mathcal{V})$ be a partial graph-decomposition of G with support Y. The *outer-radial width* of $(H, \mathcal{V})$ is

$$\mathrm{oradw}_G(\mathcal{H}) := \sup_{h \in H} \mathrm{rad}_G(V_h),$$

and the *(inner-)radial spread* of $(H, \mathcal{V})$ is

$$\mathrm{irads}(\mathcal{H}) := \sup_{v \in Y} \mathrm{rad}_{H_v}(H_v).$$

Additionally, we set $\operatorname{irads}(\mathcal{H}, v) := \operatorname{rad}_{H_v}(H_v)$ for a vertex v of Y. 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 $\operatorname{irads}(\mathcal{H}) \leqslant R_1$.

We remark that the inner-radial spread of $(H, \mathcal{V})$ is at least the *outer-radial spread* $\operatorname{rads}(\mathcal{H}) = \sup_{v \in Y} \operatorname{rad}_H(H_v)$ of $(H, \mathcal{V})$, which we have used in [ADEFJKW23, § 3.4]. The reason why we prefer to work with inner-radial spread in this chapter is that the inner-radial spread of a graph-decomposition $(H, \mathcal{V})$ does not change if we delete edges h_0h_1 of H whose corresponding intersection $V_{h_0} \cap V_{h_1}$ of bags is empty.

The following observation is immediate from the definitions.

Observation 5.2.2. *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 H^C and $V_h := V_h^C$ for all $h \in V(H)$ and the unique $C \in \mathscr{C}(G)$ such that $h \in V(H^C)$.*

In particular, $\operatorname{oradw}_G(\mathcal{H}) = \sup_{C \in \mathscr{C}(G)} \operatorname{oradw}_C(\mathcal{H}^C)$ *and* $\operatorname{irads}(\mathcal{H}) = \sup_{C \in \mathscr{C}(G)} \operatorname{irads}(\mathcal{H}^C)$. *Moreover, if, for some connected graph X, we have $H^C \in \operatorname{Forb}_{\preccurlyeq}(X)$ for all $C \in \mathscr{C}(G)$, then also $H \in \operatorname{Forb}_{\preccurlyeq}(X)$.* □

5.2.5 Interplay between quasi-isometries and graph-decompositions

By [ADEFJKW23, § 3.4], a graph H is quasi-isometric to a graph G if and only if G admits an honest decomposition modelled on H of bounded outer-radial width and bounded radial spread, where the constants of the quasi-isometry only depend on the outer-radial width and the radial spread of the graph-decomposition and vice versa.

Lemma 5.2.3 ([ADEFJKW23], Lemma 3.9)**.** *If a graph G admits an honest (R_0, R_1)-radial decomposition modelled on a graph H, then there is a $(\max\{2R_0, 2R_1\}, \max\{2R_0, 2R_1\})$-quasi-isometry from H to G.*

Lemma 5.2.4 ([ADEFJKW23], Lemma 3.10)**.** *If a graph H is (M, A)-quasi-isometric to some graph G for $M \in \mathbb{N}_{\geqslant 1}$ and $A \in \mathbb{N}$, then G admits an honest $(M^2A\lceil (M+A)/2 \rceil + 2A, M(4A+1) + 2A + 1)$-radial H-decomposition.*[6]

By Lemmas 5.2.3 and 5.2.4, Conjecture 24 can thus be reformulated in terms of graph-decompositions in the following qualitatively equivalent way.

Conjecture 24'. *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 $f(K) = (f_0(K), f_1(K))$ and vice versa.

[6] It is immediate from the construction in [ADEFJKW23, Lemma 3.10] that the H-decomposition is honest. Also the proof yields that the constructed H-decomposition has in fact not only outer-radial spread but even inner-radial spread at most $M(4A+1) + 2A + 1$.

5.3 Preparatory work: Sequences of graph-decompositions

As we have mentioned already in the introduction, we will prove Theorems 25 and 26 by showing the equivalent statements about graph-decompositions, that is, we will prove Conjecture 24' for $X = K_4^-$ and $X = K_4$. 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 5.3.5) 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 5.3.1. *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 $C \in \mathscr{C}(G - Y^d)$ 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 $C \in \mathscr{C}(G - Y)$, 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 $\{\mathcal{H}^C \mid C \in \mathscr{C}(G-Y)\}$ of partial graph-decompositions of G is *feasible for $\mathcal{H}$* if each $\mathcal{H}^C$ is feasible for $\mathcal{H}$ and C.

Construction 5.3.2. Assume that $\mathcal{H} = (H, \mathcal{V})$ is a partial graph-decomposition of a graph G with support Y. Let $\{\mathcal{H}^C \mid C \in \mathscr{C}(G-Y)\}$ 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 $C \in \mathscr{C}(G-Y)$. 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 $C \in \mathscr{C}(G-Y)$. Set $Y' := Y \cup \bigcup_{C \in \mathscr{C}(G-Y)} Y^C$.

Note that the graph H' from Construction 5.3.2 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[7].

The following observation about $\mathcal{H}'$ from Construction 5.3.2 is immediate from its construction.

Observation 5.3.3. *Let $G, \mathcal{H}$ and the $\mathcal{H}^C$ be as in Construction 5.3.2. Then $\mathcal{H}'$ constructed as in Construction 5.3.2 from $\mathcal{H}$ and the $\mathcal{H}^C$ is a partial graph-decomposition of G with support Y', which extends $\mathcal{H}$. Moreover,*

(i) *if $H, H^C \in \mathrm{Forb}_{\preccurlyeq}(X)$ for all $C \in \mathscr{C}(G-Y)$ 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 $C \in \mathscr{C}(G-Y)$, then $\mathcal{H}'$ is honest,*

(iii) $\mathrm{rad}_G(V'_h) = \mathrm{rad}_G(V_h)$, *if $h \in V(H)$,*

(iv) $\mathrm{rad}_G(V'_h) = \mathrm{rad}_G(V^C_h)$, *if $h \in V(H^C) \setminus \{h^C\}$ for some $C \in \mathscr{C}(G-Y)$,*

(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 $C \in \mathscr{C}(G-Y)$, and*

(vii) $\mathrm{irads}(\mathcal{H}', v) \leqslant \mathrm{irads}(\mathcal{H}, v) + 2\sup_{C \in \mathscr{C}(G-Y)} \mathrm{irads}(\mathcal{H}^C, v)$ *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 $G - B_G(w, R)$.

[7] 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 25 and 26 that the arising graph H' has no K_4^-- or K_4 minor, we will not do this here.

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 5.3.4. *For every $R \in \mathbb{N}$, every honest R-component-feasible partial graph-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 $\operatorname{oradw}(\mathcal{H}') \leqslant \max\{\operatorname{oradw}(\mathcal{H}), R\}$, $\operatorname{irads}(\mathcal{H}', v) = \operatorname{irads}(\mathcal{H}, v)$ for $v \in V(Y)$, and $\operatorname{irads}(\mathcal{H}', v) = 0$ for $v \in V(Y') \setminus V(Y)$.*

Proof. For each node h of H with $\operatorname{rad}_G(V_h) \leqslant R$, 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 $C \in \mathscr{C}(G - Y)$, a node h_C of H such that $N_G(C) \subseteq V_{h_C}$ and $\operatorname{rad}_G(V_{h_C}) \leqslant R$. Set $U := V(Y) \cup \bigcup_{C \in \mathscr{C}(G-Y)} (B_G(v_{h_C}, R) \cap V(C))$. 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 $V'_h := V_h \cup (B_G(v_h, R) \cap (\bigcup_{C \in \mathscr{C}(G-Y): h_C = h} V(C)))$ 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 5.3.1 and Lemma 5.3.4, 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 5.3.5, 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 $H^C \in \operatorname{Forb}_{\preccurlyeq}(X)$ 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 $\operatorname{irads}(\mathcal{H}^C, v) \leqslant f''_1(K)$ for all $v \in N_G(C)$.

Theorem 5.3.5. *If a 2-connected finite graph X has property $(*)$ (with respect to functions R, f'_0, f'_1, f''_1), then Conjecture 24' 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 $\operatorname{irads}(\mathcal{H}, v) \leqslant R'_1$ 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 5.2.2, 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 $C \in \mathscr{C}(G - Y^d)$, be given by property $(*)$ of X.

Then applying Construction 5.3.2 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 5.3.3, 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 5.3.4, $(H', \mathcal{V}')$ 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 5.3.1, 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. □

5.4 Proof of Theorem 26: Conjecture 24 for $X = K_4^-$

In this section we prove Theorem 26, which we restate here in the terminology of graph-decompositions.

Theorem 26'. *Every graph G with no K-fat K_4^- minor for $K \in \mathbb{N}_{\geqslant 1}$ 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 5.2.1 and 5.2.3, Theorem 26' immediately implies Theorem 26.

To prove Theorem 26', it suffices by Theorem 5.3.5 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 $\mathrm{irads}(\mathcal{H}^C, v) \leqslant 1$ 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 5.4.1 and Figure 5.1). 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

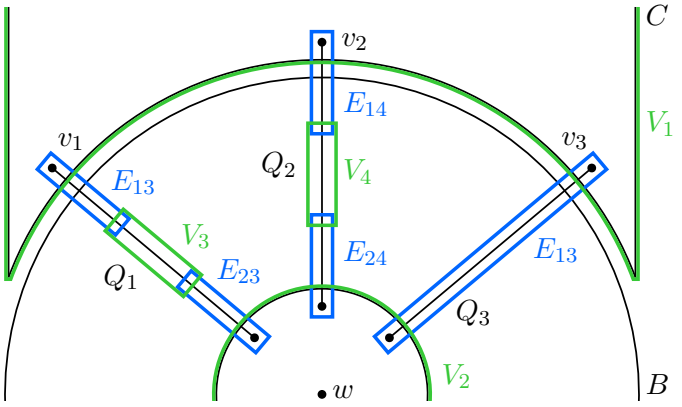

Figure 5.1: The K-fat K_4^- minor in the proof of Lemma 5.4.1.

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 5.4.3 and Figure 5.2).

Lemma 5.4.1. *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 $\mathrm{d}_G(u_i, u_j) \leqslant 2r$ for $i \neq j$ since B has radius r, and thus

$$r \geqslant \mathrm{d}_G(u_i, u_j)/2 \geqslant 7K/2 > 3K.$$

Hence, $B_G(w, r - (3K - 1))$ 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

$$V_1 := V(C),\ V_2 := B_G(w, r - (3K - 1)),\ V_3 := V(q_K^1 Q^1 q_{2K}^1),\ V_4 := V(q_K^2 Q^2 q_{2K}^2)$$

and the branch paths as

$$E_{13} := q_{2K}^1 Q^1 q_{3K}^1,\ E_{14} := q_{2K}^2 Q^2 q_{3K}^2,\ E_{23} := q_0^1 Q^1 q_K^1,\ E_{24} := q_0^2 Q^2 q_K^2,\ E_{12} := Q^3$$

we obtain a K-fat K_4^--model (see Figure 5.1). 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 $\mathrm{d}_G(V_1, V_2) \geqslant 3K \geqslant K$. In particular, as the Q^i are shortest paths, we find $\mathrm{d}_G(V_i, E_{(3-i)j}) \geqslant 2K \geqslant K$ as well as $\mathrm{d}_G(E_{1j}, E_{2j}) \geqslant K$ 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 $\mathrm{d}_G(Q^i, Q^j) \geqslant K$ for $i \neq j$. Indeed, we find

$$\mathrm{d}_G(Q^i, Q^j) \geqslant \mathrm{d}_G(u_i, u_j) - ||Q_i|| - ||Q_j|| \geqslant 7K - 2 \cdot 3K = K,$$

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. □

For the next lemma, we need the following result, which is immediate from [ADEFJKW23, Proof of Lemma 4.2].

Lemma 5.4.2. *Let G be a graph, and let P be a path that is a shortest path in G between its endvertices. 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 $2r+1$.[8] *Moreover, $V_p \subseteq B_G(p, 2r+1)$ for every $p \in V(P)$.*

Lemma 5.4.3. *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 $21K - 1 \leqslant \lfloor \mathrm{d}_G(v_1, v_2)/2 \rfloor - 1$ 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

$$\mathrm{d}_G(n_1, n_2) \geqslant \mathrm{d}_G(v_1, v_2) - \mathrm{d}_G(n_1, v_1) - \mathrm{d}_G(v_2, n_2) \geqslant (42K + 1) - 2 \cdot (7K - 1) = 28K + 3. \tag{5.1}$$

Further, since G has no K-fat K_4^- minor, Lemma 5.4.1 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 $N_{3-i} \subseteq V(C)$

[8] The radial spread is not stated explicitly in [ADEFJKW23, Lemma 4.2] but it follows easily from its proof that it is at most $4r + 2$. In fact, by optimising [ADEFJKW23, Lemma 3.10] for paths, it even follows that the radial spread is at most $2r + 1$.

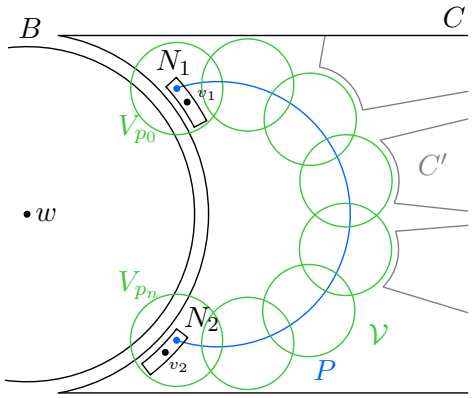

Figure 5.2: Depicted in green is the path-decomposition $(P, \mathcal{V})$ from Lemma 5.4.3 along the path P which is depicted in blue. By (ii), the neighbourhood of every component C' of $C - Y$ is contained in some bag in $\mathcal{V}$.

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

$$B_G(P, r_P) \subseteq B_G(C, r_P) = B_G(N_1, r_P) \cup V(C) \cup B_G(N_2, r_P) \subseteq B_1 \cup V(C) \cup B_2 = V(G'), \tag{5.2}$$

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,

$$B_{G'}(p, r_P) = B_G(p, r_P) \tag{5.3}$$

for all $p \in P$. So by Lemma 5.4.2, 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, $\mathrm{d}_G(p, v) \leqslant 28K + 1$ 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 $\mathrm{d}_G(u, p_0) \geqslant 7K$. Since also $\mathrm{d}_G(u, v_2), \mathrm{d}_G(p_0, v_2) \geqslant \mathrm{d}_G(N_1, N_2) \geqslant 7K$ by (5.1), Lemma 5.4.1

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 $\mathrm{d}_G(p_0, N_2), \mathrm{d}_G(p_n, N_1) \geqslant 28K+3$ by (5.1) but $\mathrm{d}_G(p, v) \leqslant 28K+1$ 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 $j_2 - j_1 > 2 \cdot (r_P + 1)$.

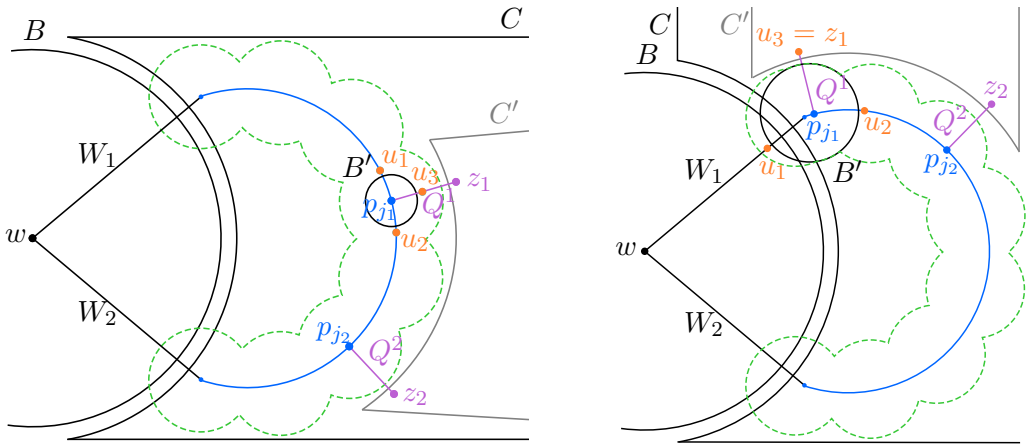

(a) Case 1: $j_1 \geqslant 7K$ and thus $u_1 \in V(P)$

(b) Case 2: $j_1 < 7K$ and thus $u_1 \in V(W_1)$

Figure 5.3: The situation in the proof of Lemma 5.4.3 if $(P, \mathcal{V})$ does not satisfy (ii).

To derive a contradiction, we find a K-fat K_4^- minor in G as follows. Let $Q^i = q_0^i \dots q_{r_P+1}^i$ be 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 $\mathrm{d}_G(v_1, v_2)/2 \geqslant 14K+1$.

Now first assume that $j_1 \geqslant 7K$ (see Figure 5.3a). 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 (5.3) 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 5.4.1 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 5.3b). 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 5.4.1, since u_1, u_2, u_3 lie in the same component of G'_B. Indeed, $Z := W^1u_1 \cup W^2 \cup u_2Pp_n \cup Q^2 \cup C'$ is connected and contains u_1, u_2, u_3, and one may follow a similar reasoning as above together with $\mathrm{d}_G(p_0, p_n) > 2 \cdot (14K + 1)$ by (5.1) to show that Z is also disjoint from B'. □

We are now in a position to prove Theorem 26'. For this, by Theorem 5.3.5, it suffices to show that K_4^- has property $(*)$.

Lemma 5.4.4. *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 $\leqslant 42K + 1$, 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 5.4.3 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 26'. By Theorem 5.3.5, Lemma 5.4.4 immediately yields Theorem 26'. □

Proof of Theorem 26. Let G be a graph with no K-fat K_4^- minor. By Lemma 5.2.3, Theorem 26' 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 5.2.1 that G is $(84K + 2, 3 \cdot (84K + 2)^2)$-quasi-isometric to H. Thus, Theorem 26 holds with $f(K) := (84K + 2, 3 \cdot (84K + 2)^2)$. □

5.5 Proof of Theorem 25: Conjecture 24 for $X = K_4$

In this section we prove Theorem 25, which we restate here in the terminology of graph-decompositions.

Theorem 25'. *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 5.2.1 and 5.2.3, Theorem 25' immediately implies Theorem 25.

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 $H + h_1h_2 \in \mathrm{Forb}_{\preccurlyeq}(K_4)$.[9] 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 5.5.1. *$\mathcal{H}_{SP}$ 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 $H \in \mathcal{H}_{SP}$. □

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 25'. For this, let $K \in \mathbb{N}_{\geqslant 1}$, and let G be a graph with no K-fat K_4 minor. By Theorem 5.3.5 (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 5.5.4). So let such a ball B and component

[9] 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 $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$.

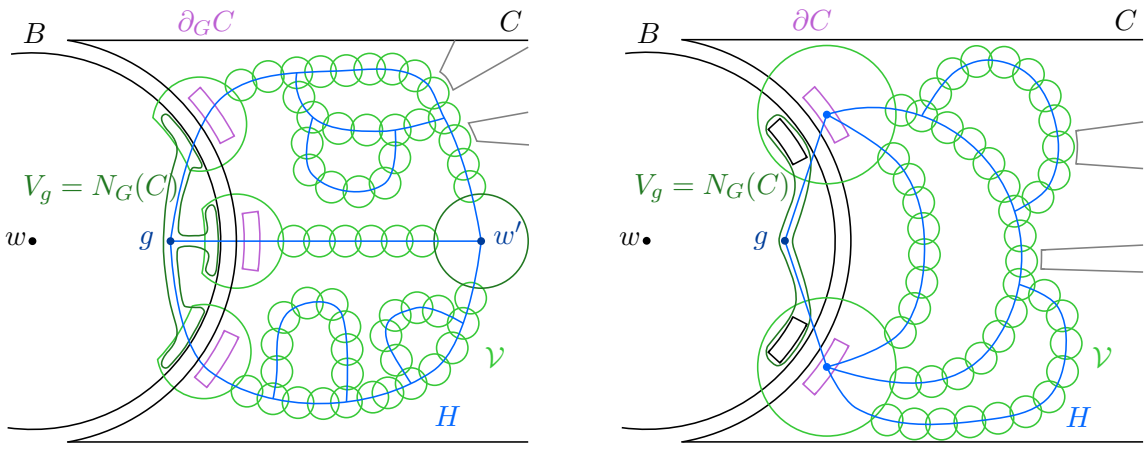

(a) Lemma 5.5.2

(b) Lemma 5.5.3

Figure 5.4: Depicted are the graph-decompositions $(H, \mathcal{V})$ with support Y in Lemmas 5.5.2 and 5.5.3 if G has no K-fat K_4 minor. By Lemma 5.5.2 (iv) and Lemma 5.5.3 (iv) the neighbourhood of every component of $C - Y$ is contained in some bag in $\mathcal{V}$.

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 partial decomposition with support $Y^C \subseteq G[C, 1]$ by applying the following Lemma 5.5.2 to C and $C^* := D$.

Lemma 5.5.2. *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 exists 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 $V_g := N_G(C)$ and $V_{g'} := (B_G(B, 22K + 1) \cap V(C)) \cup N_G(C)$. 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 5.5.3 (applied to $C := D$), and we can then identify the node of H^D given by Lemma 5.5.3 (i) with g'. This then yields the desired partial decomposition $\mathcal{H}^C$.

Lemma 5.5.3. *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 exists 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 25' from Lemmas 5.5.2 and 5.5.3. For this, as noted earlier, it suffices to prove that K_4 has property $(*)$.

Lemma 5.5.4. *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 $\mathscr{D}$ be the set of components of $C - B_G(B, 22K + 1)$, and first assume that there is some $D \in \mathscr{D}$ 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 5.5.2 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 $\mathscr{D}' \subseteq \mathscr{D}$ 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 $D \in \mathscr{D}'$ satisfies the premise of Lemma 5.5.3; let $(H^D, \mathcal{V}^D)$ be the partial decomposition with support Y^D obtained from applying Lemma 5.5.3 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 $D \in \mathscr{D} \setminus \mathscr{D}'$, 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 $D \in \mathscr{D}'$ 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 $h \in V(H) \setminus \{g, g'\}$ where D is the unique component in $\mathscr{D}'$ 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 $V_{g^D}^D \subseteq V_{g'}$ for all $D \in \mathscr{D}'$.

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 $D \in \mathscr{D} \setminus \mathscr{D}'$ and $\mathrm{rad}(V_h) \leqslant R_2(K)+22K+1$ for all other $h \in V(H)$ by (ii) of Lemma 5.5.3. Note that the radial spread of $(H, \mathcal{V})$ is $\leqslant 2 \cdot \max_{D \in \mathscr{D}} \mathrm{irads}((H^D, \mathcal{V}^D)) \leqslant 6 \leqslant f_1'(K)$, since $V_g \subseteq V_{g'}$, or it is 1 if $\mathscr{D} = \emptyset$.

By construction, it follows immediately from Lemma 5.5.3 (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 25'. By Theorem 5.3.5, Lemma 5.5.4 immediately yields Theorem 25'. □

Proof of Theorem 25. Let G be a graph with no K-fat K_4 minor. By Lemma 5.2.3, Theorem 25' 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 5.2.1 that G is $(50470K+142, 3 \cdot (50470K+142)^2)$-quasi-isometric to H. Thus, Theorem 26 holds with $f(K) := (50470K + 142, 3 \cdot (50470K + 142)^2)$. □

The remainder of this section is devoted to the proofs of Lemmas 5.5.2 and 5.5.3. For this, we first show in Section 5.5.1 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 5.5.2 and 5.5.3, we prove Lemmas 5.5.2 and 5.5.3, respectively.

5.5.1 Finding a K-fat K_4 minor

Lemma 5.5.5. *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 $q_K^2 Q^2 q_{n_2-K}^2 = Q^2 \cap G[P_2, K]$, 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 $w_{m-K} W' w_m \subseteq G[P_2, 2K]$.

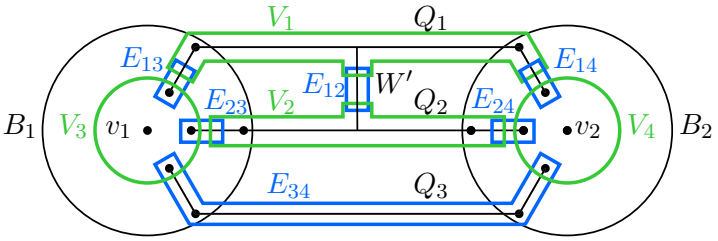

Figure 5.5: The K-fat K_4 minor in the proof of Lemma 5.5.5.

By defining the branch sets as

$$V_1 := V(q_K^1 Q^1 q_{n_1-K}^1) \cup V(w_0 W' w_K),\ V_2 := V(q_K^2 Q^2 q_{n_2-K}^2 w_{m-K} W' w_m),$$
$$V_3 := B_G(v_1, r_1 - 2K) \text{ and } V_4 := B_G(v_2, r_2 - 2K).$$

and the branch paths as

$$E_{12} := w_K W' w_{m-K},\ E_{i3} := q_0^i Q^i q_K^i,\ E_{i4} := q_{n_i-K}^i Q^i q_{n_i}^i \text{ and } E_{34} := Q^3$$

for $i \in [2]$ we obtain a K-fat K_4-model in G (see Figure 5.5). 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

$$\mathrm{d}_G(V_3, V_4), \mathrm{d}_G(E_{i3}, E_{j4}), \mathrm{d}_G(V_3, E_{i4}), \mathrm{d}_G(V_4, E_{i3}) \geqslant \mathrm{d}_G(B_1, B_2) \geqslant 5K \geqslant K$$

for $i, j \in [2]$. Moreover, $\mathrm{d}_G(V(Q^1) \cup V_1, V(Q^2) \cup V_2) \geqslant \mathrm{d}_G(P_1, P_2) - 2 \cdot 2K \geqslant 5K - 2 \cdot 2K = K$ since $V(Q^i) \cup V_i \subseteq B_G(P_i, 2K)$ for $i \in [2]$. This yields that

$$\mathrm{d}_G(V_1, V_2), \mathrm{d}_G(V_1, E_{2i}), \mathrm{d}_G(V_2, E_{1i}), \mathrm{d}_G(E_{1i}, E_{2i}) \geqslant K$$

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 $\mathrm{d}_G(P_3, P_i) \geqslant 5K$ for $i \in [2]$. Further, $\mathrm{d}_G(Q^3, W) \geqslant K$ by the choice of W. Thus,

$$\mathrm{d}_G(E_{34}, V_i), \mathrm{d}_G(E_{34}, E_{ij}), \mathrm{d}_G(E_{34}, E_{12}) \geqslant K$$

for $i \in [2]$ and $j \in \{3, 4\}$. As C is a component of $G - (B_1 \cup B_2)$, we have $\mathrm{d}_G(C, B_G(v_i, r_i - K)) = K + 1$ for $i \in [2]$. Since $E_{12} \subseteq W \subseteq C$ this implies that

$$\mathrm{d}_G(E_{12}, E_{ij}), \mathrm{d}_G(E_{12}, V_j) \geqslant \mathrm{d}_G(C, B_G(v_i, r_i - K)) = K + 1 \geqslant K$$

for $i \in [2]$ and $j \in \{3, 4\}$. Finally, we obtain

$$\mathrm{d}_G(V_i, V_j) \geqslant \mathrm{d}_G(B_G(C, K), B_G(v_{j-2}, r_{j-2} - 2K)) = 2K - K = K$$

for $i \in [2]$ and $j \in \{3, 4\}$. This completes the proof. □

5.5.2 Proof of Lemma 5.5.3

Let us briefly sketch the proof of Lemma 5.5.3. 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 $\mathrm{d}_G(B_1, B_2) \geqslant 5K + 2$. 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 $V^D_{h_i^D} = N_G(D) \cap B_i$ (Lemma 5.5.8). To obtain the desired partial graph-decomposition $(H, \mathcal{V})$ to satisfy the conclusion of Lemma 5.5.3, 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 5.5.1 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 5.5.8, we construct the partial graph-decompositions $(H^D, \mathcal{V}^D)$ recursively via Lemma 5.5.7 below. For this, we will need the following coarse version of Menger's theorem for two paths.

Theorem 5.5.6 ([AHJKW24], Theorem 2)**.** *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 exists $z \in V(Q)$ such that $B_G(z, 129d)$ intersects every $X - Y$ path.*[10]

Lemma 5.5.7. *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 $\mathrm{d}_G(P, G[C, 1]) \geqslant 5K$. 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 $B' := B_G(u, 130 \cdot 5K)$, for $i \in [2]$ and for every component C' of $G - (B_i \cup B')$ that is contained in C and attaches to B_i and B', there exists a B_i–B' path P' such that $d_G(P, G[C', 1]) \geqslant 5K$.

See Figure 5.6 for a depiction of the situation in Lemma 5.5.7.

Proof. Let Q be a shortest N_1–N_2 path through C. As $\mathrm{d}_G(v_1, v_2) \geqslant r_1 + r_2 + 5K + 2$, 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

[10] That the vertex z may be chosen on some shortest X–Y path is not stated in [AHJKW24, Theorem 1], 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 [AHJKW24, Theorem 1] if we increase the radius of the ball by a factor of 2.

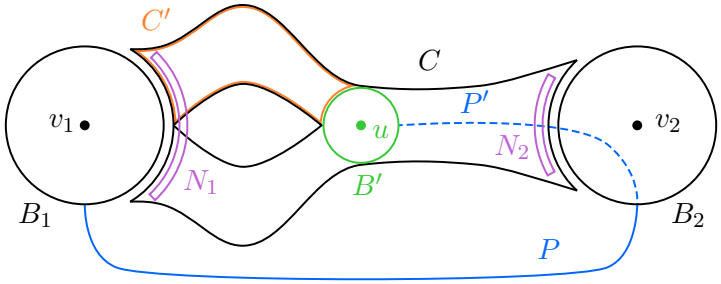

Figure 5.6: The situation in Lemma 5.5.7: A component C of $G - (B_1 \cup B_2)$ and the ball $B' \supseteq B$ around u. Depicted in blue is the path P and its dashed extension P', as required for the 'moreover' part applied to the orange component C'.

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 5.5.6 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 $\mathrm{d}_{G'}(Q_1, Q_2) \geqslant 5K$.

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, $\mathrm{d}_G(Q_1, Q_2) \geqslant 5K$ because $\mathrm{d}_{G'}(Q_1, Q_2) \geqslant 5K$ 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 $\mathrm{d}_G(Q_1 \cup Q_2, P) \geqslant \mathrm{d}_G(G[C, 1], P) \geqslant 5K$ 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 5.5.5 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 $\mathrm{d}_G(P', G[C', 1]) \geqslant 5K$. For this, let S be a shortest $G[C', 1]$–P' path in G. We show that S has length at least $5K$. Then $\mathrm{d}_G(G[C', 1], P') \geqslant ||S|| \geqslant 5K$, which yields the claim. Since $\mathrm{d}_G(G[C', 1], P) \geqslant \mathrm{d}_G(G[C, 1], P) \geqslant 5K$ 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, we have $\mathrm{d}_G(P_1 \cup P_2, B_G(w, 129 \cdot 5K)) \geqslant 5K$. 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 $\mathrm{d}_G(v_1, v_2) \geqslant 2r_1 + r_2 + 5K + 2$, we have $\mathrm{d}_G(B_1, B_2) \geqslant 5K + 2$, and thus S has length at least $5K$. □

Lemma 5.5.8. *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 $\mathrm{d}_G(P, G[C,1]) \geqslant 5K$.*

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) *$\mathrm{d}_G(v_i, v) \leqslant r_i + 2 \cdot R_0(K) + 5K + 1$ 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, $\mathrm{d}_H(h_i, h) \leqslant 3$ 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 5.5.8 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 5.5.8 for finite graphs. We may apply Lemma 5.5.7 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 5.5.8 for arbitrary graphs. We define the desired graph-decomposition $(H, \mathcal{V})$ recursively (see Figure 5.7 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 $\mathscr{C}^0 := \{C\}$, $w_1^C := h_1$ 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 $\mathscr{C}^i$ 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 $D \in \mathscr{C}^i$, an edge $w_1^D w_2^D$ of H^i,

(III) for every $D \in \mathscr{C}^{i-1} \setminus \mathscr{C}^i$, 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 $C_u \in \mathscr{C}^{i-1}$ 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 $D \in \mathscr{C}^i$ 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 $D \in \mathscr{C}^{i-1}$ we have $D \notin \mathscr{C}^i$ if and only if $\mathrm{d}_G(w_1^D, w_2^D) \geqslant r_{w_1^D} + r_{w_2^D} + 5K + 2$;

(v) for every $D \in \mathscr{C}^i \setminus \mathscr{C}^{i-1}$ and the unique component $D' \in \mathscr{C}^{i-1} \setminus \mathscr{C}^i$ 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 $D \in \mathscr{C}^i$ there exists a $B_G(w_1^D, r_{w_1^D})$–$B_G(w_2^D, r_{w_2^D})$ path P with $\mathrm{d}_G(P, G[D, 1]) \geqslant 5K$.

In the following, we will refer to components $D \in \mathscr{C}^n$ as in (iv), that is those which satisfy $\mathrm{d}_G(w_1^D, w_2^D) \geqslant r_{w_1^D} + r_{w_2^D} + 5K + 2$, as *long*.

We now obtain H^{n+1} as follows. For every long component $D \in \mathscr{C}^n$, let $u_D \in B_G(D, 1)$ be a vertex obtained from the application of Lemma 5.5.7 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 $D \in \mathscr{C}^n$ the vertex u_D and the edges $w_1^D u_D$ and $u_D w_2^D$.

Let the collection $\mathscr{C}^{n+1}$, 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 $D' \in \mathscr{C}^{n+1} \setminus \mathscr{C}^n$. Then D' is contained in a long component $D \in \mathscr{C}^n$. 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 5.5.7 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 $D \in \mathscr{C}^n$ 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 $D \in \mathscr{C}^n$ 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 5.5.1. This completes the verification that H^{n+1} is as desired.

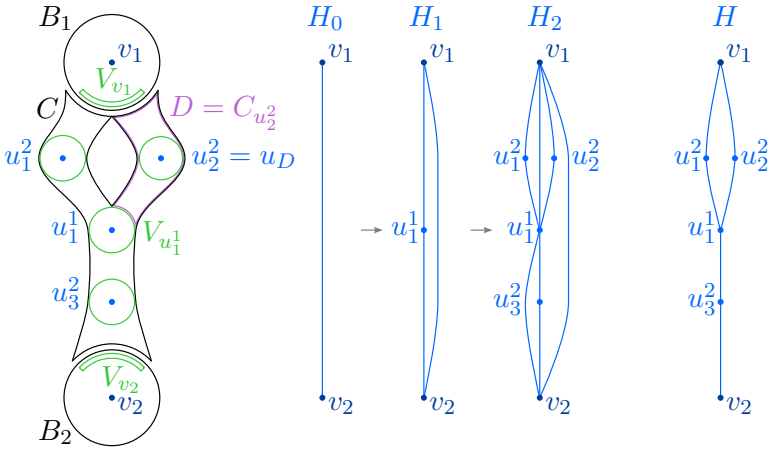

Figure 5.7: Depicted are the graphs H^0, H^1, H^2, H in blue and the bags V_u for the nodes u of H^2 for the component C of $G - (B_1 \cup B_2)$ in green. In this example, we have $w_1^D = v_1, w_2^D = u_1^1$ for the component $D = C_{u_2^2}$ in purple, and $H' = H^2$.

Let $H' := \bigcup_{i\in\mathbb{N}} H^i$ be the limit of the H^i. By Proposition 5.5.1, 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 $\mathrm{d}_G(u, u') \geqslant r_u + r_{u'} + 5K + 2$ and then subdividing each remaining edge precisely once. By Proposition 5.5.1, $H \in \mathcal{H}_{SP}$ 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 $C_{h_v} \in \mathscr{C}^{n-1}$. 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 $H'_v = H'[\{h \in V(H') \mid v \in V_h\}]$ is contained in $H'[h_v, 1]$.

For this, let $h' \in V(H') \setminus \{h_v\}$ with $v \in V_{h'}$ be given, and let $m \in \mathbb{N}$ be minimal such that $C_{h'} \in \mathscr{C}^m$. 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 $\mathscr{C}^\ell$ which contains $C_{h'}$; in particular, $C_{h'} = C^m$. Since $C_{h'} \in \mathscr{C}^m \setminus \mathscr{C}^{m-1}$, it follows from (I), (iv) and (V) that no C^ℓ is in $\mathscr{C}^{\ell-1}$. 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 $\mathrm{d}_G(h_v, h') \leqslant r_{h_v} + r_{h'}$ 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 $H[h_v, 3]$.[11] 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} and 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 $\mathrm{d}_G(v, u) \leqslant r_u$ 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 $d_G(v, w) \leqslant r_w + 2r_{w'} + 5K + 1$, 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, $C' \in \mathscr{C}^n$ 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 $\mathrm{d}_G(g, g') > r_g + r_{g'} + 5K + 2$, which by (iv) contradicts that $C' \in \mathscr{C}^{m+1}$. □

Proof of Lemma 5.5.3. 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 hence W 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 $\mathrm{d}_G(P, G[C, 1]) \geqslant R_1(K) - 1 \geqslant 5K$. Since also $\mathrm{d}_G(v_1, v_2) \geqslant 2R_1(K) + 5K + 1$ by assumption, we may apply Lemma 5.5.8 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 5.5.8 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 $Y := G[(\bigcup V(Y^{C'})) \cup ((B_1 \cup B_2) \cap$

[11] We remark that H_v is not neccessarily 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.

$V(C)) \cup N_G(C)]$. 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_1g and h_2g. 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 5.5.1; 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 $h \in V(H) \setminus \{g\}$, 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. □

5.5.3 Proof of Lemma 5.5.2

Let us briefly sketch the proof of Lemma 5.5.2. 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 5.5.10). 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 5.5.5. 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 5.5.8 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 5.5.11). 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 $G[C, 1]$.

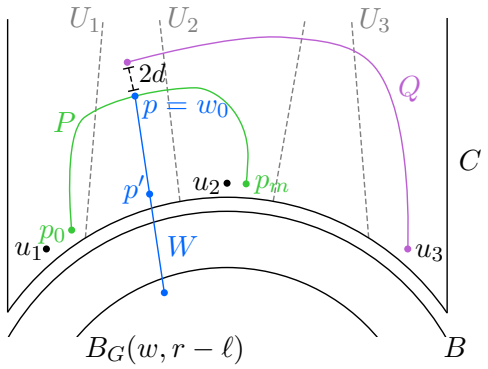

Figure 5.8: Setup in the beginning of the proof of Lemma 5.5.9

We first show that there exists a vertex w' as described above.

Lemma 5.5.9. *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 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 5.8). Since $p \notin U_3$, we may choose W so that $\mathrm{d}_G(p', u_3) \geqslant 2\ell + d + 2$. Moreover, as $\mathrm{d}_G(u_1, u_2) \geqslant 4 \cdot (2\ell + d + 2)$, we cannot have $\mathrm{d}_G(p', u_1), \mathrm{d}_G(p', u_2) < 2 \cdot (2\ell + d + 2)$; by symmetry, we may thus assume $\mathrm{d}_G(p', u_2) \geqslant 2 \cdot (2\ell + d + 2)$. In particular, this implies that

$$\mathrm{d}_G(p_m, p') \geqslant \mathrm{d}_G(p', u_2) - \mathrm{d}_G(u_2, p_m) \geqslant 2 \cdot (2\ell + d + 2) - (2\ell + d + 2) = 2\ell + d + 2. \tag{5.4}$$

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 5.9a).

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

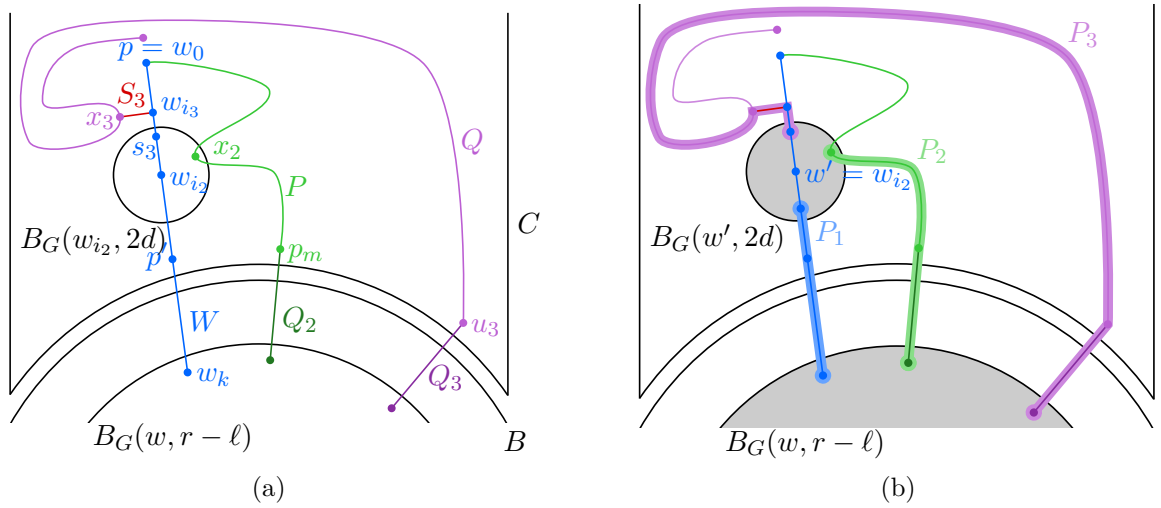

Figure 5.9: Choice of w' and construction of the paths P_1, P_2, P_3 in the proof of Lemma 5.5.9.

$S_3' := x_3 S_3 w_{i_3} W w_{i_2}$ that is contained in $B_G(w_{i_2}, 2d)$ (see Figure 5.9a). We now define our final three $B_G(w, r-\ell)$–$B_G(w', 2d)$ paths P_1, P_2, P_3 (see Figure 5.9b):

$$P_1 := w_{i_2+2d} W w_k,\ P_2 := x_2 P p_m Q_2 \text{ and } P_3 := s_3 \bar{S}_3' x_3 \bar{Q} u_3 Q_3.$$

We claim that the P_i are as desired, i.e. $\mathrm{d}_G(P_i, P_j) \geqslant d$ for every two distinct $i, j \in [3]$. First, we show $\mathrm{d}_G(P_1, P_2) \geqslant d$: By the choice of x_2, we have $\mathrm{d}_G(W, x_2 P p_m) \geqslant 2d \geqslant d$. 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

$$\mathrm{d}_G(p' W w_k, Q_2) \geqslant \mathrm{d}_G(p_m, p') - ||p' W w_k|| - ||Q_2|| \geqslant (2\ell + d + 2) - 2(\ell + 1) = d. \tag{5.5}$$

Since $\mathrm{d}_G(p_m, p') \geqslant 2\ell + d$ by (5.4), $\ell \geqslant d$, Q_2 is a shortest $\partial_G C$–$B_G(w, r-\ell)$ path and $p' \in \partial_G C$, we have $\mathrm{d}_G(Q_2, B_G(p', d) \cap V(W)) \geqslant d$. 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, $\mathrm{d}_G(W, Q_2) \geqslant d$, and hence $\mathrm{d}_G(P_1, P_2) \geqslant d$.

Next, we show $\mathrm{d}_G(P_1, P_3)$: Analogously to $\mathrm{d}_G(W, Q_2) \geqslant d$, we also find $\mathrm{d}_G(W, Q_3) \geqslant d$. Moreover, $\mathrm{d}_G(W, x_3 \bar{Q} u_3) \geqslant 2d \geqslant d$ by the choice of x_3. Thus, to see that $\mathrm{d}_G(P_1, P_3) \geqslant d$, it suffices to verify $\mathrm{d}_G(w_{i_2+2d} W w_k, S_3') \geqslant d$: Since W is a shortest path in G and $i_3 \leqslant i_2$, we have $\mathrm{d}_G(w_{i_2+2d} W w_k, w_{i_3} W w_{i_2}) = \mathrm{d}_G(w_{i_2+2d}, w_{i_2}) = 2d$. Moreover, by the choice of $x_3 \in V(Q)$ and S_3, the x_3–w_{i_3} path S_3 has length $||S_3|| = \mathrm{d}_G(W, x_3) = 2d$. Thus, the distance from the first half of S_3 to $w_{i_2+2d} W w_k$ is $\geqslant \mathrm{d}_G(x_3, W) - ||S_3||/2 = 2d - d = d$, and the distance from the second half of S_3 to $w_{i_2+2d} W w_k$ is $\geqslant \mathrm{d}_G(w_{i_3}, w_{i_2+2d} W w_k) - ||S_3||/2 \geqslant 2d - d = d$.

Finally, we show $\mathrm{d}_G(P_2, P_3) \geqslant d$: By the choice of Q, we have $\mathrm{d}_G(P, Q) \geqslant 2d$. Similarly to the previous argument, we find $\mathrm{d}_G(x_2 P p_m, S_3') \geqslant d$, as $\mathrm{d}_G(W, x_2 P p_m), \mathrm{d}_G(P, Q) \geqslant 2d$, and thus the distance from the first half of S_3 to $x_2 P p_m$ is $\geqslant \mathrm{d}_G(x_3, P) - ||S_3||/2 \geqslant 2d - d = d$, and the distance from the second half of S_3 to $x_2 P p_m$ is $\geqslant \mathrm{d}_G(w_{i_3}, x_2 P p_m) - ||S_3||/2 \geqslant 2d - d = d$.

Moreover, analogously to (5.5), we find

$$\mathrm{d}_G(Q_2, Q_3) \geqslant \mathrm{d}_G(p_m, u_3) - ||Q_2|| - ||Q_3|| \geqslant \mathrm{d}_G(U_2 \cap \partial_G C, u_3) - 2(\ell+1) \geqslant (2\ell+d+2) - 2(\ell+1) = d.$$

Further, $\mathrm{d}_G(x_2 P p_m, Q_3), \mathrm{d}_G(u_3 Q x_3, Q_2) \geqslant d$ since $\mathrm{d}_G(P, u_3), \mathrm{d}_G(Q, p_m) \geqslant \mathrm{d}_G(P, Q) \geqslant 2d$ 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 $\mathrm{d}_G(S_3, Q_2) \geqslant d$: For this, it suffices to prove $\mathrm{d}_G(x_3, Q_2) \geqslant 3d$, as S_3 has length $2d$ and starts in x_3. If $\mathrm{d}_G(x_3, B) > 3d$, then $\mathrm{d}_G(x_3, Q_2) \geqslant 3d$ follows immediately, as $Q_2 \subseteq G[B, 1]$. Thus, we may assume $\mathrm{d}_G(x_3, B) \leqslant 3d$. Suppose for a contradiction that $\mathrm{d}_G(x_3, Q_2) < 3d$. Then $\mathrm{d}_G(x_3, p_m) \leqslant 2\,\mathrm{d}_G(x_3, Q_2) < 6d$, since $x_3 \in V(C)$ and Q_2 is a shortest p_m–$B_G(w, r-\ell)$ path. Also $\mathrm{d}_G(w_{i_3}, p') \leqslant \mathrm{d}_G(w_{i_3}, x_3) + \mathrm{d}_G(x_3, B) \leqslant 5d$, since W is a shortest path. Thus, $\mathrm{d}_G(p_m, p') \leqslant \mathrm{d}_G(p_m, x_3) + \mathrm{d}_G(x_3, w_{i_3}) + \mathrm{d}_G(w_{i_3}, p') < 13d$, which contradicts $\mathrm{d}_G(p_m, p') \geqslant 2\ell + d$ as $\ell \geqslant 6d$. □

Corollary 5.5.10. *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 $\mathrm{d}_G(P, G[D, 1]) \geqslant 5K$.

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 5.5.9 (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 $\mathrm{d}_G(P_i, P_j) \geqslant 11K$ 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 $\mathrm{d}_G(P_i, G[D, 1]) \geqslant 5K$. Towards a contradiction, suppose that $\mathrm{d}_G(P_i, D) \leqslant 5K$ 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 $\mathrm{d}_G(P_i', P_j') \geqslant 11K - 2 \cdot 3K = 5K$, which by Lemma 5.5.5 ($\mathrm{d}_G(w, w') > r_1 + r_2 + 5K + 1$ with $r_1 = r - \ell - 3K, r_2 = 19K$, since $\mathrm{d}_G(w', w) \geqslant r$ and $\ell \geqslant 23K$) implies that G has a K-fat K_4 minor, which contradicts our assumption. □

Lemma 5.5.11. *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 $\mathrm{d}_G(P, G[D, 1]) \geqslant 5K$.*

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) $\mathrm{rad}_G(V_h) \leqslant R'_0(K)$ *for all nodes* $h \neq h_2$ *of* H,

(γ) $\mathrm{rad}_G(V_{h_2}) \leqslant r_2 + 2R'_0(K) + 1$,

(δ) $\mathrm{irads}((H, \mathcal{V})) \leqslant 3$, *and*

(ε) *for every component* D' *of* $G - Y$ *that meets* C *and* D, *we have* $\mathrm{rad}_G(N_G(D')) \leqslant R_2(K) - 1$ *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 5.5.8 to $v_1, v_2, r_1, r'_2 := r_2 - \ell(K) - 1, K$ 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 5.5.8.

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 $B_1 \setminus V_{h_1}^1$ 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 $N_G(D) \cap B_1 \subseteq V_{h_1}^2$ 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 $\mathrm{d}_G(v_1, v_2) \geqslant r_1 + 2 + r_2$.

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 $\mathrm{d}_G(v, v_2) \leqslant r_2 \leqslant r_2 + 2R'_0(K) + 1$ 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 $\mathrm{d}_G(v, v_2) \leqslant r'_2 + 2R_0(K) + 5K + 1$ by (c). Otherwise, by (b), we have $\mathrm{d}_G(v_2, v) \leqslant r_2 + 1 + 2\,\mathrm{rad}_G(V^1_h) \leqslant r_2 + 1 + 2R'_0(K)$ 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

$$\mathrm{d}_G(q, v_2) \geqslant \mathrm{d}_G(C, v_2) > r_2 \geqslant r'_2 + \ell(K) = r'_2 + 2 \cdot R_0(K) + 5K + 2,$$

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

$$N_G(D') = (N_G(D') \cap V(Y^1)) \cup (N_G(D') \cap B_2) \subseteq N_G(\hat{D}) \cup (N_G(D') \cap B_2) \tag{5.6}$$

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 $\ell(K)$ to V^1_h. Since also $N_G(\hat{D}) \subseteq V^1_h$, we have by (5.6) for all $u, w \in N_G(D')$ that

$$\mathrm{d}_G(u, w) \leqslant \mathrm{d}_G(u, V^1_h) + 2 \cdot \mathrm{rad}_G(V^1_h) + \mathrm{d}_G(V^1_h, w) \leqslant \ell(K) + 2 \cdot R'_0(K) + \ell(K) \leqslant R_2(K) - 1.$$

This shows that $(H, \mathcal{V})$ satisfies (ε) and thus concludes the proof. □

Proof of Lemma 5.5.2. 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 5.5.10 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 $\mathrm{d}_G(G[D, 1], P) \geqslant 5K$. Let $\mathscr{D}$ 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 5.5.11 to $v_1 := w'$, $v_2 := w$, $r_1 := 22K$, $r_2 := r$ and every component $D \in \mathscr{D}$. 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 $Y^1 := G[\bigcup_{D\in\mathscr{D}} V(Y^D)] \subseteq G[B_1 \cup (\bigcup_{D\in\mathscr{D}} V(D)) \cup B_2] \subseteq G[V(C) \cup B_2]$. 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_2g. Then we assign the nodes of H^1 bags $V_g^1 := N_G(C)$, $V_{h_2}^1 := \bigcup_{D\in\mathscr{D}} V_{h_2^D}^D \supseteq B_2$, and $V_{h_1}^1 := \bigcup_{D\in\mathscr{D}} V_{h_1^D}^D = B_1$ and $V_h^1 := V_h^D$ for all other $h \in V(H^1)$ where D is the unique component in $\mathscr{D}$ 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 $D \in \mathscr{D}$ 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 $\mathscr{C}$ 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 $C' \in \mathscr{C}$ we have that $D_{C'} \in \mathscr{D}$ 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 $\mathrm{rad}_G(N_G(C')) \leqslant R_2(K) - 1$.

Now we obtain H from H^1 by adding for each component $C' \in \mathscr{C}$ 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 $C' \in \mathscr{C}$. 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 $Y^2 := G[V(Y^1) \cup (\bigcup_{C'\in\mathscr{C}} \partial_G(C'))]$. 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 $C' \in \mathscr{C}$, 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_2g with terminals h_2 and g; so H^1 is again in $\mathcal{H}_{SP}$ with terminals h_1 and g by Proposition 5.5.1. 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 $C' \in \mathscr{C}'$, 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. $C' \in \mathscr{C}$. 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 $Y^2 = G[V(Y^1) \cup \bigcup_{C' \in \mathscr{C}} \partial_G C']$. 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 $\mathrm{rad}_G(V_{h'_{C'}}) \leqslant R_2(K)$. □

CHAPTER 6

A Menger-type theorem for two induced paths

This chapter is based on [AHJKW24], which is joint work with Sandra Albrechtsen, Tony Huynh, Raphael W. Jacobs, Paul Wollan.

6.1 Introduction

All graphs in this chapter are finite and simple.

Let G be a graph, and $X, Y \subseteq V(G)$. An $X - Y$ *path* is a path with one end in X, one end in Y, and no internal vertex in $X \cup Y$. Note that by the definition, a single vertex of $X \cap Y$ is an $X - Y$ path. If either X or Y is empty, we will use the convention that no $X - Y$ path exists. If v is a single vertex in G, we will refer to a $v - X$ path as shorthand for a $\{v\} - X$ path. Let $r \geqslant 0$ be an integer, and let $u, v \in V(G)$. The *distance between u and v*, denoted $\mathrm{d}_G(u, v)$, is the length of a shortest $u - v$ path in G. The *ball of radius r around v* is $B_G(v, r) := \{u \in V(G) : \mathrm{d}_G(u, v) \leqslant r\}$. More generally, if H is a subgraph of G, we define $B_G(H, r) = \bigcup_{v \in V(H)} B_G(v, r)$. For two subgraphs H_1, H_2 of G, the *distance between H_1 and H_2*, denoted $\mathrm{d}_G(H_1, H_2)$, is the minimum length of a $V(H_1) - V(H_2)$ path.

Given a graph G and $X, Y \subseteq V(G)$, Menger's classic result says that the maximum number of pairwise vertex disjoint $X - Y$ paths is equal to the minimum size of a set Z of vertices intersecting every $X - Y$ path [Men27]. 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 end in $V(H_1)$ and one end in $V(H_2)$. No analog of Menger's theorem is known for pairwise anti-complete paths linking two sets of vertices. Such an exact characterization 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 [Bie91; Bie92]. 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 $|V(P_1)| \cup |V(P_2)|$ minimal).

We will show the following weak characterization of the induced two paths problem.

Theorem 6.1.1. *There exists a constant c 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.*

Theorem 6.1.1 immediately follows from a more general result where we ask that the two paths be at distance at least d.

Theorem 6.1.2. *Let $c = 129$. For all graphs G, $X, Y \subseteq V(G)$, and integers $d \geqslant 1$, either there exist two disjoint $X - Y$ paths P_1, P_2 such that $\mathrm{d}_G(P_1, P_2) \geqslant d$ or there exists $z \in V(G)$ such that $B_G(z, cd)$ intersects every $X - Y$ path.*

Independently, Georgakopoulos and Papasoglu have proved Theorem 6.1.2 (with a slightly worse constant of 272) [GP23].

We conjecture that Theorem 6.1.2 should generalize to an arbitrary number of $X - Y$ paths at pairwise distance d.

Conjecture 6.1.3. *There exists a constant c satisfying the following. For all graphs G, $X, Y \subseteq V(G)$, and integers $d, k \geqslant 1$, either there exist k disjoint $X - Y$ paths $P_1, \ldots, P_k$ such that $\mathrm{d}_G(P_i, P_j) \geqslant d$ 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.*

Rose McCarty and Paul Seymour proved that the $d = 3$ case of Conjecture 6.1.3 actually implies the general case.

Theorem 6.1.4 (McCarty and Seymour, 2023)**.** *If Conjecture 6.1.3 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 $\mathrm{d}_G(u, v) \leqslant d$. 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 $\mathrm{d}_H(P_i, P_j) \geqslant 3$ 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 $\mathrm{d}_G(P_i^*, P_j^*) \geqslant d$ for all distinct i, j. Towards a contradiction, suppose $\mathrm{d}_G(P_i^*, P_j^*) \leqslant d$ 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, $\mathrm{d}_H(P_i, P_j) \leqslant 2$, which is a contradiction. □

We would like to point out that Nguyen, Scott, and Seymour [NSS24] have recently found a counterexample to Conjecture 6.1.3.

Our proof shows that we may take $c = 129$ in Theorem 6.1.2. On the other hand, the following lemma shows that c must be at least $3/2$.

Lemma 6.1.5. *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 $\mathrm{dist}_G(P_1, P_2) \geqslant d$, 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 6.1. 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 $\mathrm{d}_H(P_1, P_2) \geqslant 2$. 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 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

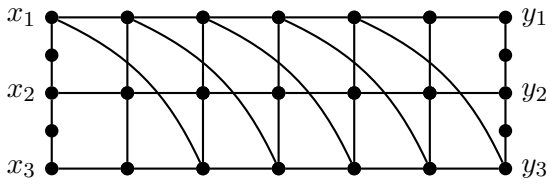

Figure 6.1: A graph without two disjoint anti-complete $\{x_1, x_2, x_3\} - \{y_1, y_2, y_3\}$ paths and no ball of radius one hitting all $X - Y$ paths.

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. □

We conclude this section with some necessary notation.

Let G be a graph and H_1, H_2 two subgraphs 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 $E(H_1) \cup E(H_2)$. An *H_1-path* is a path with both ends in H_1 and no internal vertex or edge in H_1. Let $X \subseteq V(G)$. An *X-path* is any J-path where J is the subgraph of G with vertex set X and no edges. For a path $P = v_0 v_1 \dots v_n$ and $0 \leqslant i \leqslant j \leqslant n$, we write $v_i P := v_i v_{i+1} \dots v_n$, $P v_i := v_0 v_1 \dots v_i$ and $v_i P v_j := v_i v_{i+1} \dots v_j$ for the corresponding subpaths of P.

6.2 Interlaced systems of intervals

In this section we prove a weak characterization of when a sequence of intervals contains an "interlaced" subsequence. We begin with the definition of interlaced. See Figure 6.2 for a visualisation.

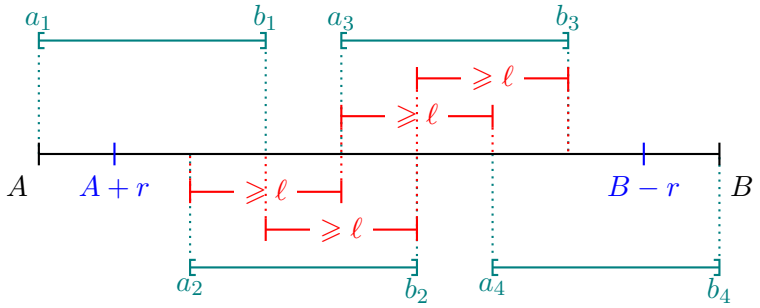

Figure 6.2: A clean (A, B, r)-system $([a_i, b_i])_{i \in [4]}$ indicated in teal which is interlaced with buffer ℓ.

Definition 6.2.1. 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[1] such that for all $i \in [t]$,

[1]We will in fact only consider the integers in such intervals.

- $a_i, b_i \in \mathbb{Z}$,
- $[a_i, b_i] \subseteq [A, B]$,
- 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 6.3).

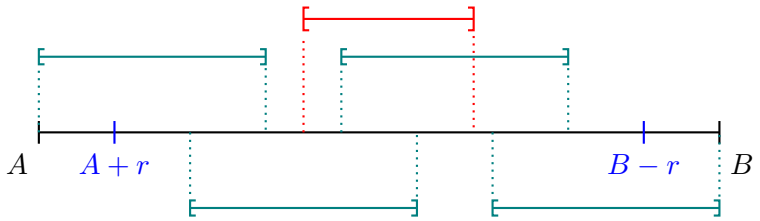

Figure 6.3: The (A, B, r)-system indicated in teal is clean. However, adding the red interval gives an (A, B, r)-system which is not clean.

Lemma 6.2.2. *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. Note that the lemma trivially holds if $t = 1$. Assume that $[a_1, b_1], \ldots, [a_t, b_t]$ 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, $[a_1, b_1], \ldots, [a_i, b_i], [a_j, b_j], \ldots, [a_t, b_t]$ 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. □

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 6.2.3. *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. □

6.3 Proof of Theorem 6.1.1

We now prove our main theorem, which we restate for convenience.

Theorem 6.1.2. *Let $c = 129$. For all graphs G, $X, Y \subseteq V(G)$, and integers $d \geqslant 1$, either there exist two disjoint $X - Y$ paths P_1, P_2 such that $\mathrm{d}_G(P_1, P_2) \geqslant d$ 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 $\mathrm{d}_G(P_1, P_2) \geqslant d_1 \geqslant d$. 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 $\mathrm{d}_G(v_j, H) = d_1 + 1$. 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 $\mathrm{d}_G(v_j, H) = d_1 + 1$. For each $H \in \mathcal{H}$, we define $I(H)$ to be the interval $[a(H), b(H)]$ (see Figure 6.4). Let $\mathcal{H}' \subseteq \mathcal{H}$ be the components which correspond to the maximal intervals in $I(\mathcal{H}) := \{I(H) : H \in \mathcal{H}\}$. Observe that $I(\mathcal{H}') := \{I(H) : H \in \mathcal{H}'\}$ is a $(-d_2, d_2 + n, d_2)$-system.

By Lemma 6.2.3, 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 exists 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 $Z = [j-2d_2, j+2d_2] \cap (-d_2, d_2+n)$.

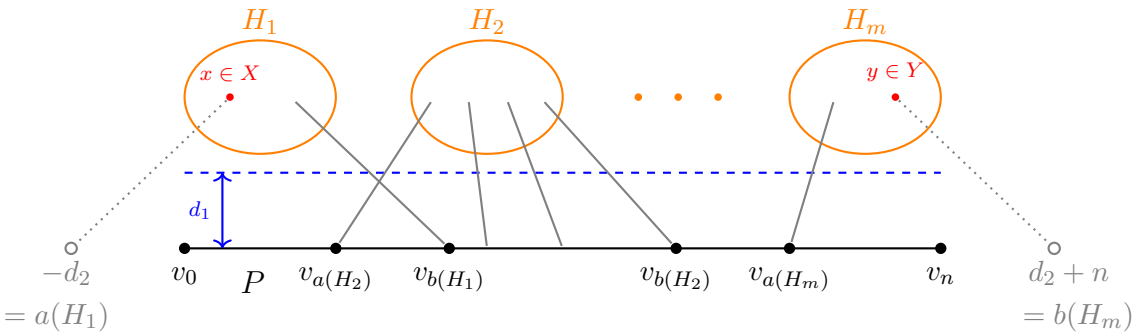

Figure 6.4: The components H_i along with their respective $a(H_i)$ and $b(H_i)$ represented by the corresponding vertex in P or a vertex in X or Y.

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 $\mathrm{d}_G(v_h, H) = d_1+1$, 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, $\mathrm{d}_G(v_k, v_\ell) = |k-\ell|$ for all $0 \leqslant k, \ell \leqslant n$. Hence, $\mathrm{d}_G(R_1, R_2) = 4d_2+2$. Since an edge in G from $B_G(R_1, d_1)$ to $B_G(R_2, d_1)$ would witness that $\mathrm{d}_G(R_1, R_2) \leqslant 2d_1+1$ 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 end 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 6.2.2, 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. $\mathrm{d}_G(v_{b(H'_j)}, v_{a(H'_{j+1})}) \geqslant 30d$ *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. $\mathrm{d}_G(v_{a(H'_j)}, v_{a(H'_{j+2})}) \geqslant d_2$ *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* $\mathrm{d}_G(H'_p, H'_q) \geqslant d_2 - 2d_1 + 2 - 30d$. *Moreover, if* $q \geqslant p+7$, *then* $\mathrm{d}_G(H'_p, H'_q) \geqslant d_2 - 2d_1 + 2$.

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 ends 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

$$\mathrm{d}_G(v_\ell, v_r) \leqslant \mathrm{d}_G(v_\ell, h'_p) + \mathrm{d}_G(h'_p, h'_q) + \mathrm{d}_G(h'_q, v_r) = 2d_1 + (||Q|| - 2).$$

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 $\mathrm{d}_G(v_\ell, v_r) = r - \ell \geqslant d_2$. Together with $\mathrm{d}_G(v_\ell, v_r) \leqslant 2d_1 + (||Q|| - 2)$, 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 $\mathrm{d}_G(v_{b(H'_{p+1})}, v_{a(H'_{p+3})}) \leqslant 30d$. Therefore, $\mathrm{d}_G(v_\ell, v_r) \geqslant d_2 - 30d$. On the other hand, $\mathrm{d}_G(v_\ell, v_r) \leqslant 2d_1 + (||Q|| - 2)$, as above. Thus, $||Q|| \geqslant d_2 - 2d_1 + 2 - 30d$, as required. ■

Let $H' = \bigcup_{i=1}^{m+s} H'_i$. A *fruit tree* is a subgraph $\mathcal{W} := \bigcup_{i \in I} H'_i \cup W_1 \cup \cdots \cup W_k$ of G (see Figure 6.5) 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 ends 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 ends 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 $v \in V(C)$, $\mathrm{d}_{\mathcal{W}}(v, V(H')) \leqslant \ell d$.*

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 ends 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 ends 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 end 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, $\mathrm{d}_{C_1}(w_1, V(H')) \leqslant \ell_1 d$ and $\mathrm{d}_{C_2}(w_2, V(H')) \leqslant \ell_2 d$. Therefore, we are done unless $\mathrm{d}_{W_{i_\ell}}(v, w_1) > (\ell - \ell_1)d$ and $\mathrm{d}_{W_{i_\ell}}(v, w_2) > (\ell - \ell_2)d$. 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 $W_1, \dots, W_k$. 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 $W_{i_1}, \dots, W_{i_\ell}$ with $\ell \geqslant 5$ and increasing $i_j \in [k]$. Let $\ell' \in [\ell]$ be the smallest index such that $W_{i_1} \cup \dots \cup W_{i_{\ell'}}$ 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 ends 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

$$\mathrm{d}_C(u_p, u_q) \leqslant \sum_{i=1}^{c_1}(i + 1)d + (4 + 2)d + \sum_{i=1}^{c_2}(i + 1)d$$

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

$$\mathrm{d}_C(u_p, u_q) \leqslant \sum_{i=1}^{4}(i+1)d + (4+2)d + \sum_{i=1}^{4}(i+1)d = 34d < d_2 - 2d_1 + 2.$$

However, this contradicts Claim 3.

If $c \in \{5, 6\}$, then it follows from $c_1 + c_2 + 1 = c \leqslant 6$ that

$$\mathrm{d}_C(u_p, u_q) \leqslant \sum_{i=1}^{4}(i+1)d + (4+2)d + \sum_{i=1}^{1}(i+1)d = 22d < d_2 - 2d_1 + 2 - 30d.$$

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 define $\min(\mathcal{W}) := \min\ I = \min\{i \in [m+s] : H'_i \subseteq \mathcal{W}\}$ 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 end of B_i in $\mathcal{W}^i$ and ℓ_{i+1} be the end 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 6.5.

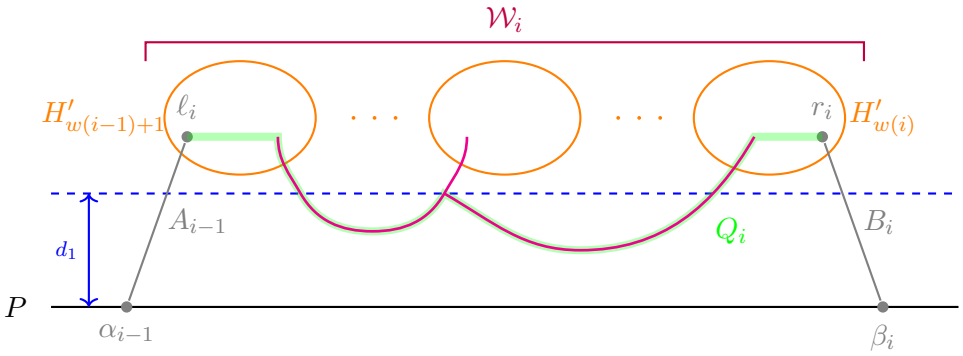

Figure 6.5: The fruit tree $\mathcal{W}_i$ of the orchard with its corresponding components (orange) and its composite paths (magenta).

We will now recursively define two $X - Y$ walks P_1 and P_2 such that $\mathrm{d}_G(P_1, P_2) \geqslant d$. These walks then contain the desired paths. In order to show that our construction indeed yields $\mathrm{d}_G(P_1, P_2) \geqslant d$, we will need to ensure that $\mathrm{d}_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

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 end of this path),

(iv) the other end of the component P' is some β_j with $j < i$,

(v) either α_i occurs after β_j along P, or $\mathrm{d}_P(\alpha_i, \beta_j) \leqslant 20d$ and if additionally $\mathrm{d}_P(\alpha_i, \beta_j) > 10d$, then α_{i+1} occurs between α_i and β_j along P, and

(vi) $\mathrm{d}_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

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 end 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 $\mathrm{d}_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$ 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 6.6).

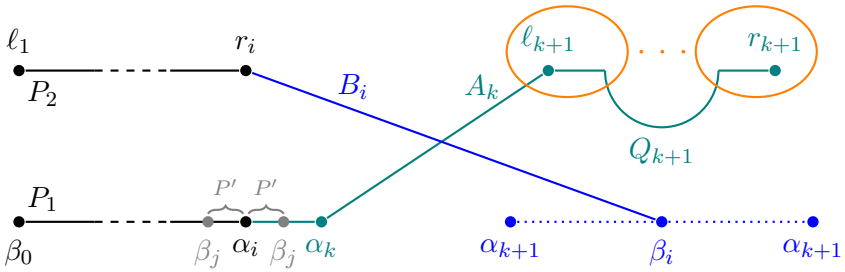

Figure 6.6: The old paths P_1 and P_2 in black together with their updates in teal and blue, respectively. Note that α_{k+1} may occur before or after β_i and that β_j may occur before or after α_i.

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 ends 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 $\mathrm{d}_P(\beta_i, \{\alpha_i, \beta_j\}) \geqslant 30d - 20d = 10d$ by assumption and Claim 1, it suffices to ensure that $\mathrm{d}_P(\beta_i, \alpha_k) \geqslant 10d$ if $k \neq i$ and that $\mathrm{d}_P(\alpha_{k+1}, \{\alpha_k, \beta_j\}) \geqslant 10d$ 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 6.7 to 6.10. 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 $\mathrm{d}_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant \mathrm{d}_P(\beta_i, \{\alpha_i, \beta_j\}) \geqslant 10d$, and we set $k := i$ (see Figure 6.7). By assumption, (v) and (vi) are clear.

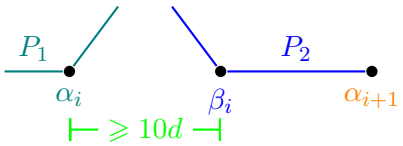

Figure 6.7: α_{i+1} occurs after β_i along P.

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 6.8). In this case, we proceed as follows. If $\mathrm{d}_P(\alpha_{i+1}, \beta_i) \geqslant 10d$, then we set $k := i+1$ (see Figure 6.8a). Again, (v) and (vi) are clear from the assumption. We remark that by assumption this in particular is the case if $\mathrm{d}_P(\alpha_i, \beta_j) > 10d$ and β_j occurs after α_i along P, since then α_{i+1} occurs between α_i and β_j along P and, as shown above, $\mathrm{d}_P(\beta_i, \beta_j) \geqslant 10d$.

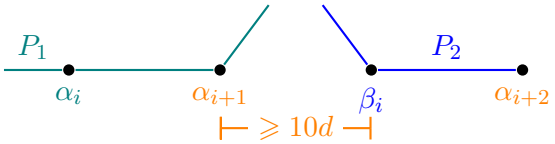

(a) $\mathrm{d}_P(\alpha_{i+1}, \beta_i) \geqslant 10d$

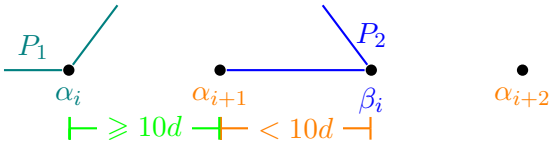

(b) $\mathrm{d}_P(\alpha_{i+1}, \beta_i) < 10d$

Figure 6.8: α_{i+1} occurs before β_i along P and α_{i+2} occurs after β_i along P.

Otherwise, $\mathrm{d}_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant 10d$. Indeed, either β_j occurs before α_i along P and thus we have $\mathrm{d}_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) = \mathrm{d}_P(\alpha_{i+1}, \alpha_i) = \mathrm{d}_P(\alpha_i, \beta_i) - \mathrm{d}_P(\alpha_{i+1}, \beta_i) \geqslant 30d - 10d \geqslant 10d$ by Claim 1, or β_j occurs after α_i along P and thus we have $\mathrm{d}_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) = \mathrm{d}_P(\alpha_{i+1}, \beta_j) = \mathrm{d}_P(\alpha_i, \beta_i) - \mathrm{d}_P(\alpha_{i+1}, \beta_i) - \mathrm{d}_P(\alpha_i, \beta_j) \geqslant 30d - 10d - 10d \geqslant 10d$. In this case we set $k := i$ (see Figure 6.8b). Then (vi) holds by the previous calculation. Moreover, recall that $\mathrm{d}_P(\alpha_{k+1}, \beta_i) < 10d$ 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 6.9). In this case, we proceed as follows. If $\mathrm{d}_P(\alpha_{i+2}, \beta_i) \geqslant 10d$, then we set $k := i + 2$ (see Figure 6.9a). Then (v) and (vi) are clear from the assumption.

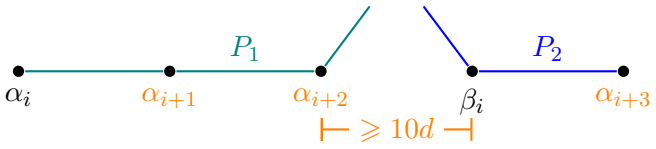

(a) $\mathrm{d}_P(\alpha_{i+2}, \beta_i) \geqslant 10d$

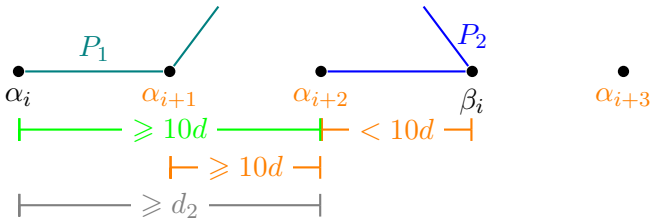

(b) $\mathrm{d}_P(\alpha_{i+2}, \beta_i) < 10d$ and $\mathrm{d}_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant 10d$

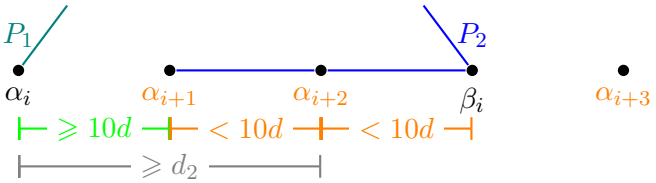

(c) $\mathrm{d}_P(\alpha_{i+2}, \beta_i) < 10d$ and $\mathrm{d}_P(\alpha_{i+1}, \alpha_{i+2}) < 10d$

Figure 6.9: α_{i+1} and α_{i+2} occur before β_i along P and α_{i+3} occurs after β_i along P.

Otherwise, if $\mathrm{d}_P(\alpha_{i+2}, \beta_i) < 10d$, we have $\mathrm{d}_P(\alpha_{i+2}, \{\alpha_i, \beta_j\}) \geqslant \mathrm{d}_P(\alpha_{i+2}, \alpha_i) - 20d \geqslant d_2 - 20d \geqslant 10d$ by assumption and Claim 2. If $\mathrm{d}_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant 10d$, then we set $k := i + 1$ (see Figure 6.9b). Then (vi) holds by assumption. Since also $\mathrm{d}_P(\alpha_{k+1}, \beta_i) < 10d$ by assumption, (v) holds too.

Otherwise, $\mathrm{d}_P(\alpha_{i+1}, \beta_i) = \mathrm{d}_P(\alpha_{i+1}, \alpha_{i+2}) + \mathrm{d}_P(\alpha_{i+2}, \beta_i) \leqslant 20d$ and $\mathrm{d}_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant \mathrm{d}_P(\alpha_i, \alpha_{i+2}) - \mathrm{d}_P(\alpha_{i+1}, \alpha_{i+2}) - 20d \geqslant d_2 - 10d - 20d \geqslant 10d$ by assumption and Claim 2, and we then set $k := i$ (see Figure 6.9c). Then (vi) holds by the previous calculation. Moreover, (v) 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 $\mathrm{d}_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant d_2$. 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 $\mathrm{d}_P(\alpha_{i+1}, \alpha_{i+2}) = \mathrm{d}_P(v_{a(H'_{w(i)+2})}, v_{a(H'_{w(i)+3})}) \geqslant d_2$ 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.

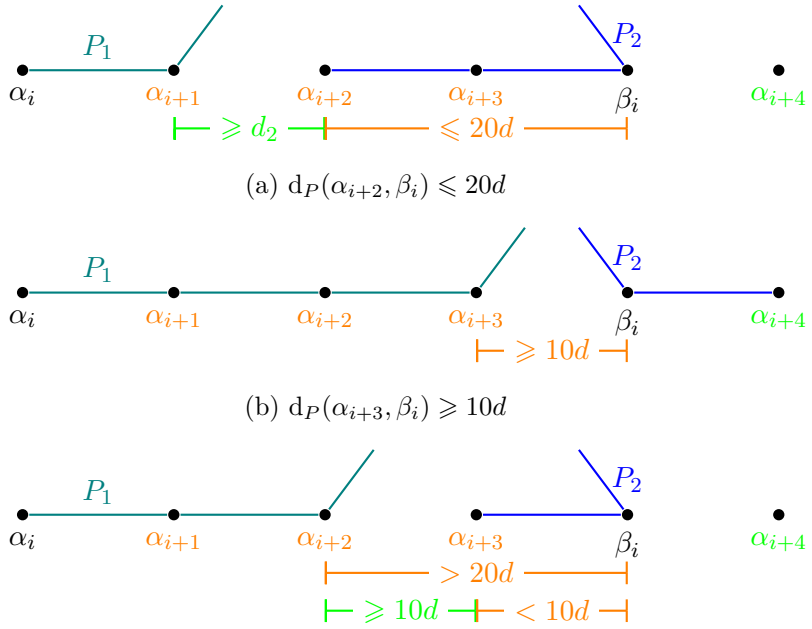

(a) $\mathrm{d}_P(\alpha_{i+2}, \beta_i) \leqslant 20d$

(b) $\mathrm{d}_P(\alpha_{i+3}, \beta_i) \geqslant 10d$

(c) $\mathrm{d}_P(\alpha_{i+2}, \beta_i) > 20d$ and $\mathrm{d}_P(\alpha_{i+3}, \beta_i) < 10d$

Figure 6.10: α_{i+1}, α_{i+2} and α_{i+3} occur before β_i along P.

If $\mathrm{d}_P(\alpha_{i+2}, \beta_i) \leqslant 20d$, then we set $k := i + 1$ (see Figure 6.10a). Since $\mathrm{d}_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant d_2 \geqslant 10d$ 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 $\mathrm{d}_P(\alpha_{i+3}, \beta_i) \geqslant 10d$, then we set $k := i + 3$ (see Figure 6.10b). Then (v) and (vi) are clear by assumption.

Otherwise, $\mathrm{d}_P(\alpha_{i+2}, \alpha_{i+3}) = \mathrm{d}_P(\alpha_{i+2}, \beta_i) - \mathrm{d}_P(\alpha_{i+3}, \beta_i) \geqslant 20d - 10d = 10d$. In this case we set $k := i + 2$ (see Figure 6.10c). Then (vi) holds by the previous calculation. Moreover, (v) holds since $\mathrm{d}_P(\alpha_{i+3}, \beta_i) < 10d$ by assumption.

Claim 6. $\mathrm{d}_G(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

Proof. By construction, we have $\mathrm{d}_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$. Since P is a shortest $X - Y$ path, the claim immediately follows. ■

Claim 7. $\mathrm{d}_G(P_1, P_2) \geqslant d$.

Proof. Suppose for a contradiction that $p_1 \in V(P_1)$, $p_2 \in V(P_2)$, and $\mathrm{d}_G(p_1, p_2) \leqslant d - 1$. If $p_1 \in V(Q_i)$ and $p_2 \in V(Q_j)$ for some $i \neq j$, then $\mathrm{d}_G(p_1, p_2) \geqslant d$ 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, $\mathrm{d}_G(p_2, H') \leqslant 4d$. Since $\mathrm{d}_G(P, H') \geqslant d_1$, we have $\mathrm{d}_G(p_1, p_2) \geqslant d_1 - 4d \geqslant d$. Suppose $p_1 \in V(P)$ and $p_2 \in V(A_i)$ for

some i. By Claim 6, $\mathrm{d}_G(p_1, \alpha_i) \geqslant 10d$. This contradicts $\mathrm{d}_G(p_1, \alpha_i) \leqslant \mathrm{d}_G(p_1, p_2) + \mathrm{d}_G(p_2, \alpha_i) \leqslant d - 1 + d_1 < 10d$. 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 $\mathrm{d}_G(p_1, \ell_i) + \mathrm{d}_G(p_2, \ell_j) \leqslant d + 1$, then

$$\mathrm{d}_G(\ell_i, \ell_j) \leqslant \mathrm{d}_G(\ell_i, p_1) + \mathrm{d}_G(p_1, p_2) + \mathrm{d}_G(p_2, \ell_j) \leqslant (d + 1) + (d - 1) = 2d,$$

which contradicts the minimality of t. Otherwise,

$$\begin{aligned}\mathrm{d}_G(\alpha_i, \alpha_j) &\leqslant \mathrm{d}_G(\alpha_i, p_1) + \mathrm{d}_G(p_1, p_2) + \mathrm{d}_G(p_2, \alpha_j) \\ &= (\mathrm{d}_G(\alpha_i, \ell_i) - \mathrm{d}_G(\ell_i, p_1)) + \mathrm{d}_G(p_1, p_2) + (\mathrm{d}_G(\alpha_j, \ell_j) - \mathrm{d}_G(\ell_j, p_2)) \\ &\leqslant 2d_1 + 2 - (d + 2) + (d - 1) < 10d,\end{aligned}$$

which contradicts $\mathrm{d}_G(\alpha_i, \alpha_j) \geqslant 10d$ 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 $\mathrm{d}_G(H', P) = d_1 + 1$ and $||A_j|| \leqslant d_1 + 1$,

$$\mathrm{d}_G(P, p_2) + \mathrm{d}_G(p_2, \ell_j) = ||A_j|| \leqslant \mathrm{d}_G(H', P) \leqslant \mathrm{d}_G(H', p_1) + \mathrm{d}_G(p_1, p_2) + \mathrm{d}_G(p_2, P).$$

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, $\mathrm{d}_{\mathcal{W}^i}(p_1, V(H')) \leqslant cd$. Altogether, we have $\mathrm{d}_G(p_2, \ell_j) \leqslant \mathrm{d}_G(H', p_1) + \mathrm{d}_G(p_1, p_2) \leqslant cd + d - 1$. Therefore,

$$\mathrm{d}_G(p_1, \ell_j) \leqslant \mathrm{d}_G(p_1, p_2) + \mathrm{d}_G(p_2, \ell_j) \leqslant cd + 2d - 2 \leqslant (c + 2)d.$$

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. □

PART III

Tree-like structures and tangles

CHAPTER 7

Linked tree-decompositions into finite parts

This chapter is based on [AJKP24a], which is joint work with Sandra Albrechtsen, Raphael W. Jacobs and Max Pitz.

7.1 Introduction

7.1.1 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 [RS90a] on well-quasi-ordering finite graphs, and in Thomas's result [Tho89] 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 7.1.1 (Thomas 1990 [Tho90], Kříž and Thomas 1991 [KT91])**.** *Every (finite or infinite) graph of tree-width $< k$ has a linked rooted tree-decomposition*[1] *of width $< k$.*

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 [Tho89],

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

[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 [Die25; RST95] done.

- *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 7.1.1 for every integer $k \in \mathbb{N}$. Our main result is that Theorem 7.1.1 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 28. *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 7.1.1, the linked rooted tree-decomposition in Theorem 28 is additionally tight and componental.

Note that achieving just a subset of the properties in Theorem 28 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 7.2.2 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 tree-decomposition clearly fails to be tight. Following an observation by Diestel [Die94], 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 28 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 [AJKP24b] we provide examples showing that Theorem 28 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 [AJKP24b, Example 2].

7.1.2 Displaying end structure and excluding infinite minors

As our first application, we show that every tree-decomposition as in Theorem 28 displays the combinatorial and topological end structure of the underlying graph as follows, resolving a question by Halin from 1977 [Hal77, §6]. (Halin's original work yields proofs of Theorems 28 and 29 for locally finite connected graphs with at most two ends.)

Theorem 29. *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 [Car19],
- *displays the ends homeomorphically* if φ is a homeomorphism [KKP23],
- *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 [KKP23, Lemma 3.1], every tight such tree-decomposition displays all dominating vertices (Lemma 7.3.2), and every such linked tight tree-decomposition displays all combined degrees (Lemma 7.3.3).

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 [Die94] 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 29 to provide short, unified proofs for two other related results by Robertson, Seymour and Thomas.

Corollary 7.1.2 (Robertson, Seymour & Thomas 1995 [RST95], (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 [Hal78; Pit20], 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 29.

By Halin's grid theorem [Hal65; KMP22], every end of infinite degree contains a half-grid minor. And by a routine exercise, every end dominated by infinitely many vertices contains a

[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 [RST95; ST93] for details about the 'easy' implications.

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 29. 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 7.1.3 (Seymour & Thomas 1993 [ST93], (1.5))**.** *A graph contains no subdivision of 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 the binary tree T_2.*

Proof. Again, we only prove the hard implication. 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 29. 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.

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] [Jun69, §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 28 shows that one can always require the witnessing tree-decompositions for Corollaries 7.1.2 and 7.1.3 to be linked, tight, and componental.

7.1.3 Lean tree-decompositions

We already mentioned that there exist even stronger versions of the Kříž-Thomas Theorem 7.1.1. See the articles by Bellenbaum and Diestel [BD02] and by Erde [Erd18] for a modern treatment of the finite case.

Theorem 7.1.4 (Thomas 1990 [Tho90], Kříž and Thomas 1991 [KT91])**.** *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'.

[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.

Now it is natural to ask whether Theorem 7.1.4 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 [AJKP24b, Example 1] 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.

Yet, we can extend Theorem 7.1.4 in the optimal way to graphs without half-grid minor by post-processing the tree-decomposition from Theorem 28 using the finite case of Theorem 7.1.4.

Theorem 30. *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* $\mathrm{tw}(G)$.

The first part of Theorem 30 is a new result, while the 'moreover'-part reobtains the infinite case in Theorem 7.1.4. We recall that Robertson, Seymour and Thomas's strongest version [RST95, (12.11)] of Corollary 7.1.2 yields witnessing tree-decompositions which are additionally linked (even in its unrooted version). An immediate application of our detailed version of Theorem 30 (see Theorem 30' in Section 7.8) strengthens their result by showing that the witnessing tree-decomposition in Corollary 7.1.2 may even be chosen to be lean.

7.1.4 Displaying all infinities

A crucial tool for the proof of Theorem 28 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 [Pol96] 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 28 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 $G\mathring{\uparrow}e$.

- *displays the infinities* if it displays the ends homeomorphically, their combined degree, their dominating vertices, the critical vertex sets and their tight components cofinitely.

A tree-decomposition that cofinitely displays the tight components of every critical vertex set can no longer in general 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 31. *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 [Car19, Theorem 5.12] showed that for every graph there is a nested set of separations which efficiently distinguishes all its ends. Elm and Kurkofka [EK22, 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 [JK24, Theorem 1]. In [AJKP24b, Construction 3.1] we present a planar graph witnessing that this result cannot be extended to graphs without $K^{\aleph_0}$ minor [AJKP24b, 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 32. *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 32 as an application of Theorem 30: Our proof of Theorem 30 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 32.

An equivalent way of stating Corollary 32 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 [CHM22; EKT22; EK22; JK24] extending the 'tree-of-tangles' theorem from Robertson and Seymour [RS91, (10.3)] to infinite graphs.

Additionally, Albrechtsen [Alb24c, Theorem 3] applied Theorem 31 to obtain a 'tangle-tree duality' theorem for infinite graphs[8], which extends Robertson and Seymour's [RS91, (4.3) & (5.1)] other fundamental theorem about tangles to infinite graphs.

7.1.5 An open problem

Thomas famously conjectured that the class of countable graphs is well-quasi-ordered under the minor relation [Tho89, (10.3)]. In light of the observation that all known counterexamples to well-quasi-orderings of infinite graphs [DL01; Kom95; Pit23; Tho88] do not have finite

[7] See Section 7.9 for definitions.

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

tree-width by the characterisation of finite tree-width graphs (i.e. graphs with normal spanning trees, Theorem 7.2.2 below) in [Pit21], it might be interesting to also consider the following, even stronger conjecture:

Conjecture 7.1.5. *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.

7.1.6 How this chapter is organised

In Section 7.2 we give a short introduction into tree-decompositions, ends and critical vertex sets. We show in Section 7.3 that the tree-decomposition from Theorem 28 already satisfies Theorem 29. Section 7.4 consists of three parts. First, we give in Section 7.4.1 a brief overview over the proof of Theorem 28, which includes the statement of Theorems 33 and 34, our two main ingredients to the proof of Theorem 28. In Section 7.4.2, we state detailed versions of Theorems 28 and 31 which assert further properties of the respective tree-decompositions. In Section 7.4.3, we then reduce Theorems 28 and 31 to Theorems 33 and 34. We prove Theorem 33 in Section 7.5. In Section 7.6, we show some basic behaviour of paths and rays in torsos to prove Theorem 34 in Section 7.7. Finally, in Section 7.8, we prove Theorem 30 and Corollary 32. Furthermore, we compare in Section 7.10 our Theorem 33 to Elm and Kurkofka's [EK22, Theorem 2].

7.2 Preliminaries

In this section, we gather all concepts needed for the remainder of the chapter. We repeat here concepts already defined in the introduction, in the hope that it may be convenient for future reference to have all definitions gathered in one place.

Throughout this chapter, graphs may be infinite. The notation and definitions follow [Die25] unless otherwise specified; in particular, $\mathbb{N} = \{0, 1, 2, \dots\}$, and we may speak of a vertex $v \in G$ (rather than $v \in V(G)$), an edge $e \in G$, and so on.

In this chapter, a tree T often contains a special vertex $\mathrm{root}(T)$, its *root*. A rooted tree T has a natural partial order $\leqslant_T$ on its vertices and edges, which depends on $r = \mathrm{root}(T)$: for two $x, y \in V(T) \cup E(T)$, we write xTy for the unique $\subseteq$-minimal path in T which contains x and y, and we then set $t \leqslant_T t'$ if $t \in rTt'$. 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 $\lceil x \rceil \setminus x$ and $\lfloor x \rfloor \setminus x$.

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

An *(oriented) separation* of a graph G is a tuple (A, B) of vertex sets A, B of G, its *sides*, such that $A \cup B = V(G)$ and there are no edges in G joining $A \setminus B$ and $B \setminus A$. Its *separator*

is $A \cap B$, and its *order* is the size of its separator. Note that if (A, B) is a separation of G, then so is (B, A).

A *star* of separations is a set σ of separations of G such that $A \subseteq D$ and $B \subseteq C$ for every two distinct $(A, B), (C, D) \in \sigma$. The *interior* of a star σ is $\operatorname{int}(\sigma) := \bigcap_{(A,B)\in\sigma} B$.

For a star σ of separations of G, we let $\operatorname{torso}(\sigma)$ be the graph arising from $G[\operatorname{int}(\sigma)]$ by making each separator of a separation in σ complete. 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.

7.2.1 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 vertex sets of 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 $v \in G$, the subgraph of T induced by $\{t \in T \mid v \in V_t\}$ is connected.

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*. Whenever a tree-decomposition is introduced as $(T, \mathcal{V})$ in this chapter, we tacitly assume that $\mathcal{V} = (V_t)_{t\in T}$. 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 $t \in T'$, 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 $\operatorname{root}(T')$ is the node of T' containing $\operatorname{root}(T)$.

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 $(\bigcup_{t\in T_0} V_t, \bigcup_{t\in T_1} V_t)$ is a separation of G [Die25, Lemma 12.3.1]. We say that $(\bigcup_{t\in T_0} V_t, \bigcup_{t\in T_1} V_t)$ is *induced* by $\vec{e}$, or more generally by $(T, \mathcal{V})$. Its separator is $V_e := V_{t_0} \cap V_{t_1}$, the *adhesion set* of $(T, \mathcal{V})$ corresponding to e, where e is the undirected edge of T underlying $\vec{e}$.

For every node $t \in T$, we write σ_t for the set of separations of G induced by the oriented edges $\vec{e} = (s, t)$ for $s \in N_T(t)$. It is easy to see that such σ_t are stars of separations and that their interior is precisely V_t. We refer to $\operatorname{torso}(\sigma_t)$ 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.

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 7.2.1. *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 $v \notin G\mathring{\uparrow}e$. □

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

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* [Jun69, 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 $C_G(X, \varepsilon)$ (see the next subsection for a background on ends).

Theorem 7.2.2. *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 $\mathrm{root}(T_C)$. 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 $R^\varepsilon = v_0, e_0, v_1, e_1, v_2, \dots$ 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. □

7.2.2 Ends and tree-decompositions

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 $X \subseteq V(G)$ both rays have tails in the same component of $G - X$. The set of all ends of a graph G is denoted by $\Omega(G)$.

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, ray in ε; we denote this component by $C_G(X, \varepsilon)$ and say that ε *lives* in $C_G(X, \varepsilon)$. We denote by $\Omega_G(X, \varepsilon)$ all end of G which live in $C_G(C, \varepsilon)$. Now 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* an 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 the end ε. 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

$$\deg(\varepsilon) := \sup\{|\mathscr{R}| \mid \mathscr{R} \text{ is a family of disjoint rays in } \varepsilon\}.$$

Halin [Hal65, Satz 1] showed that this supremum is always attained. Together with the number $\mathrm{dom}(\varepsilon) := |\mathrm{Dom}(\varepsilon)|$ of vertices dominating ε, this sums up to the *combined degree* $\Delta_G(\varepsilon) := \deg(\varepsilon) + \mathrm{dom}(\varepsilon)$ of ε. Every end ε of a graph with normal spanning tree has countable combined degree: By [Die25, Lemma 8.2.3] $\deg(\varepsilon)$ is countable, and by [Die25, Lemma 1.5.5], every vertex in $\mathrm{Dom}(\varepsilon)$ lies on the unique normal ray in ε, so $\mathrm{dom}(\varepsilon)$ is countable, too.

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 $\mathrm{root}(T)$; 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 $R = v_0e_0v_1e_1\dots$ 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 $\liminf_{e\in R} V_e = \mathrm{Dom}(\varepsilon)$ 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 $G{\restriction}e$ and to which V_e is linked. The following lemma is clear.

Lemma 7.2.3. *Every end-linked, rooted tree-decomposition $(T, \mathcal{V})$ is tight.* □

7.2.3 Critical vertex sets and tree-decompositions

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

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 7.2.4. *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.* □

The following theorem was first proved by Polat [Pol96, Theorems 3.3 & 3.8]; we here present a short proof using normal spanning trees.

Theorem 7.2.5. *Every tough, rayless graph is finite.*

Proof. Let G be a tough, rayless graph. Since G is rayless, $V(G)$ is trivially dispersed, so every component C of G admits a normal spanning tree T_C by Jung's normal spanning tree criterion. By normality, for every node $t \in T_C$ all its successors s are contained in distinct components $\lfloor s \rfloor$ of $C - \lceil t \rceil$. Hence, since G is tough and $\lceil t \rceil$ is finite, T_C has no vertices of infinite degree. As every locally finite, rayless tree, such as the T_C, is finite [Die25, Proposition 8.2.1], the components C of G are finite as well. Since G is tough, it has only finitely many components. Hence, G is finite. □

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. 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.

7.2.4 Critical vertex sets of torsos

The following lemma, which we will use in the proofs of Theorems 30 and 31, 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 7.2.6. *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 $C \in \breve{\mathscr{C}}(G - X)$ 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 $\text{torso}(\sigma) - X$ meet U. For every component C' of $\text{torso}(\sigma) - X$ 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 $\text{torso}(\sigma)$, 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 $\text{torso}(\sigma) - X$ than C'. It thus suffices to show for all infinitely many components C' of $\text{torso}(\sigma) - X$ 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 $X \subseteq \text{int}(\sigma)$. It suffices to show that $C' := C \cap \text{int}(\sigma)$ is a tight component of $\text{torso}(\sigma) - X$ for every tight component C of $G - X$ which meets $\text{int}(\sigma)$ but avoids the finite set U. The definition of torso immediately yields that C' induces a connected subgraph of $\text{torso}(\sigma) - X$ with $N_{\text{torso}(\sigma)}(C') \supseteq X$ because $X \subseteq \text{int}(\sigma)$. It remains to show that $N_{\text{torso}(\sigma)}(C') \subseteq X$. 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. □

7.3 Displaying ends

In this short section we show that Theorem 29 follows from Theorem 28, 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 [KKP23, Lemma 3.1] applied to the tree-decomposition induced by contracting all edges which violate componental.

Lemma 7.3.1. *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 7.3.2. *Let $(T, \mathcal{V})$ be a tight, componental, rooted tree-decomposition of a graph G which has finite adhesion. Then $\liminf_{e \in R} V_e = \text{Dom}(\varepsilon)$ for every rooted R of T and the unique end ε of G which gives rise to R.*

[10] Our definition of *componental* agrees with the definition of *upwards connected* from [KKP23].

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 7.3.3. *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 $R = v_0e_0v_1e_1\dots$ 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 [GH22, 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 every $n \in \mathbb{N}$. Then $\mathscr{R} = \bigcup_{n\in\mathbb{N}} \mathcal{P}_n$ is a family of pairwise disjoint paths and rays; the rays in $\mathscr{R}$ are necessarily ε-rays since ε is the unique end of G that gives rise to R, and the finite paths in $\mathscr{R}$ necessarily have a vertex of $\mathrm{Dom}(\varepsilon)$ as their last vertex. Hence, $|\mathscr{R}| \leqslant \Delta(\varepsilon)$. Moreover, we clearly have $|\mathscr{R}| = \liminf_{n\in\mathbb{N}} |S_n|$, which implies that $\liminf_{n\in\mathbb{N}} |S_n| = |\mathscr{R}| \leqslant \Delta_G(\varepsilon)$, as desired. □

Proof of Theorem 29 given Theorem 28. By Lemmas 7.3.1 to 7.3.3 the tree-decomposition from Theorem 28 is as desired. □

7.4 A high-level proof of the main result

7.4.1 A two-step approach to Theorem 28

By Theorem 7.2.5, 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 7.2.5.

But how to guarantee that the resulting tree-decomposition is linked? Lemma 7.2.4 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 33. *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:

Theorem 34. *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 33 and Theorem 34 can be combined to yield a proof of Theorem 28. The subsequent Sections 7.5 and 7.7 are then concerned with proving Theorems 33 and 34 respectively.

7.4.2 Detailed versions of the main theorems

In fact, we will prove a version of Theorem 28 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 28' (Detailed version of Theorem 28)**.** *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*

[11] This is not completely true: In order to ensure that the arising tree-decomposition is linked we need a 'refinement version' of Theorem 34 (see Theorem 34') 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 34.

(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 $G{\uparrow}e$, and*

(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 7.5.1). 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 [Hal75, 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; [AJKP24b, Example 5.2] 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.

Now as already indicated in the introduction, we first show Theorem 31 and then derive Theorem 28 from it. In fact, we derive the detailed version Theorem 28' from the following detailed version of Theorem 31.

Theorem 31' (Detailed version of Theorem 31)**.** *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 28'. Moreover,*

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

(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 $V_t \in \operatorname{crit}(G)$.*

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 31' already 'nearly' satisfies (L3), that is, has 'almost' distinct bags.

Lemma 7.4.1. *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 7.2.1 implies that $V_{s'} = V_t = V_s$, a contradiction. □

Proof of Theorem 28' given Theorem 31'. Let $(T', \mathcal{V}')$ be the tree-decomposition obtained from Theorem 31'. 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 28', 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 7.4.1, we only need to check that every critical vertex set appears as at most one bag of $(T, \mathcal{V})$. By Lemma 7.4.1, every $X \in \mathrm{crit}(G)$ 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 31', we still follow the promised two-step approach, but need the following detailed versions of Theorems 33 and 34.

Theorem 33' (Detailed version of Theorem 33)**.** *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 31'.*

In order to state the detailed version of Theorem 34, we need one more definition. 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 34' (Detailed version of Theorem 34)**.** *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.

7.4.3 Proof of the main result of this chapter

According to our two-step approach, we prove Theorem 31' by applying Theorem 34' to the torsos of the tree-decomposition given by Theorem 33'. For this, we need to ensure that all torsos again have finite tree-width:

Lemma 7.4.2. *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 31' given Theorem 33' and Theorem 34'. Let $(T^1, \mathcal{V}^1)$ be the rooted tree-decomposition from Theorem 33' 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 31'. Let t be a node of T^1. We describe how we refine the torso at t in T^1 using Theorem 34':

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 $V^1_t \in \mathrm{crit}(G)$; 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 $V^1_t \in \mathrm{crit}(G)$. Thus, $(T^t, \mathcal{V}^t)$ is a rooted tree-decomposition as in the conclusion of Theorem 34'.

Secondly, we now assume that t is a node of T^1 with $V^1_t \notin \mathrm{crit}(G)$. By construction, the torso of σ'_t in G^t is equal to the torso of σ_t in G; so by Lemma 7.4.2, this torso has finite tree-width. We claim that G^t, X_t and σ'_t are as required to apply Theorem 34'. 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 $V_t \notin \mathrm{crit}(G)$, $t \neq t_X$; 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 7.2.1 yields $V_t = X \in \mathrm{crit}(G)$ which contradicts the assumptions on t. All in all, we may now apply Theorem 34' 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 $\mathrm{root}(T^{\mathrm{root}(T^1)})$. 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 31', since $(T^1, \mathcal{V}^1)$ and all the $(T^t, \mathcal{V}^t)$ satisfy (I1). It also satisfies (I2). Indeed, by Theorem 33', $(T^1, \mathcal{V}^1)$ satisfies (I2). Moreover, by (R4) from Theorem 34', 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 $V^1_e \subsetneq V^s_{\mathrm{root}(T^s)}$ 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 34' satisfies $X_t = V^1_{st} \subsetneq V^1_t = \mathrm{int}(\sigma) = \mathrm{int}(\sigma'_t)$ by (I2) from Theorem 33' and also $G^t - X_t$ is connected with $N^t_G(G^t - X) = X$ by (I1) and since $(T^1, \mathcal{V}^1)$ is fully tight. It follows that $V^2_x \supsetneq V^2_e \subsetneq V^2_y$ for all edges $e = xy$ with $x <_{T^2} y$ except those where $V^2_x = V^1_x = V^1_y = V^2_y \in \mathrm{crit}(G)$. 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 7.2.5, 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^2_s at $s \in T^2$ in $(T^2, \mathcal{V}^2)$ is equal to the torso G^t_s at s in $(T^t, \mathcal{V}^t)$; in particular, these torsos are rayless by (R1). Suppose for a contradiction that the torso G^2_s at $s \in T^2$ in $(T^2, \mathcal{V}^2)$ contains a critical vertex set X. By Lemma 7.2.6 (i), infinitely many tight components of $G - X$ meet the torso G^2_s; in particular, they meet the torso G^1_t at $t \in T^1$ in $(T^1, \mathcal{V}^1)$. Now by Lemma 7.2.6 (ii), cofinitely many of these tight components of $G - X$ restrict to tight components of the torso G^1_t at t in $(T^1, \mathcal{V}^1)$. Thus, X is a critical vertex set of the tough torso G^1_t 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 34' is linked and X_t-linked, where $X_t = V^1_{st}$ 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 $k := |V^t_f|$ disjoint V^t_e–$V^t_{e'}$ 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^2_e–$V^2_{e'}$ 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 $V^2_{f_i}$–$V^2_{f_{i+1}}$ 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 $V^2_{g_1}$ and $V^2_{g_3}$. Therefore, $P_i \in \mathcal{P}_i$ and $P_j \in \mathcal{P}_j$ are internally disjoint for $i \neq j$. We remark that $V^2_{f_{i+1}}$ is an adhesion set of $(T^1, \mathcal{V}^1)$ and thus critical in G. Hence, Lemma 7.2.4 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 7.3.1 to 7.3.3. It is immediate from the construction that $(T^2, \mathcal{V}^2)$ still displays the critical vertices and their tight components cofinitely. □

7.5 Tree-decomposition along critical vertex sets

In this section we prove Theorem 33', which we restate here for convenience.

Theorem 33' (Detailed version of Theorem 33)**.** *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 $V_s \supseteq V_t \in \operatorname{crit}(G)$ and $\deg(t) = \infty$, and*

(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 $V_t \in \operatorname{crit}(G)$.*

Recall that these (I1) and (I2) are the same properties as (I1) and (I2) in Theorem 31'.

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 33' 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 7.5.1. *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 $X \in \operatorname{crit}(G)$ 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 $t_X <_T t$ 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 7.2.6 (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 7.5.1:

Lemma 7.5.2. *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 7.2.6 (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 33' except that it might not display the critical vertex sets.

Lemma 7.5.3. *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 7.5.3 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 7.2.2. 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 7.5.1 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 $X \in \operatorname{crit}(G)$, 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 7.5.2, 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 7.5.3 into one that additionally displays all critical vertex sets and their tight components cofinitely.

Construction 7.5.4. Let $(T, \mathcal{V})$ be a rooted tree-decomposition of a graph G of finite adhesion. For every $X \in \operatorname{crit}(G)$, let t_X be the $\leqslant_T$-minimal node t with $X \subseteq V_t$ (which exists by Lemma 7.5.1). We obtain the tree T' from T by simultaneously adding for each $X \in \operatorname{crit}(G)$ 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 $X \in \operatorname{crit}(G)$.

To show that the tree-decomposition $(T', \mathcal{V}')$ from Construction 7.5.4 displays the critical vertex sets of G we need the following auxiliary lemma.

Lemma 7.5.5. *Let $(T, \mathcal{V})$ be a tight, rooted tree-decomposition of a graph G such that for every $X \in \operatorname{crit}(G)$ 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 $V_t \in \operatorname{crit}(G)$. So let $t \in T$ be a node with finite V_t and $V_t \notin \operatorname{crit}(G)$. 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 7.5.4 is as desired for Theorem 33' if we start with a tree-decomposition $(T, \mathcal{V})$ that is tight and componental.

Lemma 7.5.6. *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 7.5.4 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 7.5.4. Then (i), (ii) and (iii) hold for $(T', \mathcal{V}')$. Lemma 7.5.1 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 7.5.5, $(T', \mathcal{V}')$ 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 33'. Apply Construction 7.5.4 to the tree-decomposition $(T, \mathcal{V})$ obtained from Lemma 7.5.3. By Lemma 7.5.6 this tree-decomposition $(T', \mathcal{V}')$ is as desired; in particular, it satisfies (I1). Moreover, since $(T, \mathcal{V})$ satisfies (i) from Lemma 7.5.3, $(T, \mathcal{V})$ still satisfies it at all $t \in T'$ which have not been not added in Construction 7.5.4, i.e. $(T', \mathcal{V}')$ satisfies (I2). □

We remark that a Theorem similar to Theorem 33' was proven by Elm and Kurkofka [EK22, Theorem 2]. However, their result cannot be used in place our Theorem 33', as their method in general produces only a nested set of separations, and not a tree-decomposition. We refer the reader to the Section 7.10 for a detailed discussion of the differences between the two results.

7.6 Lifting paths and rays from torso

From Theorem 33' we obtain a tree-decomposition $(T, \mathcal{V})$ whose torsos are tough. Following our overall proof strategy for Theorem 28 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 7.4.1 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 7.6.3). Moreover, we prove similar results, Proposition 7.6.1 and Lemma 7.6.2, for extending rays and infinite path families in torsos at stars whose separations are just left-tight.

We recall that for two sets X, Y of vertices of a graph G, an *X–Y path* meets X precisely in its first vertex and Y precisely in its last vertex. A path *through* a set C of vertices in G is a path in G whose internal vertices are contained in C.

Proposition 7.6.1. *Let σ be a star of left-tight, finite-order separations of a graph G. Let P be a path or ray in* $\mathrm{torso}(\sigma)$*. Then there exists a path or ray P' in G, respectively, with $P' \cap G[\mathrm{int}(\sigma)] \subseteq P$ 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. [Die25, 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 $H \cap G[\mathrm{int}(\sigma)] = P \cap G[\mathrm{int}(\sigma)]$ by construction, and because $P' \subseteq H$, it follows that $P' \cap G[\mathrm{int}(\sigma)] \subseteq P$, so P' is as desired. □

We remark that later in this chapter we sometimes want to apply Proposition 7.6.1 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 7.6.1 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 $P' \cap G[\mathrm{int}(\sigma)] \subseteq P$ which meets $V(P)$ infinitely often.

Lemma 7.6.2. *Let σ be a star of left-tight, finite-order separations of a graph G. Let $X, Y \subseteq \mathrm{int}(\sigma)$ such that there are infinitely many disjoint X–Y paths in* $\mathrm{torso}(\sigma)$*. Then there are infinitely many X–Y paths in G.*

Proof. Let $\mathcal{P}$ be an infinite family of disjoint X–Y paths in $\mathrm{torso}(\sigma)$. 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 7.6.1, 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 7.6.3. *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. □

7.7 Linked tree-decompositions into rayless parts

In this section we prove Theorem 34', which we restate here for convenience.

Theorem 34' (Detailed version of Theorem 34)**.** *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* $\mathrm{torso}(\sigma)$ *has finite tree-width. Further, let $X \subseteq \mathrm{int}(\sigma)$ 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 $G{\restriction}e$ or linked to a set $A \cap B \subseteq G{\restriction}e$ 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 $G{\restriction}e'$ or is contained in $A \cap B \subseteq V(G{\restriction}e')$ 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 34' 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 34' to the following statement which differs slightly from Theorem 34'. 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 $C(A \cap B, \varepsilon) \subseteq A$.

Theorem 7.7.1. *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 $X \subseteq \operatorname{int}(\sigma)$ 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 $A \cap B \subseteq \operatorname{Dom}(G[A \setminus B])$.*

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 34'.*

Proof of Theorem 34' given Theorem 7.7.1. Let H be the graph obtained from $\operatorname{torso}(\sigma)$ 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 7.7.1 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 34', $\operatorname{torso}(\sigma)$ 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 7.7.1, 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 7.6.3 (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 $V_e = V'_e \subseteq \mathrm{int}(\sigma)$ is linked to $A \cap B \subseteq V(G{\uparrow}e)$ since $A \cap B$ separates $\mathrm{int}(\sigma)$ from $R_{A,B}$ and because of Lemma 7.6.3 (ii) as the $H[A \cap B]$ are complete. Otherwise, Lemma 7.6.3 (i) and Lemma 7.6.2 ensure that V_e is linked to an end of G that lives in $G{\uparrow}e$. Indeed, let $\mathscr{R}$ be a family of $|V_e|$ equivalent, disjoint rays in $\mathrm{torso}(\sigma)$ that start in V_e and that witness that V_e is end-linked. Then Lemma 7.6.3 (i) yields a family $\mathscr{R}'$ of $|V_e|$ disjoint rays in $G{\uparrow}e$ that start in V_e and that each meet $V(\mathscr{R})$ infinitely often. In particular, since $\mathscr{R}$ is finite, for every $R' \in \mathscr{R}'$ there is $R \in \mathscr{R}$ such that R' meets $V(R)$ infinitely often. To see that the rays in $\mathscr{R}'$ are equivalent, let $R'_0, R'_1 \in \mathscr{R}'$ be given. Since R_0 and R_1 are equivalent in $\mathrm{torso}(\sigma)$, 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 $\mathrm{torso}(\sigma)$. Since these paths are in fact R'_0–R'_1 paths in $\mathrm{torso}(\sigma)$, Lemma 7.6.2 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 7.7.1. 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^n_e is contained in Y. We then replace the bag V^n_ℓ 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 7.7.1 the Algorithm 7.7.2 that will construct the bags of the desired tree-decomposition for Theorem 7.7.1, that is, the sets Y mentioned above. This algorithm is simple to state, but we require some tools to

properly analyse it. In Section 7.7.2 we build a set of tools centred around 'regions', which we then use in Section 7.7.3 to prove some properties of Algorithm 7.7.2 that will later on ensure that the resulting tree-decomposition is as desired. Finally, in Section 7.7.4, we follow the above described approach to construct a tree-decomposition by inductively applying Algorithm 7.7.2 and then prove with the help of the main result from Section 7.7.3 that this tree-decomposition is as desired.

7.7.1 Building the bags of the tree-decomposition

In this section we describe a transfinite recursion, Algorithm 7.7.2, that will construct the bags of the tree-decomposition for Theorem 7.7.1. 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 7.7.2. (Construction of a bag)

Input: a connected graph H; a finite set X of $k \in \mathbb{N}$ vertices of H; a set $\mathscr{D}$ 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 $D \in \mathscr{D}$.

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 $D \in \mathscr{D}$; 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 $D \in \mathscr{D}$, 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

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

with the C_j for $j < i$ and with all $D \in \mathscr{D}$ 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 7.7.1 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 7.7.2 as the part $G{\restriction}\ell$ 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 7.7.2 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 7.7.9 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 7.7.1 requires the torsos of the tree-decomposition to be rayless, so Y should not contain any end of H: for this, Algorithm 7.7.2 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 will cut off all the remaining ends of G. We will justify these intuitions in Theorem 7.7.9.

We begin by argueing that Algorithm 7.7.2 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 7.7.3. *Let H be a graph. Let $\mathscr{C}$ be a chain of regions of H. Then $C = \bigcup \mathscr{C}$ is a region, and $\mathscr{C}' := \{V(H - C) \cap N_H(C') \mid C' \in \mathscr{C}\}$ is a chain with $N_H(C) = \bigcup \mathscr{C}'$.*

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

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

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

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

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

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

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

(iv): If the neighbourhood of C is finite, the fact that $\mathscr{C}'$ is a chain with $\bigcup \mathscr{C}' = N_H(C)$ yields that there is a $C' \in \mathscr{C}$ 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 $C' \in \mathscr{C}$ 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 $C' \in \mathscr{C}$, then D is also contained in $C \supseteq C'$, as desired. So we may assume that D is contained in no $C' \in \mathscr{C}$. Note that whenever $C'_0 \in \mathscr{C}$ is contained in D or does not touch D, every $C' \in \mathscr{C}$ with $C' \subseteq C'_0$ is contained in D or does not touch D, respectively. Thus, either all $C' \in \mathscr{C}$ 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. □

7.7.2 Regions

In this section we collect some statements about regions which we then use in Section 7.7.3 to analyse the regions C_i and the set Y which Algorithm 7.7.2 outputs. Recall that we want Algorithm 7.7.2 to cut off all ends of H in that every end of H lives in some C_i. In order to prove that Algorithm 7.7.2 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 7.7.6 and 7.7.7, which will later on ensure that if at step i of the recursion in Algorithm 7.7.2 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 7.7.2 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 7.7.4 (i)) or Case B (Lemma 7.7.4 (ii)) – except that it might not be nested with $\mathscr{D}$ and with the previously chosen C_j for $j < i$.

Lemma 7.7.4. *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, [GH22, 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 [Die25, Proposition 8.4.1]), the desired paths exist.

(ii): [GH22, 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 7.7.4 to obtain one which is additionally nested with $\mathscr{D}$ 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 7.7.5. *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 [Die25, 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 $D \in \mathscr{D}$ 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 7.7.6. *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, $D \in \mathscr{D}$ meets $N(C)$.

Let $\mathcal{E}'_<$ consist of all $D \in \mathscr{D}'$ 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 $(N(C) \setminus \bigcup_{D\in\mathcal{E}'_<} V(D)) \cup \bigcup_{D\in\mathcal{E}'_<} (N(D) \setminus V(\bar{C}))$, 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 $N(D) \setminus V(\bar{C})$ and $V(C) \cap N(D)$ in $\bar{D}$, must have at least size $|N(D) \setminus V(\bar{C})|$, since D is well linked by assumption and $|V(D) \cap N(C)| < |V(C) \cap N(D)|$ by definition of $\mathcal{E}'_<$. Hence, $|N(D) \setminus V(\bar{C})| \leqslant |V(D) \cap N(C)| < |V(C) \cap N(D)|$ 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 $|N(D) \setminus V(\bar{C})|$ pairwise disjoint $(N(D) \setminus V(\bar{C}))$–$(V(C) \cap N(D))$ 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 $N(C) \setminus (\bigcup_{D\in\mathcal{E}'_<} V(D))$ 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

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

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 $D \in \mathscr{D}$ an ε-linked region which is a candidate for the region C_i in Case B.

Lemma 7.7.7. *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 7.7.4 (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 7.7.6 ensures that $\mathcal{E}'$ is finite. Thus, Lemma 7.7.4 (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 $D \in \mathscr{D}$ as C_0 avoids Z. Now Lemma 7.7.6 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 $D \in \mathcal{E}$. □

7.7.3 Analysis of Algorithm 7.7.2

In this section we analyse Algorithm 7.7.2 and provide in Theorem 7.7.9 the properties of the regions C_i and the set Y obtained from Algorithm 7.7.2 which we need for the proof of Theorem 7.7.1. First let us take note of some basic properties of Algorithm 7.7.2 that follow easily from its definition:

Observation 7.7.8. *In the setting of Algorithm 7.7.2, 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 7.7.2 terminates at some ordinal $< |H|^+$.*

Proof. (O1): It is immediate from Algorithm 7.7.2 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 7.7.2 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 7.7.2.

(O5): By (O3) we have $\ell_j \leqslant \ell_i$. So suppose for a contradiction that $\ell_j = \ell_i$. By Algorithm 7.7.2, 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 7.7.2.

(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 7.7.2, we show in Section 7.7.4 that the iterative application of Algorithm 7.7.2 yields the desired tree-decomposition for Theorem 7.7.1:

Theorem 7.7.9. *In the setting of Algorithm 7.7.2 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 $D \in \mathscr{D}$ 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 $C'_1 \supseteq C'_2 \supseteq \cdots \supseteq C'_n$ 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 7.7.2 applied to $H^m := \bar{C}'_m$, the finite set $X^m := N_H(C'_m)$, the set $\mathscr{D}^m := \{D \in \mathscr{D} \mid D \subseteq C'_m\}$ and an arbitrary vertex $x_m \in C'_m - (\bigcup \mathscr{D})$. 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 $\mathscr{C}_x$ be the collection of all those C_i that contain x. Then $\mathscr{C}_x$ is non-empty since by the definition of Y at least one C_i contains x. We claim that $C := \bigcup \mathscr{C}_x$ is a component of $H - Y$ with $|N(C)| < k = |X|$. Indeed, it is immediate from Algorithm 7.7.2 that every $C_i \in \mathscr{C}_x$ is a region as in Case A, since the regions that where chosen in Case B avoid x; in particular, every $C_i \in \mathscr{C}_x$ is a $(< k)$-region. Since all the $C_i \in \mathscr{C}_x$ meet in x and are nested, $\mathscr{C}_x$ is a chain (with respect to inclusion). Thus, Lemma 7.7.3 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 $C = \bigcup \mathscr{C}_x$, we have $N(C) \subseteq \bigcup\{N(C_i) \mid C_i \in \mathscr{C}_x\}$. Since all the C_i are nested, every C_j which meets the neighbourhood of some $C_i \in \mathscr{C}_x$ already contains C_i; in particular, such C_j are also in $\mathscr{C}_x$ 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 $\mathscr{D}$ and the choice of the C_i, in every step i of Algorithm 7.7.2 every $D \in \mathscr{D}$ 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 $\mathscr{D}$, the region $D \in \mathscr{D}$ is ε-linked for some end ε of H. So, since Algorithm 7.7.2 terminates because of Case C by Observation 7.7.8 (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 7.7.2. 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 7.7.2, 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 7.7.2 and the assumptions on $\mathscr{D}$ that in step i, the region $D \in \mathscr{D}$ 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 $\mathscr{C}^{<\ell}_{\max}$ be the set of all C_i with $\ell_i < \ell$ that are $\subseteq$-maximal among all $(< \ell)$-regions C_j. By Observation 7.7.8 (O5) every C_i with $\ell_i < \ell$ is included in some $C_j \in \mathscr{C}^{<\ell}_{\max}$. Moreover, since all C_i are nested, the regions in $\mathscr{C}^{<\ell}_{\max}$ are pairwise non-touching. Hence, by Lemma 7.7.7 applied to $Z := X \cup \{x\}$ and the set $\mathcal{E}$ of all $\subseteq$-maximal elements in $\mathscr{D} \cup \mathscr{C}^{<k}_{\max}$, 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 $\mathscr{D} \cup \mathscr{C}^{<k}_{\max}$. Set $\ell := |N(C)|$. Then Lemma 7.7.6 applied to C and the set $\mathcal{E}$ of all $\subseteq$-maximal regions in $\mathscr{D} \cup \mathscr{C}^{<\ell+1}_{\max}$ yields a ε-linked $(\leqslant \ell)$-region C' which is nested with all regions in $\mathscr{D} \cup \mathscr{C}^{<\ell+1}_{\max}$ 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 $D \in \mathscr{C}^{<\ell+1}_{\max} \setminus \mathscr{C}^{<k}_{\max}$ since any such region D was chosen in Case B. Moreover, by Lemma 7.7.6, C' is not strictly contained in any region in $\mathscr{D} \cup \mathscr{C}^{<\ell+1}_{\max}$. In particular, C' is nested with all $D \in \mathscr{D}$ and all C_i with $\ell_i \leqslant \ell$. But this contradicts Algorithm 7.7.2: If there is

no C_i with $\ell_i > \ell$, then C' contradicts that Algorithm 7.7.2 terminated because of Case C by Observation 7.7.8 (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 7.7.8 (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 [Die25, 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 [Die25, 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 7.7.2 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 7.7.4 (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_i^N be the regions obtained from the application of Algorithm 7.7.2 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 $\mathscr{C}_{\max}^{\leqslant |S|}$ be the $\subseteq$-maximal elements of the C_i^N with $\ell_i^N \leqslant |S|$. Note that by Observation 7.7.8 (O5) every C_i^N with $\ell_i^N < |S|$ is contained in some $C_j^N \in \mathscr{C}_{\max}^{\leqslant |S|}$. Since all regions in $\mathscr{C}_{\max}^{\leqslant |S|} \cup \mathscr{D}$ are end-linked, they are well linked by Lemma 7.7.5. Moreover, if ε lives in $D \in \mathscr{D}$, then $D \subseteq C'_n \subseteq C$ by (A2). To be able to apply Lemma 7.7.6 to C and the set $\mathcal{E}$ of $\subseteq$-maximal regions in $\mathscr{C}_{\max}^{\leqslant |S|} \cup \mathscr{D}$ it suffices that ε lives in no C_j^N with $\ell_j^N \leqslant |S|$.

So suppose for a contradiction that ε lives in some C_j^N with $\ell_j^N \leqslant |S|$. We claim that then

$C'' := C_j^N \cup C'_n$ is an ($\leqslant |S|$)-region. Indeed, C'' is a region because C_j^N and C'_n are both connected and intersect as ε lives in both of them. Moreover, $N(C'') = (N(C_j^N) \cup N(C'_n)) \setminus V(C'') = (N(C_j^N) \setminus V(C'_n)) \dot\cup (N(C'_n) \setminus V(\bar C_j^N))$ has size at most $|N(C_j^N)| = \ell_j^N \leqslant |S|$: if not, one easily checks by doubling counting that $Z := ((N(C_j^N) \cup N(C'_n)) \cap V(C'')) \cup (N(C_j^N) \cap N(C'_n)) = (N(C'_n) \cap V(C_j^N)) \dot\cup (N(C_j^N) \cap V(C'_n)) \dot\cup (N(C_j^N) \cap N(C'_n))$ has size less than $|N(C'_n)|$. But this contradicts that C'_n is ε-linked as ε lives in a component of $C_j^N \cap C'_n$, 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_j^N and C'_{N+1} both contain C'_n by assumption and thus touch. So since C_j^N and C'_{N+1} are nested by Algorithm 7.7.2 (because C'_{N+1} is some C_i^N by (A4)), Observation 7.7.8 (O4) implies that $C_j^N \subseteq C'_{N+1}$, and therefore $C'' \subseteq C'_{N+1}$. But this contradicts our choice of C. Thus, ε lives in no region in $\mathscr{C}_{\max}^{\leqslant |S|} \cup \mathscr{D}$.

So we may apply Lemma 7.7.6 to C and the set $\mathcal{E}$ of $\subseteq$-maximal separations in $\mathscr{C}_{\max}^{\leqslant |S|} \cup \mathscr{D}$ to obtain an ε-linked ($\leqslant |S|$)-region C' which is disjoint from X and which, for every region in $\mathscr{C}_{\max}^{\leqslant |S|} \cup \mathscr{D}$, 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 $D \in \mathscr{D}$. 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. □

7.7.4 Proof of Theorem 7.7.1

Using Theorem 7.7.9 we now show that the tree-decomposition obtained from iteratively applying Algorithm 7.7.2 is as desired for Theorem 7.7.1.

Proof of Theorem 7.7.1. First, we define the desired tree-decomposition $(T, \mathcal{V})$ for the case $X = \operatorname{int}(\sigma)$. 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 7.2.2, we may fix a normal spanning tree T_{NST} of G whose root is in X. We denote by $\mathscr{D}$ the set of subgraphs $G[A \setminus B]$ of G with $(A, B) \in \sigma$. The assumptions on σ ensure that $\mathscr{D}$ satisfies the assumptions in Algorithm 7.7.2. Since G admits a normal spanning tree, all its ends have countable combined degree. Hence, G satisfies the assumptions of Theorem 7.7.9. 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 $G[V_\ell^n]$ as follows. Set $X' := X$ if $n = 0$ and ℓ is the unique node of T^0, and $X' := V_f^n$ 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 $V_\ell^n \setminus (X' \cup \bigcup \mathscr{D})$. 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 $G[V_\ell^n] - (X' \cup \bigcup \mathscr{D})$ is connected by the assumption on σ.

Now apply Algorithm 7.7.2 to $H := G[V_\ell^n]$ with the finite set X', the vertex x and $\{D \in \mathscr{D} \mid D \subseteq H\}$ to obtain $Y \subseteq V(H)$. Then let T^ℓ be the star with centre ℓ and whose set of leaves is the set $\mathscr{C}$ of all components of $H - Y$. Further, set $V_\ell^\ell := Y$ and $V_C^\ell := \bar{C}$ for every $C \in \mathscr{C}$. Then Theorem 7.7.9 (A4) yields that for every edge $e = \ell C$ of T^ℓ, the graph $H\mathring{\uparrow}e$ is a C_i in Algorithm 7.7.2 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 7.6.1. But this contradicts Theorem 7.7.9 (A3). Furthermore, it satisfies the following (which will ensure (R3') of $(T^{n+1}, \mathcal{V}^{n+1})$) by Theorem 7.7.9 (A5): if $|V_e^\ell| \geqslant |V_f^n|$ for an edge $e \in T^\ell$, then $V_e^\ell \cap V_f^n \subseteq \operatorname{Dom}(H\uparrow e)$.

By construction, the T^ℓ are pairwise disjoint and only share their centre ℓ with T^n. We set $T^{n+1} := T^n \cup \bigcup_{\ell \in L} T_\ell^{n+1}$ with $\operatorname{root}(T^{n+1}) := \operatorname{root}(T^n)(= r)$, and bags $V_t^{n+1} := V_t^n$ for every node $t \in T^n - \ell$ and $V_t^{n+1} := V_t^\ell$ 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 7.7.9 (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 $V_t := V_t^N = V_t^n$ 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 $v \in V_t^n$ 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 $D \in \mathscr{D}$. Then $(*)$ is ensured by considering Theorem 7.7.9 (A1) along $\mathrm{root}(T_{NST})T_{NST}v = v_0v_1 \dots v_m$: 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 $\mathrm{root}(T_{NST}) \in X \subseteq V_r^1$ 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 $v_i = \min_{\leqslant_S}(V(G\mathring{\uparrow}e^n) \setminus (\bigcup V(\mathscr{D})))$ for every $n \geqslant N$. Hence, Theorem 7.7.9 (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 $D \in \mathscr{D}$. By Theorem 7.7.9 (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 $\mathscr{D}$ are nested, the leaf separation at ℓ is $(V(G) \setminus V(D), V(D) \cup N(D))$.

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 $V(D) \cup N(D) = V_D^N$ 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] □

7.8 Lean tree-decompositions

In this section we prove Theorem 30, the strengthening of Theorem 28 for graphs without half-grid minor in which we replace the linkedness of the tree-decomposition with leanness. We

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

restate it here in its more detailed version:

Theorem 30' (Detailed version of Theorem 30)**.** *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 30' is structured as follows. We start our construction with the tree-decomposition $(T, \mathcal{V})$ of the graph G from Theorem 31', 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 30', that is Thomas's Theorem 7.1.4. 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 30' 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 [Alb24c, Corollary 6.3] (see Lemma 7.8.1 below). Moreover, it also yields the well-linkedness of the separations on their left side (see Lemma 7.8.2 below). This property is crucial to the proof of Theorem 30' 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 7.6.3 (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 30', 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.

Following [Alb24c], we call a finite-order separation (A, B) of a graph G *left-ℓ-robust* 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 (A, B) *ℓ-robust* if it is both left- and right-ℓ-robust.

Lemma 7.8.1 ([Alb24c], Corollary 6.3)**.** *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* $\operatorname{torso}(\sigma)$ *has tree-width at most w.*

Lemma 7.8.2. *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 for 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 [Die25, 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 30': the pre-processing step in which we contract certain edges of the tree-decomposition from Theorem 31. This ensures that the assumptions of both Lemmas 7.8.1 and 7.8.2 are met at all edges incident with nodes whose (finite) torso we later aim to refine using Thomas's Theorem 7.1.4.

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 $w = \operatorname{tw}(G) \in \mathbb{N}$ 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 7.8.1 and 7.8.2 in the respective contexts.

Setting out from Theorem 31', our pre-processing step yields the following result:

Lemma 7.8.3. *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 31'. 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 7.8.4. *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 independent U–v paths $P_1^v, \dots, P_\ell^v$ in G avoiding $\mathrm{Dom}(\varepsilon) \setminus \{v\}$ that only meet in v.*

Proof. Since $R_1, \dots, R_d$ belong to the same end 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 $\mathrm{Dom}(\varepsilon) \setminus \{v\}$ 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 $u_{j-1}R_1u_j$ 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 $P_j^i := u_jR_1q_k^iQ_k^i$ and $P_j^v := u_jR_1q_k^vQ_k^v$ for the respective k as given by the choice of u_j. □

Lemma 7.8.5. *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 $R = r_0r_1 \dots$ 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 7.8.4. 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 $P^i_\ell, \dots, P^i_\ell$ and $P^v_1, \dots, P^v_\ell$ (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 7.8.4 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 7.8.6. *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 7.2.5 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 7.5.2 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 7.6.1 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 7.8.3. Let F be the set of edges in T that we contracted. By Lemma 7.8.5, the set F satisfies the premise of Lemma 7.8.6, 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 31' 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 $V_t, V_s \in \operatorname{crit}(G)$ 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 $V_s =: X \in \operatorname{crit}(G)$. 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 30'. 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 [Hal78]. Let $(T, \mathcal{V})$ be the (rooted) tree-decomposition of G from Theorem 31', and let $(T', \mathcal{V}')$ be the rooted tree-decomposition obtained from $(T, \mathcal{V})$ by applying Lemma 7.8.3 with $m = (w+1)^2(w+2)+w+1$ if $w := \operatorname{tw}(G) \in \mathbb{N}$ for finitely bounded tree-width G or with $m = 0$ for other G. Then Lemma 7.8.3 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 7.8.1 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 [Tho90, Theorem 5] (cf. Theorem 7.1.1 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[14] 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 $(T^t, \mathcal{V}^t)$

[14] We remark that this is not explicitly stated in [Tho90, Theorem 5] but follows directly from its proof.

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 torso(σ'_s) and torso(σ'_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 7.8.3 and because the T^t are finite. Moreover, since $(T', \mathcal{V}')$ displays the infinities by Lemma 7.8.3, it follows that $(\tilde{T}, \tilde{\mathcal{V}})$ also does so, as the T^t are finite and consist of a single node, if $V'_t \in \text{crit}(G)$ 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 torso(σ'_t), either there exists an edge $e \in t_1T^tt_2$ with $|V^t_e| < k$ or there are k disjoint Z_1–Z_2 paths in torso(σ'_t). In the first case, the construction of $(\tilde{T}, \tilde{V})$ yields $t_1\tilde{T}t_2 = t_1T^tt_2$ 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. And if V'_t is not critical, then all separations in σ'_t are left-good by (i) from Lemma 7.8.3 and because $(T', \mathcal{V}')$ displays the critical vertex sets of G. Hence, they are left-well-linked by Lemma 7.8.2. This allows us to apply Lemma 7.6.3 (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 $f = s_1s_2$ 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 $\tilde{V}_{s_1} \in \text{crit}(G)$, Lemma 7.2.4 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 $\tilde{V}_{s_2} \in \text{crit}(G)$.

Suppose now that $\tilde{V}_{s_1}, \tilde{V}_{s_2} \notin \text{crit}(G)$. Since there are no k disjoint Z_1–Z_2 paths in G, Menger's theorem (see for example [Die25, 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 $\tilde{V}_{s_1} \notin \operatorname{crit}(G)$, (A, B) is left-good by the construction of $(\tilde{T}, \tilde{\mathcal{V}})$ and (i) from Lemma 7.8.3. Hence, (A, B) is left-well-linked by Lemma 7.8.2. 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. □

7.9 Tree-decompositions distinguishing infinite tangles

An *infinite tangle* of a graph G is a set τ of (oriented) finite-order separations of G such that

- τ contains precisely one orientation of every finite-order separation of G, and
- there are no three (not necessarily distinct) separations $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$ with $G[A_1] \cup G[A_2] \cup G[A_3] = G$.

A separation of G *distinguishes* two infinite tangles if they orient it differently. It distinguishes them *efficiently* if they are not distinguished by any separation of smaller order. A tree-decomposition $(T, \mathcal{V})$ of G *distinguishes* some set of infinite tangles of G *efficiently* if every two tangles in this set are distinguished efficiently by a separation that is induced by an edge of $(T, \mathcal{V})$.

By [Die17, 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 $C \in \mathscr{C}(G - X)$. If the first case holds for τ, then we say that τ is *induced* by ε.

Following [EK22], 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 $C \in \mathscr{C}_X$ and such that $(V(G - C), V(C) \cup X) \in \tau'$ for a component $C \in \mathscr{C}_X$.

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

Observation 7.9.1. *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 [AJKP24b, Construction 3.1] has finite tree-width and thus by Theorem 31 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 [AJKP24b, Lemma 3.2].

However, Elm and Kurkofka [EK22, 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 [EK22, Corollary 3.4]. In the special case where G has no half-grid minor we obtain the following strengthening of their result:

Corollary 32' (Tangle version of Corollary 32)**.** *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 30' is as desired. For this, let any pair τ_1, τ_2 of combinatorially distinguishable infinite tangle be given. Then by Observation 7.9.1, 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 $C \in \mathscr{C}(G - X_i)$, 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 32. Apply Corollary 32'. □

7.10 A closer look at Theorem Theorem 33'

A theorem which is similar to Theorem 33' was proven by Elm and Kurkofka [EK22, Theorem 2]:

Theorem 7.10.1. *Every connected graph G has a nested*[15] *set of separations whose separators are precisely the critical vertex sets of G and all whose torsos*[16] *are tough.*

[15] A set N of separations is *nested* if for every two separations $\{A, B\}, \{C, D\}$ one of $A \subseteq C$ & $B \supseteq D$, $A \subseteq D$ & $B \supseteq C$, $B \subseteq C$ & $A \supseteq D$ and $B \subseteq D$ & $A \supseteq C$ holds.

[16] We refer the reader to [EK22, Section 2.6] for a definition of *torsos* in this context. We remark that if the

In the remainder of this section we compare Theorem 33' and Theorem 7.10.1. For this, let us first note that the set of separations induced by a tree-decomposition is always nested, so for graphs of finite tree-width, our Theorem 33' immediately implies Theorem 7.10.1 from Elm and Kurkofka. However, not every nested set of separations is induced by a tree-decomposition. In fact, in [EK22, Example 5.12] Elm and Kurkofka describe a graph which does not admit a tree-decomposition whose induced nested of separation is of their form[17]. We remark that this graph does not have finite tree-width.

But also for graphs of finite tree-width it is not only not obvious that their result does not yield a tree-decomposition but also not true: We present in the following Example 7.10.2 a graph G of finite tree-width such that no tree-decomposition of G induces their constructed nested set of separations. For this let us recall the nested set that Elm and Kurkofka constructed in more detail. Their nested set [EK22, Proof of Theorem 2, i.e. Theorem 5.10 & Theorem 5.11] is of the form

$$\{\{V(G) \setminus \bigcup \mathscr{K}(X), X \cup \bigcup \mathscr{K}(X)\} \mid X \in \operatorname{crit}(G)\} \\ \cup \{\{V(G) \setminus C, X \cup C\} \mid X \in \operatorname{crit}(G), C \in \mathscr{K}(X)\} \quad (7.1)$$

for some choice of $\mathscr{K}(X) \subseteq \breve{\mathscr{C}}(G - X)$ containing all but at most one element of $\breve{\mathscr{C}}(G - X)$.

The nested set induced by the tree-decomposition $(T', \mathcal{V}')$ from Theorem 33' is also of the form as in (7.1) but for some choice of $\mathscr{K}(X) \subseteq \mathscr{C}(G - X)$ consisting of cofinitely many (instead of one) elements of $\breve{\mathscr{C}}(G - X)$. The following example shows that there are graphs of finite tree-width that have no tree-decomposition inducing a nested set as in (7.1) for some choice of $\mathscr{K}(X)$ containing all but at most one element of $\mathscr{C}(G - X)$. In particular, we cannot strengthen the definition of 'displaying the tight components of all critical vertex sets cofinitely' by dropping the 'cofinitely'.

Example 7.10.2. There is a graph G of finite tree-width such that no tree-decomposition of G induces a set of separations of the form as in (7.1) for a choice of $\mathscr{K}(X) \subseteq \breve{\mathscr{C}}(G - X)$ containing all but at most one element of $\breve{\mathscr{C}}(G - X)$.

Proof. Let $R = dr_0r_1, \dots$ be a ray, and let H be the graph obtained from R by first adding the edges dr_i for all $i \in 2\mathbb{N}$ and the edges r_ir_{i+2} and r_ir_{i+3} for all $i \in 2\mathbb{N} + 1$, and then adding vertices u_{ij} for $i \in 2\mathbb{N} \setminus \{0\}$ and $j \in \mathbb{N}$ and joining them to R with edges r_iu_{ij}. To obtain G, we further add the edges

$$\{u_{ij}r_{i-1}, u_{ij}r_k \mid i \in 2\mathbb{N} \setminus \{0\}, j \in \mathbb{N}, k \in 2\mathbb{N} \cap \{0, 1, \dots, i-2\}\}$$

(see Figure 7.1). Note that G has a normal spanning tree as depicted in blue in Figure 7.1.

We claim that G is as desired, i.e. that no tree-decomposition of G induces a set of separations of the form as in (7.1) for a choice of $\mathscr{K}(X) \subseteq \breve{\mathscr{C}}(G - X)$ containing all but at

nested set is the set of induced separations of some tree-decomposition, then these torso are precisely the torsos at nodes of the decomposition tree.

[17] Moreover, one can show that their example does not even admit a tree-decomposition whose adhesion sets are the critical vertex sets of G and whose torsos are tough.

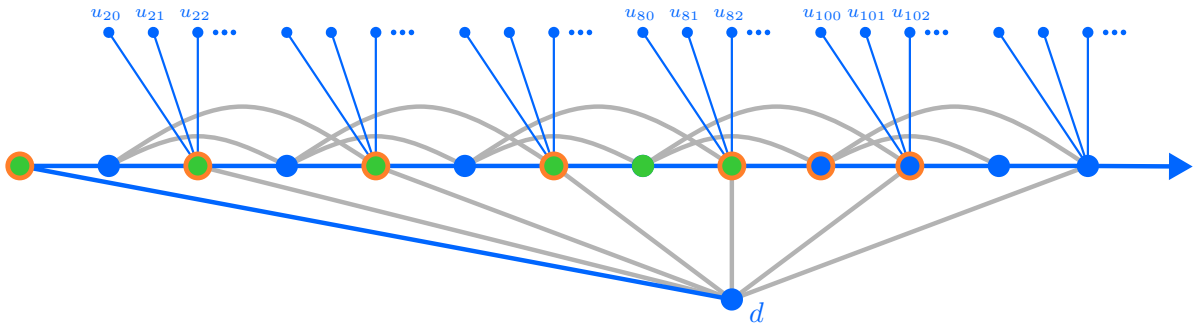

Figure 7.1: Depicted is the graph H. The green (orange) vertices are precisely the neighbours in G of u_{8j} (u_{10j}) for all $j \in \mathbb{N}$. The blue edges induce a normal spanning tree of G with root d.

most one element of $\breve{\mathscr{C}}(G - X)$. For this, let us first observe that all critical vertex sets of G are of the form

$$X_i := \{r_j \mid j \in 2\mathbb{N}, j \leqslant i\} \cup \{r_{i-1}\},$$

and the components of $G - X_i$ either consist of a single vertex u_{ij} or are of the form

$$C_i := G[\{r_j \mid j \in 2\mathbb{N}+1, j \leqslant i-2\} \cup \{u_{jk} \mid j < i, k \in \mathbb{N}\}] \text{ or}$$
$$C'_i := G[\{r_j \mid j \geqslant i\} \cup \{u_{jk} \mid j > i, k \in \mathbb{N}\} \cup \{d\}].$$

In particular, for all X_i all components of $G - X_i$ are tight (cf. Figure 7.2). Hence, the unique

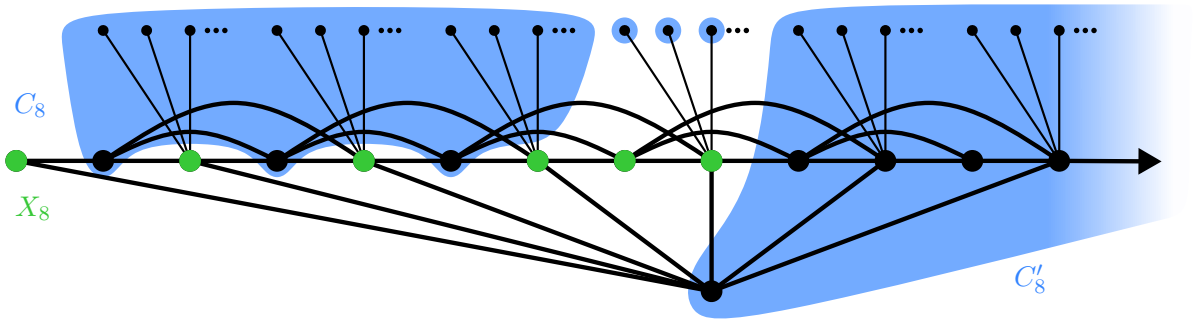

Figure 7.2: Indicated in green is the critical vertex set X_8 of G. The components of $G - X_8$ are indicated in blue.

nested set that is of the form as in (7.1) for a choice of $\mathscr{K}(X_i) \subseteq \breve{\mathscr{C}}(G - X_i)$ containing all but at most one element of $\breve{\mathscr{C}}(G - X_i)$ is the set

$$N := \{\{V(C) \cup X_i, V(G - C)\} \mid X_i \in \operatorname{crit}(G), C \in \mathscr{C}(G - X_i)\}.$$

But N is not induced by any tree-decomposition of G. Indeed, suppose for a contradiction that N is induced by a tree-decomposition $(T, \mathcal{V})$ of G, and consider the separations $\{V(G - C'_i), V(C'_i) \cup X_i\}$. Since $C'_i \supseteq C'_j$ for all $i \leqslant j$, it is easy to see that these separations have to be induced by edges e_i of T that all lie on a common ray in T. In particular, there is some $j \in 2\mathbb{N}$ such that the e_i all lie on a common rooted ray in T and that $G\mathring{\restriction}e_i = V(C'_i)$. But since

$d \in C'_i$ for all $i \in 2\mathbb{N} \setminus \{0\}$, one can conclude that d does not lie in any bag of $(T, \mathcal{V})$, which contradicts property (T1) of tree-decompositions. □

CHAPTER 8

Counterexamples regarding linked and lean tree-decompositions of infinite graphs

This chapter is based on [AJKP24b], which is joint work with Sandra Albrechtsen, Raphael W. Jacobs and Max Pitz.

8.1 Introduction

8.1.1 Lean tree-decompositions

A cornerstone in both Robertson and Seymour's work [RS90a] on well-quasi-ordering finite graphs, and in Thomas's result [Tho89] 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 8.1.1 (Thomas 1990 [Tho90], Kříž and Thomas 1991 [KT91])**.** *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 8.1.1 from finite k to arbitrary infinite cardinalities? In the following 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 35. *There is a planar, locally finite, connected graph that admits 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 35 witnesses that the answers to both questions (i) and (ii) are in the negative.

On the positive side, we provide in [AJKP24a, Theorem 3] a sufficient criterion that guarantees the existence of a lean tree-decomposition into finite parts:

Theorem 8.1.2. *Every graph without half-grid minor admits 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 35 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 8.1.2 and our negative result Example 35 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 $< \kappa$).

8.1.2 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 [BD02; Erd18]: 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 [Tho89] 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 8.1.1 by Robertson and Seymour [RS90a] and by Thomas [Tho89] 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 (cf. [AJKP24a, §1.2-1.4] and [Alb24c]; see paragraph after Theorem 8.1.3 below for more details). 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 35:

Example 36. *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 $\kappa = \aleph_0$ [AJKP24a, Theorem 1]:

Theorem 8.1.3. *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 [AB09, 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 8.1.3, are particularly useful (cf. [AJKP24a, §1.2-1.4] and [Alb24c]): In [AJKP24a, §3], we show 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 [Hal77, §6] but also allows us to deduce from Theorem 8.1.3, by means of short and unified proofs, the characterisations due to Robertson, Seymour and Thomas of graphs without half-grid minor [RST95, Theorem 2.6], and of graphs without binary tree subdivision [ST93, (1.5)]. Also the proof of Theorem 8.1.2 in [AJKP24a, §8] heavily relies on post-processing the tree-decomposition from Theorem 8.1.3. Beside these, there are more applications of rooted tree-decomposition into finite parts which are linked, tight and componental in [AJKP24a, §1.4] and [Alb24c].

In fact, we show in [AJKP24a, Theorem 1'] a more detailed version of Theorem 8.1.3 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 8.5 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 37. *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.

8.1.3 How this chapter is organised

In Section 8.2 we recall some important definitions and facts about ends and their interplay with tree-decompositions. In Section 8.3 we prove Examples 35 and 36, and in Section 8.4

we prove Example 37. Finally, in Section 8.5, we discuss whether it is possible to strengthen Theorem 8.1.3 so that the adhesion sets of the tree-decomposition are 'upwards' disjoint.

8.2 Preliminaries

We refer the reader to [AJKP24a, §2] for all the relevant definitions. Besides that, our notation and definitions follow [Die25]. In particular, the *neighbourhood* $N_G(X)$ of a set X of vertices in a graph G is the set of all vertices in $V(G) \setminus X$ that are neighbours of some $x \in X$. For convenience of the reader we recall the ones which are most important in this chapter.

Let $(T, \mathcal{V})$ be a rooted tree-decomposition of a graph G. Given an edge $e = t_1t_2$ of T with $t_1 <_T t_2$, we let T_i be the component of $T - e$ containing t_i for $i = 1, 2$. We abbreviate the sides of its induced separation by $G{\downarrow}e := G\left[\bigcup_{t\in T_1} V_t\right]$ and $G{\uparrow}e := G\left[\bigcup_{t\in T_2} V_t\right]$. Further, we write $G\mathring{\downarrow}e := G{\downarrow}e - V_e$ and $G\mathring{\uparrow}e := G{\uparrow}e - V_e$. It is easy to see that $G{\uparrow}e \supseteq G{\uparrow}f$, $G\mathring{\uparrow}e \supseteq G\mathring{\uparrow}f$, $G{\downarrow}e \subseteq G{\downarrow}f$ and $G\mathring{\downarrow}e \subseteq G\mathring{\downarrow}f$ for every two edges $e, f \in E(T)$ with $e \leqslant_T f$, and also $\bigcap_{e\in R} G\mathring{\uparrow}e = \emptyset$ for every rooted ray R in T.

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. Note that for every end ε of G and every set X of vertices of G which meets every ε-ray at most finitely often, there is a unique component of $G - X$ that contains a tail of every ε-ray. We refer to this component as $C_G(X, \varepsilon)$, which coincides with above definition for finite X.

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 remark that locally finite graphs have no dominating vertices. 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 $\Delta(\varepsilon) := \deg(\varepsilon) + |\mathrm{Dom}(\varepsilon)|$.

Given a rooted tree-decomposition $(T, \mathcal{V})$ of G, an end ε of G *gives rise* to a rooted ray R in T if every ε-ray meets every bag V_t with $t \in R$ at most finitely often, and for every $e = st \in T$ with $s < t$ we have $C_G(V_s, \varepsilon) \subseteq G\mathring{\uparrow}e$[1]. Note that if such a ray exists, then it is unique, and we denote it by R_ε. Moreover, if every ray in an end ε of G meets every bag of a tree-decomposition $(T, \mathcal{V})$ at most finitely often, then ε gives rise to a ray in T. If the bags of a rooted tree-decomposition $(T, \mathcal{V})$ meet all rays in G at most finitely often, the above yields a map $\varphi\colon \Omega(G) \to \Omega(T)$, $\varepsilon \mapsto R_\varepsilon$. Such a rooted tree-decomposition $(T, \mathcal{V})$

- *displays* the ends of G if φ is a bijection [Car19],
- *displays all dominating vertices* if $\liminf_{e\in E(R_\varepsilon)} V_e = \mathrm{Dom}(\varepsilon)$[2] for all $\varepsilon \in \Omega(G)$, and
- *displays all combined degrees* if $\liminf_{e\in E(R_\varepsilon)} |V_e| = \Delta(\varepsilon)$ for all $\varepsilon \in \Omega(G)$.

[1] We remark that this definition extends the one given in [AJKP24a, §2.2] to a larger class of tree-decompositions which do not necessarily have finite adhesion.

[2] 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 $R = v_0e_0v_1e_1v_1\ldots$ in T, one gets $\liminf_{e\in E(R_\varepsilon)} V_e = \bigcup_{n\in\mathbb{N}} \bigcap_{i\geqslant n} V_{e_i}$.

Moreover, we recall the following lemma [AJKP24a, Lemma 3.3]:

Lemma 8.2.1. *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 in T which arises from no other end of G and that $\liminf_{e \in R} V_e = \mathrm{Dom}(\varepsilon)$. Then $\liminf_{e \in R} |V_e| = \Delta(\varepsilon)$.*

The following lemma follows immediately from [AJKP24a, Lemmas 3.1, 3.2 & 3.3]:

Lemma 8.2.2. *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.* □

8.3 Examples 35 and 36 – negative answers to questions (i), (ii) and (iv)

In this section we explain Examples 35 and 36, which we restate here for convenience:

Example 35. *There is a planar, locally finite, connected graph that admits no lean tree-decomposition.*

Example 36. *There is a planar, locally finite, connected graph which admits no tight, componental rooted tree-decomposition into finite parts that is strongly linked.*

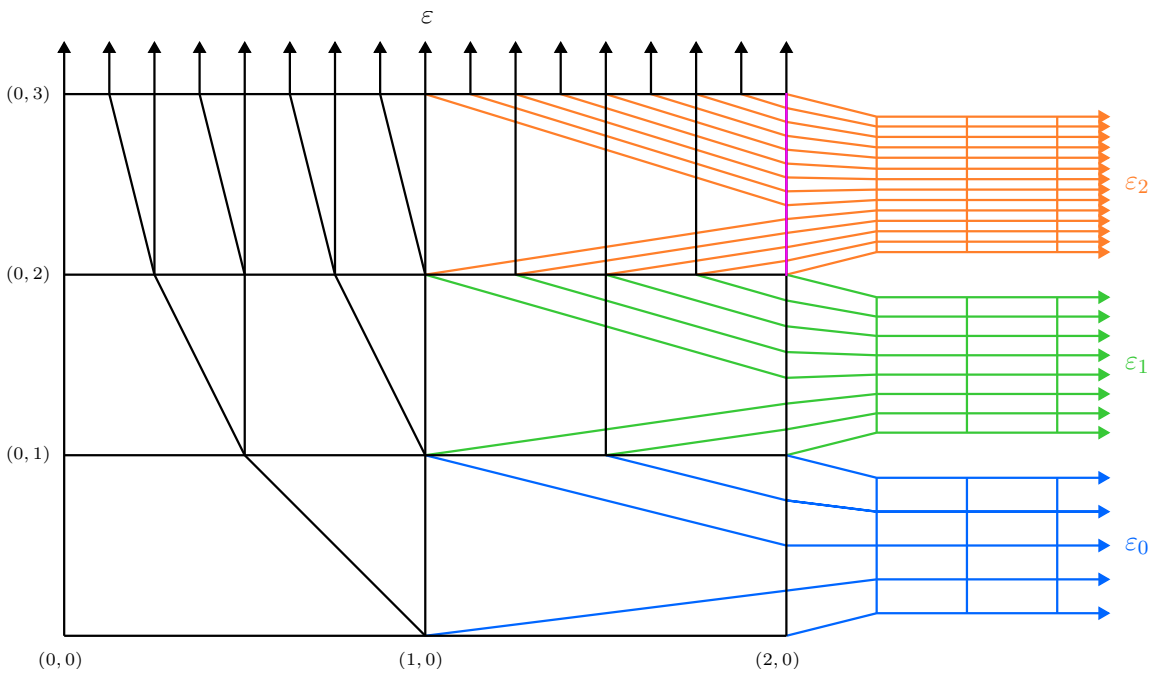

Figure 8.1: Depicted is the graph G from Construction 8.3.1. The subgraph induced by the orange edges is G_2 together with the extensions of its horizontal rays.

For our proofs of Examples 35 and 36 we construct a graph G in Construction 8.3.1 below, and then show that G is already as desired for both Examples 35 and 36. The graph in this construction is inspired by [CHM22, Example 7.4].[3]

Construction 8.3.1. Let G be the graph depicted in Figure 8.1. 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 $((2i-1)/2^{j+1}, j+1)$ for $i \leqslant 2^j$ (this is the black subgraph in Figure 8.1). 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 8.1. In particular, the extended horizontal rays of each G_n are still disjoint and their initial vertices are precisely $U_n := \{(i/2^j, j) \mid j \in \{n, n+1\}, 2^j \leqslant i \leqslant 2^{j+1}\}$. To get a graph which 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 8.1 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 8.2.) This completes the construction.

In the remainder of this section we prove that the graph G from Construction 8.3.1 is as desired for Examples 35 and 36. 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 [Die25, 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 8.3.2. *Let $(T, \mathcal{V})$ be a tree-decomposition of the graph from Construction 8.3.1. Then there exist $n \in \mathbb{N}$ such that $(T, \mathcal{V})$ does not efficiently distinguish ε_n and ε.*

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

[3]The example is only presented in the arXiv version of [CHM22].

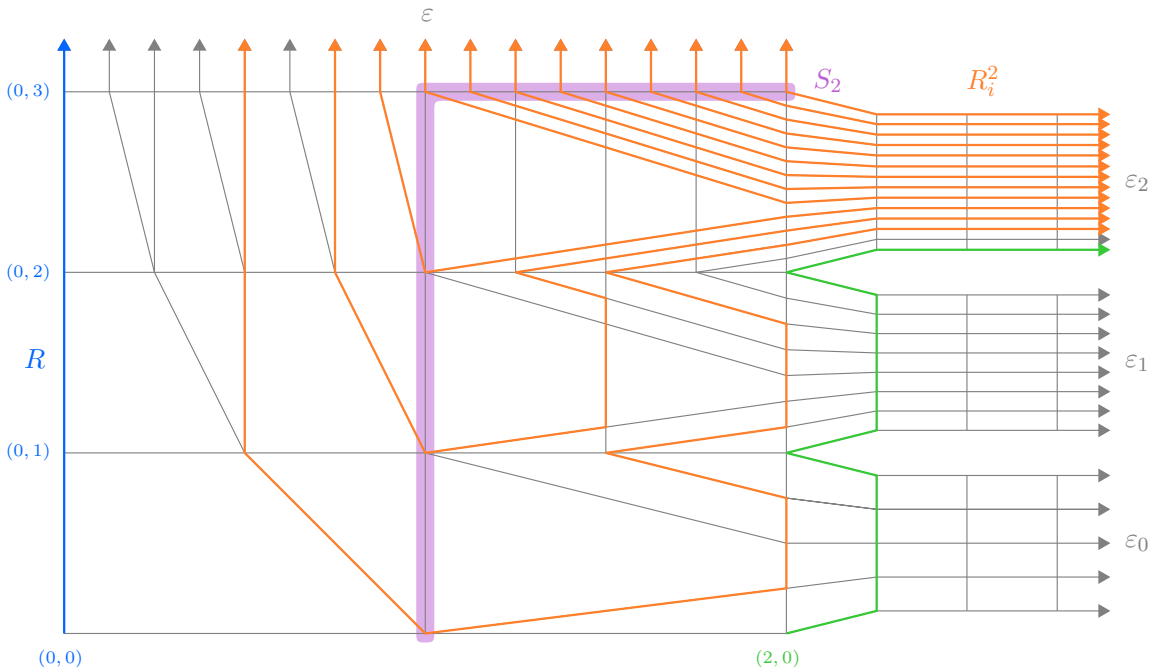

Figure 8.2: Depicted is the situation in the proof of Lemma 8.3.2 for $n = 2$.

purple in Figure 8.2). 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 8.2), 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 8.2) 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 8.2) which starts in $(2,0)$ and is disjoint from the R_i^n, 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 8.3.3. *Let $(T, \mathcal{V})$ be a rooted tree-decomposition of the graph G from Construction 8.3.1. 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 $G\mathring{\restriction}e$ avoids the first and second column of G_n. As $\liminf_{e\in R_n}|V_e| = \deg(\epsilon_n)$, 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 $G\mathring{\restriction}e_n$ avoids the first column of G_n, the component D_n of $G\mathring{\restriction}e_n$ in which ε_n lives is contained in G_n.

We claim that $G\mathring{\restriction}e_n$ 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 $G\mathring{\restriction}e_n \supseteq D_n$ contains all but finitely many vertices of G_n. Since $D_n \subseteq G\mathring{\restriction}e_n$ 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, $G{\restriction}e_n$ is connected. It follows that $G{\restriction}e_n$ is a subgraph of G_n and avoids the first column of G_n, since $G\mathring{\restriction}e_n$ 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 8.2). 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 $G{\restriction}e$ avoids H_n. For this, we note that the rays R_m are distinct from R_ϵ, since $\liminf_{e\in R_m}|V_e| = \deg(\epsilon_m)$ is finite but $\liminf_{e\in R_\epsilon}|V_e| = \deg(\epsilon)$ 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 $(G{\restriction}e) \cap (G\mathring{\restriction}e_m) = \emptyset$. As $G\mathring{\restriction}e_m$ contains all but finitely many vertices of G_m for $m \leqslant n$, the set $W_n := V(H_n) \setminus (\bigcup_{m\leqslant n} V(G\mathring{\restriction}e_n))$ is finite. Since $(T, \mathcal{V})$ is a tree-decomposition, for all but finitely many edges e of R^n_ε, the part $G\mathring{\restriction}e$ avoids W_n and hence H_n. Note that, for each edge e of R^n_ε, its adhesion set V_e avoids $G\mathring{\restriction}e_m$ 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 $G{\restriction}e$ avoids H_n. Finally, since $\liminf_{e\in R_\epsilon}|V_e| = \deg(\epsilon)$ is infinite, we may choose an edge e'_n of R_ε so that $G{\restriction}e'_n$ 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 8.2), and hence the separation induced by f_n efficiently distinguishes ε and ε_n. Since $n \in \mathbb{N}$ was chosen arbitrarily, Lemma 8.3.2 yields the desired contradiction. □

Proof of Example 36. The graph G from Construction 8.3.1 is planar, locally finite and connected. By Lemma 8.2.2, 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 8.3.3 ensures that those tree-decompositions are not strongly linked. □

We now turn to our proof of Example 35, i.e. that the graph G from Construction 8.3.1 does not admit a lean tree-decomposition. The proof consists of two steps. First, we show in Lemma 8.3.6 that G does not admit a lean tree-decomposition into finite parts. Then, we show in Lemma 8.3.7 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 8.3.3: 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 8.3.5 together with Lemma 8.2.1 implies that this holds true for *all* graphs H and lean tree-decompositions of them.

To prove Lemma 8.3.5, we first show that even though lean tree-decompositions may not be componental, they are not far away from it.

Lemma 8.3.4. *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 a trivial path. 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 8.3.5. *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 8.3.4. 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 8.3.4. □

Lemma 8.3.6. *The graph G from Construction 8.3.1 admits 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 8.3.5 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 8.2.1 implies that $\liminf_{e \in R_\omega} |V_e| = \Delta(\omega) = \deg(\omega)$ (because G is locally finite). Thus, by Lemma 8.3.3, $(T, \mathcal{V})$ is not strongly linked, which is a contradiction. □

For the proof of Example 35 it remains to show that the graph from Construction 8.3.1 neither admits a lean tree-decomposition with possibly infinite parts.

Lemma 8.3.7. *The graph G from Construction 8.3.1 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 8.3.6, $(T, \mathcal{V})$ 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 contains infinitely many vertices of $G'_n := G'[\{(i,j) \in V(G') : j > n\}] \cup \bigcup_{m>n} G_m$. 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 [Die25, Lemma 8.2.2] 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 8.2). 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^n_i in G, and hence the separation induced by f_n distinguishes ε and ε_n efficiently. Now Lemma 8.3.2 yields the desired contradiction. □

Proof of Example 35. The graph G from Construction 8.3.1 is planar, locally finite and connected. The assertion thus follows from Lemma 8.3.7. □

8.4 Example 37 – negative answer to question (vi)

In this section we construct Example 37, which we restate here for convenience:

Example 37. *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 [Car19, 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 8.3). 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 $f = t_1 t_2$ 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, $m \leqslant i$.

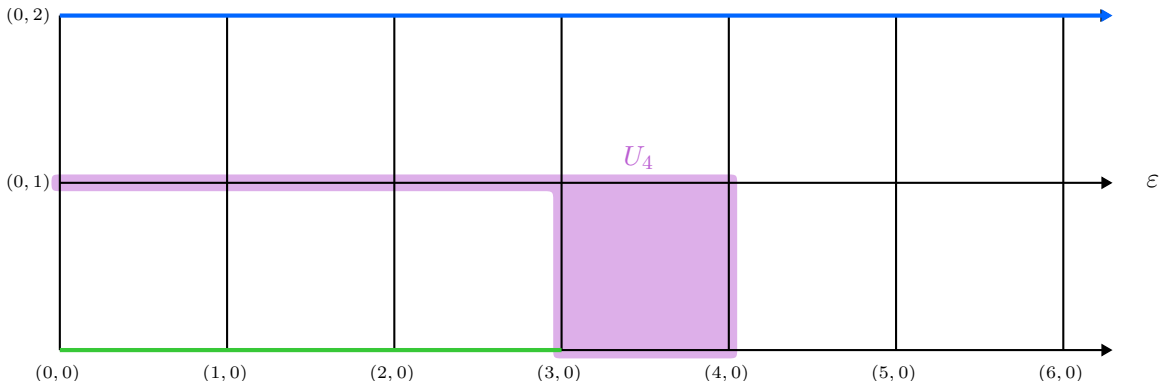

Figure 8.3: Depicted is the graph G' and the set U_4 in purple from Example 37. Indicated in blue is the ray $(0,2)(1,2)\dots$ that contains the vertex w and indicated in green is the path $(0,0)(1,0)(2,0)(3,0)$ that contains the vertex u in the case $m = 4$.

By the choice of 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 $(0,2) \in V(G{\downarrow}e) \subseteq V(G{\downarrow}f)$ 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. □

8.5 Upwards disjointness of adhesion sets

As mentioned in the introduction, we show in [AJKP24a, Theorem 1'] a more detailed version of Theorem 8.1.3 which, among others, yields that the adhesion sets of the tree-decomposition intersect 'not more than necessary':

Theorem 8.5.1. *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[4] set of G that is included in $G{\uparrow}e'$.*

Halin [Hal75, Theorem 2] showed that every locally finite, connected graph has a linked ray-decomposition[5] into finite parts with disjoint adhesion sets. He used this result in [Hal77,

[4] A set X of vertices of G is *critical* if infinitely many components of $G - X$ have neighbourhood X in G.
[5] A *ray-decomposition* is a tree-decomposition whose decomposition tree is a ray.

Satz 10] to establish Theorem 8.5.1 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 8.1.3 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 8.1.3 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 $G{\restriction}e'$, or is contained in a critical vertex set of G that is included in $G{\restriction}e'$.

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 8.1.3 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 8.5.2. *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 $c_{00}(\mathbb{N})$ denote the set of all sequences with values in $\mathbb{N}$ that are eventually zero. 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 $i < N$ and $s_i = 0$ for all $i \geqslant N$.[6] Let G be the graph depicted in Figure 8.4, 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 $k < n$. Intuitively, a sequence such as $\mathcal{S} = (3, 2, 3, 0, 0, \dots)$ encodes the specific $\mathbb{N} \times 3$ grid that one encounters when following along the (blue) $(0, 0, 0, \dots)$ grid, then turning after 3 steps into the (green) $(3, 0, 0, \dots)$ grid, then turning after another 2 steps into the (black) $(3, 2, 0, 0, \dots)$ grid, and finally turning after another 3 steps into the $(3, 2, 3, 0, 0, \dots)$ grid (not depicted). Note that G is locally finite and connected.

[6] We restrict to the set c' instead of $c_{00}(\mathbb{N})$ to ensure that G is connected.

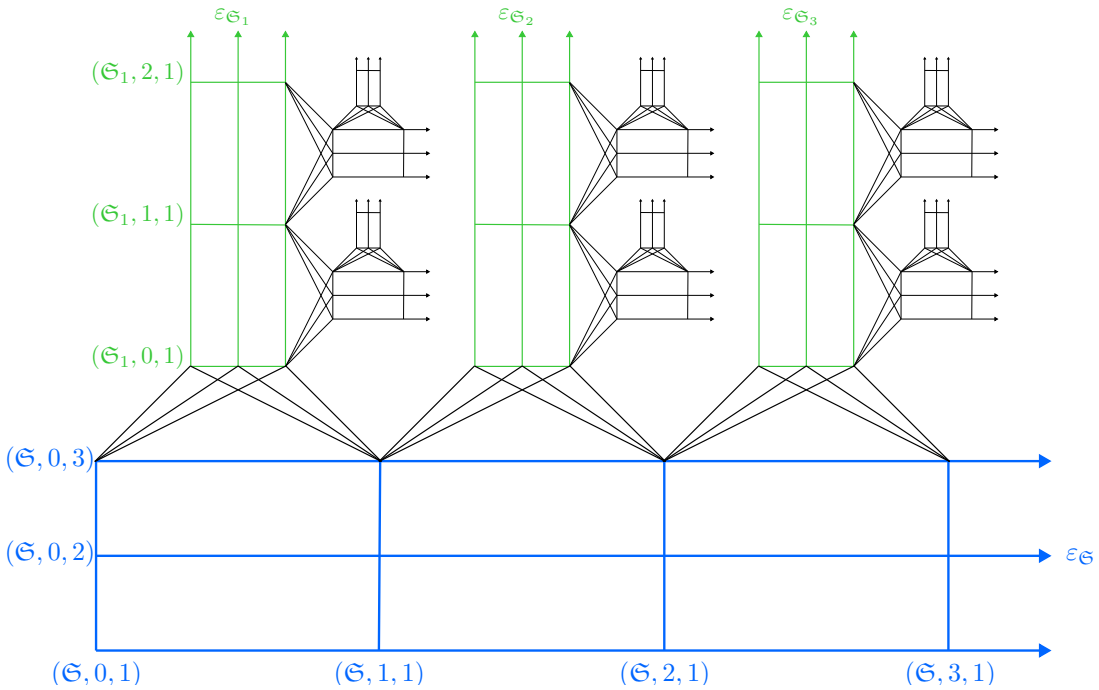

Figure 8.4: Depicted is the graph G from Example 8.5.2. The blue subgraph is induced by the vertices $(\mathfrak{S}, i, j)$ of G with $\mathfrak{S} = (0, 0, 0, \dots)$. The green subgraphs are induced by the vertices $(\mathfrak{S}_x, i, j)$ of G with $\mathfrak{S}_x = (x, 0, \dots)$ for $x \in \mathbb{N}_{\geqslant 1}$, respectively.

By Lemma 8.2.2 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 $V_{e_{\mathcal{S}}} \cap V(R_{\mathcal{S},3})$ (see Figure 8.5).

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 $\mathcal{S}_{n-1}$ is defined as $(s_0, \dots, s_{n-1}, 0, \dots)$. 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 8.5).

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. □

This example might explain why Halin never extended his result [Hal77, Satz 10] mentioned above to graphs with more than two ends: His precursor notion to a tree-decomposition, namely

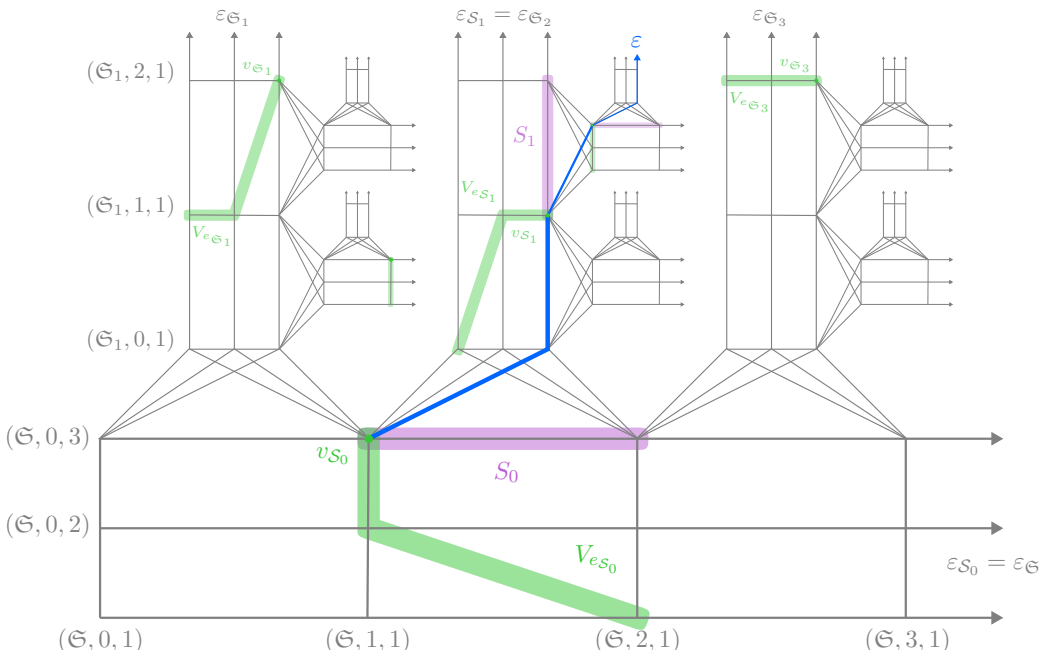

Figure 8.5: A ray in G that meets all the sets $V_{e_{\mathcal{S}_n}}$. The end to which it belongs has degree 2.

the quasi-trees and pseudo-trees discussed in [Hal91b], required upwards disjoint separators, and hence could not possibly capture the types of locally finite graphs described in Example 8.5.2.

CHAPTER 9

Efficiently distinguishing all tangles in locally finite graphs

This chapter is based on [CDHS14], which is joint work with Raphael W. Jacobs.

9.1 Introduction

In their Graph Minors Project [RS], Robertson and Seymour introduced the notion of tangles as a novel way to capture highly cohesive substructures, or 'clusters', of a graph. Tangles capture such clusters indirectly in that they encode the clusters' positions in the graph: a tangle is an orientation of all the low-order separations of the graph, orienting every such separation towards the side where the cluster is located. Robertson and Seymour's innovation was to view such 'consistent orientations' themselves as abstract clusters, thus providing a general framework for many concrete cluster types.

One of their key results about tangles is the *tree-of-tangles theorem*, which asserts that the tangles in a finite graph can be arranged in a tree-like way [RS91]. Diestel, Hundertmark and Lemanczyk found such a tree-like arrangement which is even *canonical*, that is, the automorphisms of the graph act naturally on it:

Theorem 9.1.1 ([DHL19], Theorem 3)**.** *Every connected finite graph has a canonical tree-decomposition which efficiently distinguishes all its tangles.*

Here, we say that a separation of G *distinguishes* two tangles in G if they orient that separation differently, and it does so *efficiently* if it has minimum order among all the separations of G that distinguish these two tangles. If two tangles are distinguished by some separation of G, they are called *distinguishable*. Now a tree-decomposition $(T, \mathcal{V})$ of a graph G *efficiently distinguishes* all the tangles in G if, for every two distinguishable tangles in G, there exists a separation of G which is induced by an edge of the decomposition tree T and efficiently distinguishes these two tangles.

The set of separations induced by a tree-decomposition is itself tree-like [Die18b; GK21] in that the separations are pairwise nested. In finite graphs, these two tree-like structures are equivalent, since the converse holds as well: every nested set of separations of a finite graph induces a tree-decomposition [CDHS14, Theorem 4.8] that gives rise to it. This converse, however, fails for locally finite graphs, potentially infinite graphs in which every vertex has only finitely many neighbours.

The resulting difference between nested sets of separations and tree-decompositions of locally finite graphs also appears in the context of tree-like arrangements of tangles. Elbracht, Kneip and Teegen showed on the one hand that Theorem 9.1.1 does not directly extend to such graphs:

they constructed a locally finite graph that does not have a tree-decomposition that efficiently distinguishes all its tangles, canonical or not [EKT22, Example 4.9][1]. On the other hand, they [EKT22, Theorem 6.6] and, independently, Carmesin, Hamann and Miraftab [CHM22, Theorem 7.3] proved that every locally finite graph still admits a *tree of tangles*, a nested set N of separations such that every two distinguishable tangles are efficiently distinguished by some element of N and every separation in N efficiently distinguishes some two tangles:

Theorem 9.1.2 ([CHM22; EKT22])**.** *Every connected locally finite graph has a canonical tree of tangles.*

So in order to understand to which locally finite graphs Theorem 9.1.1 extends, i.e., which locally finite graphs admit tree-decompositions that efficiently distinguish all their tangles, we have to investigate what prevents a tree of tangles, which exists by Theorem 9.1.2, from inducing a tree-decomposition. All known examples of locally finite graphs whose trees of tangles do not induce tree-decompositions share one feature: such graphs have a *thick* end, that is, an end containing infinitely many disjoint rays. In this chapter, we show that thick ends must appear in all such examples:

Theorem 38. *Every connected locally finite graph without thick ends has a canonical tree-decomposition which efficiently distinguishes all its tangles.*

Note that the converse of Theorem 38 does not hold, as the infinite grid witnesses. More generally, [DJKK22, Lemma 8.5] implies that every connected locally finite graph that is quasi-transitive and accessible admits a tree-decomposition as in Theorem 38 – no matter if it has a thick end or not.

For a proof of Theorem 38, it suffices to show that some (canonical) tree of tangles, which exists by Theorem 9.1.2, induces a (canonical) tree-decomposition if the locally finite graph has no thick end. We prove that, in fact, *every* tree of tangles of such a graph does so.

In order to determine whether a nested set N of separations, such as a tree of tangles, induces a tree-decomposition or not, we study the limit behaviour of sequences of oriented separations in N that are strictly increasing with respect to the natural partial order of oriented separations: Let the *limit* of a strictly increasing sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations be $(A, B) := (\bigcup_{i\in\mathbb{N}} A_i, \bigcap_{i\in\mathbb{N}} B_i)$, which is again a separation of the graph. Elbracht, Kneip and Teegen [EKT22, Lemma 2.7] showed that if the limit separation (A, B) of every strictly increasing sequence in a (canonical) nested set N of separations is trivial, in that $B = \emptyset$, then N induces a (canonical) tree-decomposition.

Our strategy for the proof Theorem 38 now is as follows. Let G be a locally finite graph, and let N be a tree of tangles of G. If N does not induce a tree-decomposition, then there exists a strictly increasing sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations in N whose limit (A, B) satisfies $B \neq \emptyset$. We show that the separator $A \cap B$ of this limit separation must be infinite. Then some end

[1] Carmesin, Hamann and Miraftab independently constructed such an example in their arXiv version [CHM22, Example 7.4] but it is not present in their journal version.

of G is in the closure of $A \cap B$ in the Freudenthal compactification of G. It will be easy to see that this end ω is unique.

The main effort of the proof will be to show that this end ω is thick. To do so, we carefully analyse the sequence $((A_i, B_i))_{i\in\mathbb{N}}$ and the interaction with its limit (A, B). In particular, we exploit the fact that $((A_i, B_i))_{i\in\mathbb{N}}$ consists of separations in the tree of tangles N, that is, each $\{A_i, B_i\}$ efficiently distinguishes some two tangles in G. Our *interlacing theorem*, Theorem 9.4.1, obtains from $((A_i, B_i))_{i\in\mathbb{N}}$ a strictly increasing sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ of separations in N with the same limit (A, B) and an accompanying sequence of tangles in G. This sequence of tangles not only witnesses that each $\{A'_i, B'_i\}$ efficiently distinguishes some two tangles in G, but is also 'interlaced' with $((A'_i, B'_i))_{i\in\mathbb{N}}$. With this structure at hand, we then deduce that the unique end ω in the closure of $A \cap B$ must be thick.

All in all, our proof method finds a thick end at the limit of every obstruction $((A_i, B_i))_{i\in\mathbb{N}}$ that prevents a tree of tangles from inducing a tree-decomposition:

Theorem 39. *Let N be a tree of tangles of a connected locally finite graph G, and let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence of separations in N with limit (A, B). If $B \neq \emptyset$, then G has a unique end that is in the closure of $A \cap B$, and this end is thick.*

This chapter is structured as follows. We first recall some relevant definitions and the corresponding notation in Section 9.2. In Section 9.3 we study limit separations and develop various general techniques for their investigation. Then we introduce our interlacing theorem in Section 9.4. Finally, we combine our results on limit separations with the interlacing theorem to prove our main results, Theorems 38 and 39, in Section 9.5.

9.2 Preliminaries

For the general definitions and notations around graphs, we follow [Die25]. In this section we recall the relevant definitions around separations of graphs and introduce related notation.

A *separation* of a graph G is an unordered pair $\{A, B\}$ of subsets of $V(G)$ such that $A \cup B = V(G)$ and there is no edge in G with one endvertex in $A \setminus B$ and the other one in $B \setminus A$. The *order* $|A, B|$ of a separation $\{A, B\}$ is the cardinality of its *separator* $A \cap B$.

The two *orientations* of a separation $\{A, B\}$ are the ordered pairs (A, B) and (B, A) which are *oriented separations* of G and whose *underlying* separation is $\{A, B\}$. Given a set S of separations, we write $\vec{S}$ for the set $\{(A, B), (B, A) \mid \{A, B\} \in S\}$ of all orientations of separations in S. An *orientation* of a set S of separations of G is a subset of $\vec{S}$ containing precisely one orientation of each separation in S.

We equip the set of oriented separations of a graph G with their natural partial order:

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

Two separations are *nested* if they have orientations which are comparable with respect to $\leqslant$, and a set of separations is *nested* if its elements are pairwise nested. We further say that two

oriented separations are *nested* if their underlying separations are nested. If two separations are not nested, then they *cross*. We will use the following observation which can easily be derived from the definition of the partial order of oriented separations.

Lemma 9.2.1 ([CDHS14], Statement (6))**.** *For two oriented separations (A, B) and (C, D) of a graph G, we have $(A, B) \leqslant (C, D)$ if and only if $(A \cap D) \smallsetminus S = \emptyset$ where $S := (A \cap B) \cap (C \cap D)$. In particular, two separations $\{A, B\}$ and $\{C, D\}$ of G cross if and only if all four sets $A \cap C$, $A \cap D$, $B \cap C$ and $B \cap D$ are non-empty after removing S.*

A separation $\{A, B\}$ of a graph G is *proper* if $A, B \neq V(G)$. Note that if (A, B) satisfies $(A, B) \leqslant (C, D)$ and $(A, B) \leqslant (D, C)$ for some separation $\{C, D\}$, then $B = V(G)$ and hence $\{A, B\}$ is not proper. It is immediate from the definition of separation that the separator $A \cap B$ of a proper separation $\{A, B\}$ of G is non-empty if G is connected.

For better readability, we will often use the term 'separation' for both separations and oriented separations, if the meaning is clear from the context. Analogously, we use terms such as 'order' which are defined for separations also for oriented separations if there is no ambiguity in how the definitions transfer.

We finally remark that whenever we speak of (sub)sequences, then we assume them to be infinite.

9.3 Limit separations

In this section, we study sequences of (oriented) separations and their limits. We begin by defining limit separations and then note some direct consequences of the definition. Next, we investigate the limit separations of strictly increasing sequences of separations which are 'non-exhaustive'. Finally, we describe how the relation of a finite-order separation to such a limit separation is connected to its relation with the separations in the sequence.

9.3.1 Definition and basic properties

Let $((A_i, B_i))_{i \in \mathbb{N}}$ be a sequence of separations of a graph G. Its *supremum* (A, B) (with respect to the partial order $\leqslant$ of oriented separations) is given by $A := \bigcup_{i \in \mathbb{N}} A_i$ and $B := \bigcap_{i \in \mathbb{N}} B_i$. We remark that $(A_i, B_i) \leqslant (A, B)$ for all $i \in \mathbb{N}$ by definition. The next lemma shows that (A, B) is again a separation of G. We will use it throughout this chapter tacitly.

Lemma 9.3.1. *Let $((A_i, B_i))_{i \in \mathbb{N}}$ be a sequence of separations of a graph G. Then its supremum (A, B) is a separation of G.*

Proof. Suppose for a contradiction that there exists an edge $e \in G$ with endvertices $a \in A \smallsetminus B$ and $b \in B \smallsetminus A$. It is immediate from the definition of A and B that $b \in B_i \smallsetminus A_i$ for all $i \in \mathbb{N}$. The definition of B also implies that there exists $I \in \mathbb{N}$ with $a \notin B_I$, and hence $a \in A_I \smallsetminus B_I$. So e is an edge of G joining $A_I \smallsetminus B_I$ and $B_I \smallsetminus A_I$, which contradicts that (A_I, B_I) is a separation of G. □

If a sequence of separations 'dominates' another sequence, then this transfers to their suprema. To make this precise, a sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations of a graph G *dominates* a sequence $((C_j, D_j))_{j\in\mathbb{N}}$ of separations of G if for every $j \in \mathbb{N}$ there exists $i = i(j) \in \mathbb{N}$ such that $(C_j, D_j) \leqslant (A_i, B_i)$. If two sequences of separations dominate each other, then they are *interlaced*. Note that an increasing sequence of separations is interlaced with each of its subsequences.

The next lemma is an immediate consequence of the above definitions. It describes how the suprema of two sequences of separations are related to each other if one sequence dominates the other or the sequences are interlaced.

Lemma 9.3.2. *Let $((A_i, B_i))_{i\in\mathbb{N}}$ and $((C_j, D_j))_{j\in\mathbb{N}}$ be two sequences of separations of a graph, and let (A, B) and (C, D) be their suprema, respectively. If $((A_i, B_i))_{i\in\mathbb{N}}$ dominates $((C_j, D_j))_{j\in\mathbb{N}}$, then $(C, D) \leqslant (A, B)$. In particular, if $((A_i, B_i))_{i\in\mathbb{N}}$ and $((C_j, D_j))_{j\in\mathbb{N}}$ are interlaced, then $(A, B) = (C, D)$.* □

While the above two lemmas deal with general sequences of separations, we are mainly interested in sequences which are increasing or even strictly increasing with respect to the partial order of oriented separations. For an increasing sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations of a graph G, we refer to its supremum (A, B) as its *limit*, as this separation is indeed the limit of the sequence with respect to pointwise convergence.

It is immediate from this definition of the limit separation (A, B) that, for every finite set $X \subseteq A$, there exists $I \in \mathbb{N}$ such that $X \subseteq A_i$ for all $i \geqslant I$, and that, for every finite set $X \subseteq A \setminus B = V(G) \setminus B$, there exists $I \in \mathbb{N}$ such that $X \subseteq V(G) \setminus B_i = A_i \setminus B_i$ for all $i \geqslant I$. Combining these two observations directly yields the following interplay of finite sets with the limit separation and the corresponding sequence.

Lemma 9.3.3. *Let $((A_i, B_i))_{i\in\mathbb{N}}$ be an increasing sequence of separations of a graph, and let (A, B) be its limit. For every finite set $X \subseteq A$, there exists $I \in \mathbb{N}$ such that, for all $i \geqslant I$, the intersections of X with $A \cap B$ and $A \setminus B$ equal the respective intersections with $A_i \cap B_i$ and $A_i \setminus B_i$.* □

9.3.2 Non-exhaustive sequences and their limit separations

Our main motivation for the study of limit separations of strictly increasing sequences comes from the next lemma by Elbracht, Kneip and Teegen [EKT22, Lemma 2.7] which establishes a condition for when a nested set of separations induces a tree-decomposition. We will later use this lemma in the proof of Theorem 38 to obtain a tree-decomposition as in Theorem 38 from the nested set given by Theorem 9.1.2.

To state this lemma, let us call a strictly increasing sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations of a graph G *exhaustive* if its limit (A, B) satisfies $B = \emptyset$ [DJKK22, Section 7.1]. Furthermore, a nested set N of separations of G is *(ω-)exhaustive* if every strictly increasing sequence in $\overrightarrow{N}$ is exhaustive.

Lemma 9.3.4 ([EKT22],Lemma 2.7)**.** *Let N be a nested set of proper finite-order separations of a connected graph G. If N is exhaustive, then it induces a tree-decomposition $(T, \mathcal{V})$ of G. Moreover, if N is canonical, then $(T, \mathcal{V})$ is canonical as well.*

In order to prove that a nested set N of separations is exhaustive, it suffices to consider certain 'representative' sequences in $\overrightarrow{N}$ with strong structural properties instead of just arbitrary ones. Lemma 9.3.2 provides us a tool for this: when we investigate the limit of a strictly increasing sequence in $\overrightarrow{N}$, then Lemma 9.3.2 allows us to instead consider any interlaced strictly increasing sequence as they have the same limit. As an important example of this, we will demonstrate such a use of Lemma 9.3.2 in combination with the next Lemma 9.3.5 in the paragraph after stating the latter.

To state Lemma 9.3.5, recall that, given a set X of vertices of a graph G, a component K of $G - X$ is *tight* if $N_G(K) = X$, and a separation $\{A, B\}$ of G is *tight* if $A \setminus B$ and $B \setminus A$ each contain the vertex set of a tight component of $G - (A \cap B)$ (cf. [EKT22, Section 6]). Note that tight separations are proper. We remark that the property 'tight' of a component K of $G - X$ heavily depends on the knowledge of the deleted set X via context.

Lemma 9.3.5 ([DJKK22], Lemma 7.5)**.** *Let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence of tight separations of a connected locally finite graph G. If there exists $K \in \mathbb{N}$ such that $|A_i, B_i| \leqslant K$ for all $i \in \mathbb{N}$, then $((A_i, B_i))_{i\in\mathbb{N}}$ is exhaustive.*

We use Lemma 9.3.2 to show that if a nested set N of separations consists of tight separations, then in order to determine whether $\overrightarrow{N}$ is exhaustive, it suffices to investigate strictly increasing sequences in $\overrightarrow{N}$ whose orders are strictly increasing. Note that by applying Lemma 9.3.2 again on a suitable subsequence we see that it suffices that the orders of the elements of every non-exhaustive strictly increasing sequence $(A_i, B_i)_{i\in\mathbb{N}}$ in $\overrightarrow{N}$ *cofinitely exceed every integer*, that is for every $k \in \mathbb{N}$ there exists $I \in \mathbb{N}$ such that $|A_i, B_i| \geqslant k$ for every $i \geqslant I$. Indeed, suppose that for some $k \in \mathbb{N}$ there exists no such $I \in \mathbb{N}$. Then there exists a subsequence of $(A_i, B_i)_{i\in\mathbb{N}}$ whose orders are bounded from above by k. This subsequence is exhaustive by Lemma 9.3.5 which contradicts by Lemma 9.3.2 that $(A_i, B_i)_{i\in\mathbb{N}}$ is not exhaustive.

In the remainder of this section, we show that if a strictly increasing sequence of tight separations is non-exhaustive, then its limit separation has infinite order.

Theorem 9.3.6. *Let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence of tight separations of a connected locally finite graph G, and let (A, B) be its limit. If $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive, then $A \cap B$ is infinite.*

Towards the proof of Theorem 9.3.6, we show two auxiliary results. The first lemma observes that $A \cap B$ is non-empty even under weaker assumptions.

Lemma 9.3.7. *Let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence of separations of a connected graph, and let (A, B) be its limit. If $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive, then $A \cap B$ is non-empty.*

Proof. Since the sequence $((A_i, B_i))_{i\in\mathbb{N}}$ is strictly increasing, some A_i, and thus A, is non-empty. The set B is non-empty, as $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive. So since G is connected and (A, B) is a separation of G, $A \cap B$ has to be non-empty as well. □

The second lemma shows that if the separations in a sequence are tight, then its supremum is again close to being tight.

Lemma 9.3.8. *Let $((A_i, B_i))_{i\in\mathbb{N}}$ be an increasing sequence of tight separations of a locally finite graph G, and let (A, B) be its supremum. If B is non-empty, then every $v \in A \cap B$ has a neighbour $w \in B \setminus A$ such that, for infinitely many $i \in \mathbb{N}$, w is contained in a tight component K_i of $G - (A_i \cap B_i)$ with $V(K_i) \subseteq B_i \setminus A_i$. In particular, we have $N_G(B \setminus A) = A \cap B$.*

Proof. Consider an arbitrary vertex $v \in A \cap B$. For every $i \in \mathbb{N}$, the separation $\{A_i, B_i\}$ is tight by assumption and hence there exists a tight component K_i of $G-(A_i \cap B_i)$ with $V(K_i) \subseteq B_i \setminus A_i$. Since G is locally finite, v has only finitely many neighbours. As $v \in A \cap B$, all but at most finitely many K_i contain v in their neighbourhood. Thus, the pigeonhole principle yields a neighbour w of v which is contained in infinitely many K_i; let us denote the collection of all such i by $\mathcal{I}$. Then we have $w \in \bigcap_{i\in\mathcal{I}} V(K_i) \subseteq \bigcap_{i\in\mathcal{I}}(B_i \setminus A_i) = B \setminus A$, where the last equality follows from Lemma 9.3.2 since $((A_i, B_i))_{i\in\mathcal{I}}$ is a subsequence of increasing sequence $((A_i, B_i))_{i\in\mathbb{N}}$. The 'in particular'-part follows immediately. □

We briefly remark that the above proof of Lemma 9.3.8 did not use the fact that there also exists a tight component of $G - (A_i \cap B_i)$ in $G[A_i \setminus B_i]$ for all $i \in \mathbb{N}$. Indeed, it is enough to assume that the separation $\{A_i, B_i\}$ is proper and that there is a tight component in $G[B_i \setminus A_i]$.

With Lemma 9.3.7 and Lemma 9.3.8 at hand, we are ready to prove Theorem 9.3.6.

Proof of Theorem 9.3.6. Suppose for a contradiction that $A \cap B$ is finite. Then Lemma 9.3.3 yields $I_1 \in \mathbb{N}$ such that $A \cap B \subseteq A_i \cap B_i$ for all $i \geqslant I_1$. Moreover, since $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive, Lemma 9.3.2 and Lemma 9.3.5 together imply that there is an integer $I_2 \geqslant I_1$ such that $|A_i, B_i| > |A, B|$ for all $i \geqslant I_2$; in particular, we have $A \cap B \subsetneq A_i \cap B_i$ for all $i \geqslant I_2$. Furthermore, $A \cap B \neq \emptyset$ by Lemma 9.3.7, and Lemma 9.3.8 thus yields a neighbour $w \in B \setminus A$ of $A \cap B$ and an infinite set $\mathcal{I} \subseteq \mathbb{N}$ such that w is contained in a tight component K_i of $G - (A_i \cap B_i)$ with $V(K_i) \subseteq B_i \setminus A_i$ for all $i \in \mathcal{I}$.

Fix any integer $i \in \mathcal{I}$ with $i \geqslant I_2$. Since $i \geqslant I_2$, there exists some vertex $u \in (A_i \cap B_i) \setminus (A \cap B)$. So $u \in A_i \cap B_i = N_G(K_i)$, since K_i is a tight component of $G - (A_i \cap B_i)$, and $w \in K_i$. Thus, there exists a u–w path P in G such that all vertices of P except u are contained in K_i. We claim that P contradicts that (A, B) is a separation of G.

To see this, recall that $w \in B \setminus A$. We also have $u \in A \setminus B$, as $u \in A_i \setminus (A \cap B)$ by its choice and $A_i \subseteq A$. Moreover, K_i is a component of $G - (A_i \cap B_i)$ and hence disjoint from $A_i \cap B_i$ and in particular from $A \cap B$, since $A \cap B \subseteq A_i \cap B_i$ for $i \geqslant I_2 \geqslant I_1$. Altogether, the path P joins $u \in A \setminus B$ and $w \in B \setminus A$, but avoids $A \cap B$ which contradicts that (A, B) is a separation of G. □

By Theorem 9.3.6, the separator $A \cap B$ of the limit separation (A, B) of a non-exhaustive strictly increasing sequence of separations is infinite. As we will see in Section 9.5, this is already the first step towards Theorem 39 which asserts that the graph has a thick end which is in the closure of $A \cap B$ in the Freudenthal compactification of G. Indeed, for every infinite

set of vertices in a connected locally finite graph, such as $A \cap B$, there exists an end which is in the closure of this set. So it remains to show that the end arising in this way is thick which will be the main challenge in Section 9.5.

9.3.3 Interplay with the partial order of oriented separations

Consider the limit separation (A, B) of an increasing sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations, and a further fixed separation $\{C, D\}$. Given the relation of (A, B) and the orientations of $\{C, D\}$ with respect to the partial order of oriented separations, what can we say about the relation of the (A_i, B_i) and the orientations of $\{C, D\}$? This question is of particular relevance for the analysis of limit separations that we will perform in Section 9.5, and we shall address this question here.

We start with the case that no orientation of a separation $\{C, D\}$ is comparable with the limit separation (A, B), i.e. that $\{C, D\}$ crosses $\{A, B\}$.

Lemma 9.3.9. *Let $((A_i, B_i))_{i\in\mathbb{N}}$ be an increasing sequence of separations of a graph G, and let (A, B) be its limit. If a separation $\{C, D\}$ of G crosses $\{A, B\}$, then there exists $I \in \mathbb{N}$ such that $\{C, D\}$ crosses $\{A_i, B_i\}$ for all $i \geqslant I$.*

Proof. We prove the contrapositive of this statement. Assume that a separation $\{C, D\}$ of G is nested with infinitely many $\{A_i, B_i\}$. Denote the infinite set of these indices i by I. First, assume that there is some $j \in I$ such that $(C, D) \leqslant (A_j, B_j)$ or $(D, C) \leqslant (A_j, B_j)$. Since $((A_i, B_i))_{i\in\mathbb{N}}$ is increasing, $(C, D) \leqslant (A_i, B_i)$ for every $i \in \mathbb{N}$ with $i \geqslant j$. Thus, $(C, D) \leqslant (A, B)$ or $(D, C) \leqslant (A, B)$, respectively. Otherwise, we have $(A_j, B_j) \leqslant (C, D)$ or $(A_j, B_j) \leqslant (D, C)$ for every $j \in I$. Since I is infinite, we may assume by symmetry that there are infinitely many $j \in I$ with $(A_j, B_j) \leqslant (C, D)$. Thus, $(A, B) \leqslant (C, D)$. □

We now turn to the case that some orientation of a separation $\{C, D\}$ is comparable with the limit separation (A, B), i.e. that $\{C, D\}$ is nested with $\{A, B\}$. It is immediate from the definition of limit separation that if an orientation of $\{C, D\}$, say (C, D), is greater than (A, B), then (C, D) is also greater than all the (A_i, B_i) in the sequence defining (A, B).

It remains to consider the case that an orientation of a separation $\{C, D\}$, say (C, D), is smaller than the limit separation (A, B). For this case, we restrict our attention to finite-order separations (C, D), as infinite-order separations do not allow any conclusion on their relation with the (A_i, B_i) in the sequence defining (A, B). The next two lemmas describe the relation of (C, D) with the (A_i, B_i), depending on whether the sequence $((A_i, B_i))_{i\in\mathbb{N}}$ is exhaustive or not.

First, let us assume that $((A_i, B_i))_{i\in\mathbb{N}}$ is exhaustive. Then both orientations of $\{C, D\}$ are smaller than (A, B). Lemma 9.3.10 asserts that if all the B_i are connected, then some orientation of $\{C, D\}$ is also smaller than (A_i, B_i) for all but finitely many $i \in \mathbb{N}$.

Lemma 9.3.10. *Let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence of separations of a graph G such that $G[B_i]$ is connected for every $i \in \mathbb{N}$, and let $\{C, D\}$ be a finite-order separation of G.*

If $((A_i, B_i))_{i\in\mathbb{N}}$ is exhaustive, then there is an $I \in \mathbb{N}$ such that either $(C, D) \leqslant (A_i, B_i)$ for all $i \geqslant I$ or $(D, C) \leqslant (A_i, B_i)$ for all $i \geqslant I$.

Proof. Since $((A_i, B_i))_{i\in\mathbb{N}}$ is exhaustive, we have $B = \emptyset$ and hence $C \cap D \subseteq A = A \setminus B$. By assumption, $C \cap D$ is finite, so we can apply Lemma 9.3.3 to $C \cap D$ and obtain an $I \in \mathbb{N}$ with $C \cap D \subseteq A_I \setminus B_I$. Since $G[B_I]$ is connected by assumption and $C \cap D$ is disjoint from B_I, this implies that B_I is contained either in $C \setminus D$ or in $D \setminus C$. If $B_I \subseteq C \setminus D$, then $A_I \supseteq V(G) \setminus B_I \supseteq D$ and hence $(D, C) \leqslant (A_I, B_I)$; analogously, we have $(C, D) \leqslant (A_I, B_I)$ if $B_I \subseteq D \setminus C$. The claim then follows directly from the fact that the sequence $((A_i, B_i))_{i\in\mathbb{N}}$ is increasing. □

Secondly, we consider the case that $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive. For sequences of tight separations in connected locally finite graphs, Lemma 9.3.11 then asserts that (C, D) is smaller than all but finitely many (A_i, B_i).

Lemma 9.3.11. *Let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence of tight separations of a connected locally finite graph G, let (A, B) be its limit, and let (C, D) be any finite-order separation of G. If $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive and $(C, D) \leqslant (A, B)$, then there exists an $I \in \mathbb{N}$ such that $(C, D) \leqslant (A_I, B_I)$.*

Proof. The sequence $((A_i, B_i))_{i\in\mathbb{N}}$ is increasing, so it suffices to find a single $I \in \mathbb{N}$ with $(C, D) \leqslant (A_I, B_I)$. To do so, let $S_i := (A_i \cap B_i) \cap (C \cap D)$ for every $i \in \mathbb{N}$. We show the existence of some $I \in \mathbb{N}$ with $(C \cap B_I) \setminus S_I = \emptyset$, which then yields the claim by Lemma 9.2.1. More precisely, we successively find conditions on $i \in \mathbb{N}$ for the emptiness of each of the three sets $X_i := ((C \cap B_i) \setminus S_i) \cap D$, $Y_i := ((C \cap B_i) \setminus S_i) \cap A_i$ and $Z_i := ((C \cap B_i) \setminus S_i) \setminus (A_i \cup D)$ (see Figure 9.1), which partition $(C \cap B_i) \setminus S_i$, and finally give some $I \in \mathbb{N}$ satisfying all these conditions.

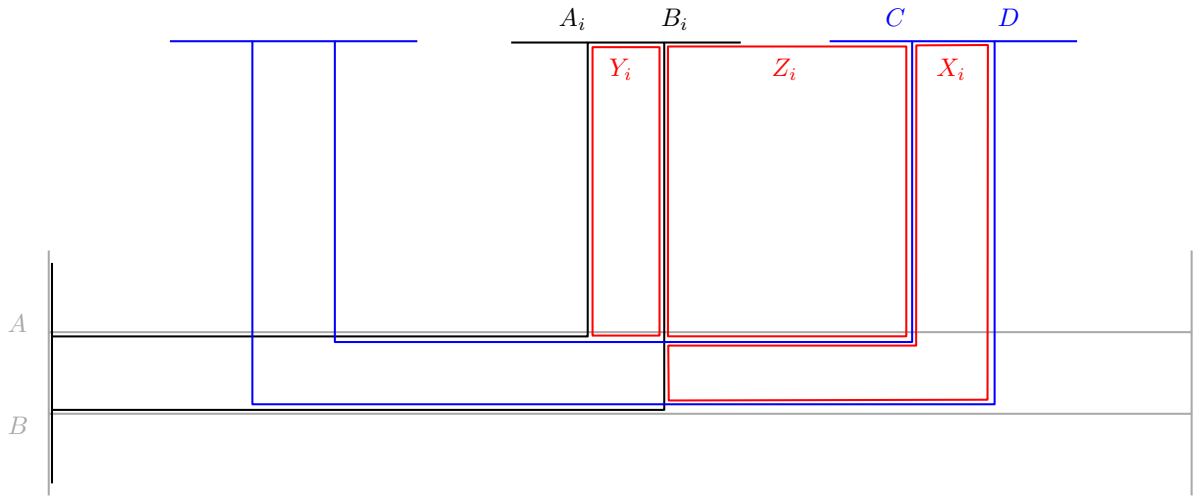

Figure 9.1: A diagram depicting the three sets X_i, Y_i and Z_i from the proof of Lemma 9.3.11.

First, we consider the set $X_i = ((C \cap B_i) \setminus S_i) \cap D$. Since $(C, D) \leqslant (A, B)$, we have $C \cap D \subseteq A$.

So since $C \cap D$ is finite, Lemma 9.3.3 yields $I_1 \in \mathbb{N}$ such that $C \cap D \subseteq A_i$ for all $i \geqslant I_1$. This immediately implies $X_i = (C \cap D) \cap (B_i \setminus A_i) = \emptyset$ for all $i \geqslant I_1$.

Next, we turn to $Y_i = ((C \cap B_i) \setminus S_i) \cap A_i$. Recall that $A \cap B \neq \emptyset$ by Lemma 9.3.7. So Lemma 9.3.8 yields a neighbour $w \in B \setminus A$ of $A \cap B$ and an infinite set $\mathcal{I} \subseteq \mathbb{N}$ such that w is contained in a tight component K_i of $G - (A_i \cap B_i)$ with $V(K_i) \subseteq B_i \setminus A_i$ for all $i \in \mathcal{I}$. Now let $i \in \mathcal{I}$ with $i \geqslant I_1$, and suppose for a contradiction that Y_i is non-empty. Since $N_G(K_i) = A_i \cap B_i$, there thus exists a path P from w to the set $Y_i \subseteq A_i \cap B_i$ whose vertices lie in K_i except its last vertex u. Note that $u \in C \setminus D$, as $u \in Y_i = (C \cap B_i \cap A_i) \setminus D$, and $w \in D \setminus C$, as $w \in B \setminus A$ and $(C, D) \leqslant (A, B)$. Moreover, the path P does not meet $C \cap D$ as $(C \cap D) \cap K_i \subseteq (C \cap D) \cap (B_i \setminus A_i) = X_i$ is empty for $i \geqslant I_1$, as shown above. Altogether, P starts in $w \in D \setminus C$ and ends in $u \in C \setminus D$, but P avoids $C \cap D$, which contradicts that (C, D) is a separation of G.

It remains to consider $Z_i = ((C \cap B_i) \setminus S_i) \setminus (A_i \cup D) = (C \cap B_i) \setminus (A_i \cup D)$. We first prove that if Z_i is non-empty for $i \in \mathcal{I}$ with $i \geqslant I_1$, then it must contain a neighbour of S_i. Secondly, we show that Z_i cannot contain a neighbour of S_i for large enough i, which yields the desired emptiness of Z_i for large enough $i \in \mathcal{I}$.

First, let us prove that if Z_i is non-empty for $i \in \mathcal{I}$ with $i \geqslant I_1$, then it contains a neighbour of S_i. It is easy to check that $(A_i \cup D, C \cap B_i)$ is again a separation of G with separator $(C \cap B_i) \cap (A_i \cup D)$. As shown above, the two sets $X_i = ((C \cap B_i) \setminus S_i) \cap A_i$ and $Y_i = ((C \cap B_i) \setminus S_i) \cap D$ are empty since $i \in \mathcal{I}$ with $i \geqslant I_1$. Thus, the separator of $(A_i \cup D, C \cap B_i)$ is in fact S_i. Both $Z_i = (C \cap B_i) \setminus (A_i \cup D)$ and $(A_i \cup D) \setminus (C \cap B_i)$ are non-empty, where the latter follows from (A_i, B_i) being tight and hence $A_i \setminus B_i$ being non-empty. So since $(A_i \cup D, C \cap B_i)$ is a separation of the connected graph G, its separator S_i has to be non-empty and Z_i contains a neighbour of S_i.

Secondly, let us show that no neighbour of S_i is contained in Z_i for large enough i. By applying Lemma 9.3.3 to the set $C \cap D$, which is finite since (C, D) has finite order, we obtain $I_2 \in \mathbb{N}$ such that S_i equals $S := (A \cap B) \cap (C \cap D)$ for all $i \geqslant I_2$. Thus, it suffices to argue that no neighbour of S is contained in Z_i for large i. Consider any neighbour u of S which is contained in C, and hence in A, since $(C, D) \leqslant (A, B)$. Then there exists $I_u \in \mathbb{N}$ such that $u \in A_i$ for all $i \geqslant I_u$, and hence $u \notin B_i \setminus A_i \supseteq (C \cap B_i) \setminus (A_i \cup D) = Z_i$ for $i \geqslant I_u$. Let I_3 be the maximum over all I_u, and note that this maximum exists since the finite set S has only finitely many neighbours in the locally finite graph G. Then Z_i contains no neighbour of S for all $i \geqslant I_3$.

Choosing $I \in \mathcal{I}$ with $I \geqslant \max\{I_1, I_2, I_3\}$ thus completes the proof. □

9.4 The interlacing theorem

Interlaced strictly increasing sequences of separations have the same limit, as Lemma 9.3.2 shows. So when we investigate the limit separation of a strictly increasing sequence of separations, then we may instead consider it as the limit of an interlaced sequence with stronger structural properties. One example for this approach is the combination of Lemma 9.3.2 and Lemma 9.3.5

which shows that, given a non-exhaustive strictly increasing sequence, there exists an interlaced sequence, and even a subsequence of the given one, whose separations have strictly increasing order. We remark that this technique was already crucial to the proof of [DJKK22, Lemma 7.1].

In this section we introduce the *interlacing theorem* which enables us to use the above described method of interlaced sequences for strictly increasing sequences of 'tangle-distinguishing' separations. Before we are ready to state the interlacing theorem, we introduce the notion of pre-tangles; with all the wording below, we imitate the special case of tangles (cf. [Die25, Section 12.5]). A set O of oriented separations of a graph G is *consistent* if we do not have $(B, A) \leqslant (C, D)$ for any two $(A, B), (C, D) \in O$ with $\{A, B\} \neq \{C, D\}$. A set P of oriented separations of G is a *pre-tangle in* G if there exists $k \in \mathbb{N} \cup \{\aleph_0\}$ such that P is a consistent orientation of the set of all the separations of G of order less than k. It is immediate from this definition that if a pre-tangle P in G orients some separation, then it also orients every separation of at most the same order.

A separation of G *distinguishes* two pre-tangles P and P' in G if both P and P' contain one of its orientations but not the same, and the separation does so *efficiently* if it has minimum order among all the separations distinguishing P and P'. Note that every separation which distinguishes two pre-tangles has finite order by the definition of pre-tangle. Given a set $\mathcal{P}$ of pre-tangles in G, a set N of separations of G *(efficiently) distinguishes* $\mathcal{P}$ if every two pre-tangles in $\mathcal{P}$ that are distinguished by some separation of G are (efficiently) distinguished by some separation in N. Moreover, a separation of G is *$\mathcal{P}$-relevant* if it efficiently distinguishes some two pre-tangles in $\mathcal{P}$.

The interlacing theorem roughly speaking asserts that, given a set $\mathcal{P}$ of pre-tangles in a graph G and a nested set N of $\mathcal{P}$-relevant separations which distinguishes $\mathcal{P}$ efficiently, every non-exhaustive strictly increasing sequence in $\vec{N}$ admits an interlaced sequence in $\vec{N}$ whose $\mathcal{P}$-relevance is witnessed in a very strong sense by a corresponding sequence in $\mathcal{P}$. More precisely, the interlacing theorem reads as follows.

Theorem 9.4.1 (Interlacing Theorem)**.** *Let $\mathcal{P}$ be a set of pre-tangles in a connected locally finite graph G, let N be a nested set of $\mathcal{P}$-relevant proper separations of G which efficiently distinguishes $\mathcal{P}$, and let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence in $\vec{N}$. If $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive, then there exists a strictly increasing sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ in $\vec{N}$ that is interlaced with $((A_i, B_i))_{i\in\mathbb{N}}$ and a sequence $(P_i)_{i\in\mathbb{N}}$ in $\mathcal{P}$ such that*

- *(IM1) $(B'_i, A'_i) \in P_i$ and $(A'_i, B'_i) \in P_{i+1}$ for all $i \in \mathbb{N}$,*
- *(IM2) every minimal-order separation among $\{A'_i, B'_i\}, \dots, \{A'_{j-1}, B'_{j-1}\}$ efficiently distinguishes P_i and P_j for all $i, j \in \mathbb{N}$ with $i < j$.*

Moreover, the $\{A'_i, B'_i\}$ can be chosen such that $(|A'_i, B'_i|)_{i\in\mathbb{N}}$ is strictly increasing, and then $\{A'_i, B'_i\}$ efficiently distinguishes P_i from every P_j with $j > i$.

The remainder of this section is devoted to the proof of Theorem 9.4.1.

For better readability, we shall in the following denote separations by lower-case letters such as s and its orientations as $\vec{s}$ and $\overleftarrow{s}$. Once we denote one orientation of s as $\vec{s}$, we write $\overleftarrow{s}$

for the other one, and vice-versa. Analogously, we denote oriented separations as lower-case letters with a forward or backward arrow on top such as $\vec{s}$ and $\overleftarrow{s}$, and we write s for their underlying separation. The order of a separation s is denoted by $|s|$.

The proof of Theorem 9.4.1 is divided into the four Lemmas 9.4.3 to 9.4.6; each of these lemmas may be of independent interest for the analysis of strictly increasing sequences. As one ingredient for their proofs, we have the following easy observation about those separations in a nested set that distinguish two fixed pre-tangles.

Lemma 9.4.2. *Let N be a nested set of separations of a graph G. Given two pre-tangles P and Q in G, the set of all the oriented separations $\vec{s} \in P \cap \overrightarrow{N}$ with $\overleftarrow{s} \in Q$ forms a chain.*

Proof. In order to show that this set is a chain, we have to prove that for any two $\vec{r}, \vec{s} \in P \cap \overrightarrow{N}$, we have either $\vec{r} < \vec{s}$ or $\vec{s} < \vec{r}$. Since r and s are contained in N and hence nested, one of $\vec{r} < \vec{s}$, $\vec{s} < \vec{r}$, $\overleftarrow{r} < \vec{s}$ and $\vec{s} < \overleftarrow{r}$ holds. The first and the second case are as desired, so it remains to show that the third and fourth one do not occur. We cannot have $\overleftarrow{r} < \vec{s}$, as both $\vec{r}$ and $\vec{s}$ are contained in P which is consistent. For $\vec{s} < \overleftarrow{r}$, note that this implies $\vec{r} < \overleftarrow{s}$ which cannot be the case by the consistency of Q, in which both $\overleftarrow{s}$ and $\overleftarrow{r}$ are contained. □

The first Lemma 9.4.3 shows that for the proof of Theorem 9.4.1, we can replace (IM2) with a weaker property (IM2') in terms of consecutive pre-tangles.

Lemma 9.4.3. *Let $(\vec{s}_i)_{i\in\mathbb{N}}$ be a strictly increasing sequence of separations of a graph G, and let $(P_i)_{i\in\mathbb{N}}$ be a sequence of pre-tangles in G. If $(\vec{s}_i)_{i\in\mathbb{N}}$ and $(P_i)_{i\in\mathbb{N}}$ satisfy (IM1) and (IM2'), where*

(IM2') the separation s_i efficiently distinguishes P_i and P_{i+1} for all $i \in \mathbb{N}$,

then they also satisfy (IM2).

Proof. To show (IM2), consider any two integers $i, j \in \mathbb{N}$ with $i < j$, and let $\ell \in \{i, \dots, j-1\}$ such that s_ℓ has minimal order among $s_i, \dots, s_{j-1}$. Then both pre-tangles P_i and P_j in G orient s_ℓ, since P_i orients s_i by (IM1) and $|s_\ell| \leqslant |s_i|$ by the choice of ℓ and analogously for P_j and s_{j-1}. Since $\overleftarrow{s}_i \in P_i$ and $\vec{s}_{j-1} \in P_j$ by (IM1), the consistency of P_i and P_j implies that $\overleftarrow{s}_\ell \in P_i$ and $\vec{s}_\ell \in P_j$, so s_ℓ distinguishes P_i and P_j.

It remains to show that s_ℓ distinguishes P_i and P_j even efficiently. Let r be an arbitrary separation of G with $|r| < |s_\ell|$. By the choice of ℓ, we have $|r| < |s_\ell| \leqslant |s_k|$ for all $k \in \{i, \dots, j-1\}$. So for every $k \in \{i, \dots, j-1\}$, both P_k and P_{k+1} orient r, and they both contain the same orientation of r by (IM2'). Repeatedly applying this fact from $k = i$ up to $k = j-1$ implies that P_i and P_j contain the same orientation of r. Thus, r does not distinguish P_i and P_j. So no separation of G of order less than $|s_\ell|$ distinguishes P_i and P_j; in particular, s_ℓ distinguishes P_i and P_j efficiently. □

The second Lemma 9.4.4 asserts, in the setting of Theorem 9.4.1, that for every strictly increasing sequence $(\vec{s}_i)_{i\in\mathbb{N}}$ in $\overrightarrow{N}$, there exists a corresponding sequence in $\mathcal{P}$ such that they together satisfy (IM1) and (IM2') (and thus also (IM2) by Lemma 9.4.3) – if every pair $\vec{s}_i < \vec{s}_{i+1}$

is $\mathcal{P}$-relevant in a strong form. To make this precise, recall that a separation of a graph G is *$\mathcal{P}$-relevant* for a set $\mathcal{P}$ of pre-tangles in G if it efficiently distinguishes some two pre-tangles in $\mathcal{P}$. We now call a pair $\vec{s} < \vec{t}$ of two oriented separations of G *strongly $\mathcal{P}$-relevant* if there exist pre-tangles $O, P, Q \in \mathcal{P}$ such that

- s efficiently distinguishes O and P with $\vec{s} \in P$, and
- t efficiently distinguishes P and Q with $\overleftarrow{t} \in P$.

A strictly increasing sequence $(\vec{s}_i)_{i\in\mathbb{N}}$ of separations of G is then called *strongly $\mathcal{P}$-relevant* if every two consecutive separations $\vec{s}_i < \vec{s}_{i+1}$ are strongly $\mathcal{P}$-relevant.

Lemma 9.4.4. *Let $\mathcal{P}$ be a set of pre-tangles in a graph G, let N be a nested set of separations of G which efficiently distinguishes $\mathcal{P}$, and let $(\vec{s}_i)_{i\in\mathbb{N}}$ be a strictly increasing sequence of separations in $\overrightarrow{N}$. If $(\vec{s}_i)_{i\in\mathbb{N}}$ is strongly $\mathcal{P}$-relevant, then there exists a sequence $(P_i)_{i\in\mathbb{N}}$ in $\mathcal{P}$ which together with $(\vec{s}_i)_{i\in\mathbb{N}}$ satisfies (IM1) and (IM2').*

Proof. Since $(\vec{s}_i)_{i\in\mathbb{N}}$ is strongly $\mathcal{P}$-relevant, there exist, for every $i \in \mathbb{N}$, pre-tangles O_{i+1}, P_{i+1} and Q_{i+1} in $\mathcal{P}$ with

(i) s_i efficiently distinguishes O_{i+1} and P_{i+1} with $\vec{s}_i \in P_{i+1}$, and

(ii) s_{i+1} efficiently distinguishes P_{i+1} and Q_{i+1} with $\overleftarrow{s}_{i+1} \in P_{i+1}$.

Set $P_0 := O_1$. Then the sequence $(P_i)_{i\in\mathbb{N}}$ together with $(\vec{s}_i)_{i\in\mathbb{N}}$ immediately satisfies (IM1) by (i) and (ii). We prove that they also satisfy (IM2').

By (IM1), s_i distinguishes P_i and P_{i+1}. Thus, it remains to show that s_i does so efficiently. Suppose for a contradiction that this is not the case. We then have $i \geqslant 1$ by the definition of P_0 as O_1 and (i). Since N efficiently distinguishes $\mathcal{P}$, there is a separation $r \in N$ that efficiently distinguishes P_i and P_{i+1}; in particular $|r| < |s_i|$. Let $\vec{r}$ be the orientation of r in P_{i+1}, then $\overleftarrow{r} \in P_i$. Both $\vec{s}_i$ and $\vec{r}$ are contained in $P_{i+1} \cap \overrightarrow{N}$, so Lemma 9.4.2 implies that either $\vec{r} < \vec{s}_i$ or $\vec{s}_i < \vec{r}$. We show that both cases lead to contradictions.

First, suppose $\vec{r} < \vec{s}_i$. Since $i \geqslant 1$, the pre-tangle Q_i is defined, and we have $\vec{r} \in Q_i$, as $\vec{s}_i \in Q_i$ by the choice of Q_i and Q_i is consistent. But $\overleftarrow{r} \in P_i$, so r distinguishes P_i and Q_i. This then contradicts (ii) as $|r| < |s_i|$.

Secondly, assume that $\vec{s}_i < \vec{r}$. This implies $\overleftarrow{r} < \overleftarrow{s}_i$. We proceed analogously to the first case, but for O_{i+1}. We have $\overleftarrow{s}_i \in O_{i+1}$ by the choice of O_{i+1}, and hence $\overleftarrow{r} \in O_{i+1}$, since O_{i+1} is consistent. But $\vec{r} \in P_{i+1}$, so r distinguishes O_{i+1} and P_{i+1} which contradicts (i) due to $|r| < |s_i|$. □

The third Lemma 9.4.5 asserts, in the setting of Theorem 9.4.1, that a strictly increasing sequence of $\mathcal{P}$-relevant separations with strictly increasing orders is interlaced with a strongly $\mathcal{P}$-relevant strictly increasing sequence.

Lemma 9.4.5. *Let $\mathcal{P}$ be a set of pre-tangles in a graph G, and let N be a nested set of $\mathcal{P}$-relevant proper separations of G which efficiently distinguishes $\mathcal{P}$. Let $(\vec{s}_i)_{i\in\mathbb{N}}$ be a strictly*

increasing sequence in $\overrightarrow{N}$. If the orders of the s_i are strictly increasing, then there exists a strictly increasing sequence $(\overrightarrow{s'_i})_{i\in\mathbb{N}}$ in $\overrightarrow{N}$ that contains $(\vec{s}_i)_{i\in\mathbb{N}}$ as a subsequence and is strongly $\mathcal{P}$-relevant.

Proof. We recursively find a strictly increasing function $f\colon \mathbb{N}\to\mathbb{N}$ and a strictly increasing sequence $(\overrightarrow{s'_i})_{i\in\mathbb{N}}$ in $\overrightarrow{N}$ which is strongly $\mathcal{P}$-relevant such that for every $k\in\mathbb{N}$, we have

(i) $\overrightarrow{s'_{f(k)}} = \vec{s}_k$, and

(ii) the pair $\overrightarrow{s'_i} < \overrightarrow{s'_{i+1}}$ is strongly $\mathcal{P}$-relevant for all $i < f(k)$.

Clearly, such a sequence $(\overrightarrow{s'_i})_{i\in\mathbb{N}}$ is as desired.

To start the recursion, we set $f(0) := 0$ and $\overrightarrow{s'_0} := \vec{s}_0$. Let $\ell\in\mathbb{N}$ such that we have already defined $f(0),\dots,f(\ell)$ and $(\overrightarrow{s'_i})_{i\leqslant f(\ell)}$ which satisfy properties (i) and (ii) for every $k\leqslant\ell$. We now proceed by defining $f(\ell+1) > f(\ell)$ and $\overrightarrow{s'_i}$ for the remaining $i = f(\ell)+1,\dots,f(\ell+1)$ such that $(\overrightarrow{s'_i})_{i\leqslant f(\ell+1)}$ satisfies the properties (i) and (ii) for every $k\leqslant\ell+1$.

The recursion step roughly works as follows. Appending the separation $\vec{s}_{\ell+1}$ to $(\overrightarrow{s'_i})_{i\leqslant f(\ell)}$ would immediately yield (i). Then either the pair $\vec{s}_\ell < \vec{s}_{\ell+1}$ is strongly $\mathcal{P}$-relevant itself, which implies (ii), or we will insert an additional separation $\vec{r}$ in between $\vec{s}_\ell$ and $\vec{s}_{\ell+1}$ to then get two strongly $\mathcal{P}$-relevant pairs $\vec{s}_\ell < \vec{r}$ and $\vec{r} < \vec{s}_{\ell+1}$ and hence (ii), which will then conclude the proof.

Since both $\vec{s}_\ell$ and $\vec{s}_{\ell+1}$ are $\mathcal{P}$-relevant, there exist pre-tangles P_j and Q_j in $\mathcal{P}$ for $j=\ell,\ell+1$ such that s_j efficiently distinguishes P_j and Q_j with $\vec{s}_j\in Q_j$. Since we have $|s_\ell| < |s_{\ell+1}|$ by assumption, the pre-tangles $P_{\ell+1}$ and $Q_{\ell+1}$ in G orient not only $s_{\ell+1}$, but also s_ℓ, and they do so in the same way as $s_{\ell+1}$ efficiently distinguishes them. We have $\vec{s}_\ell\in Q_{\ell+1}$, as $Q_{\ell+1}$ is consistent and $\vec{s}_\ell < \vec{s}_{\ell+1}\in Q_{\ell+1}$, and hence $\vec{s}_\ell\in P_{\ell+1}$. Altogether, s_ℓ distinguishes P_ℓ and $P_{\ell+1}$.

First, let us assume that s_ℓ distinguishes P_ℓ and $P_{\ell+1}$ efficiently. Then the pre-tangles P_ℓ, $P_{\ell+1}$, $Q_{\ell+1}$ in $\mathcal{P}$ witness the strong $\mathcal{P}$-relevance of the pair $\vec{s}_\ell < \vec{s}_{\ell+1}$. So in this case, we set $f(\ell+1) := f(\ell)+1$ and $\overrightarrow{s'_{f(\ell+1)}} := \vec{s}_{\ell+1}$. Then $(\overrightarrow{s'_i})_{i\leqslant f(\ell+1)}$ satisfies (i) and (ii) for every $k\leqslant\ell+1$, completing the recursion step.

Secondly, we assume that s_ℓ does not distinguish P_ℓ and $P_{\ell+1}$ efficiently. Since N efficiently distinguishes $\mathcal{P}$, there exists a separation $r\in N$ which distinguishes P_ℓ and $P_{\ell+1}$ efficiently; in particular, $|r| < |s_\ell|$. Let $\vec{r}$ be the orientation of r in $P_{\ell+1}$ and $\overleftarrow{r}$ the one in P_ℓ. In the remainder of the proof, we show that setting $f(\ell+1) := f(\ell)+2$, $\overrightarrow{s'_{f(\ell)+1}} := \vec{r}$ and $\overrightarrow{s'_{f(\ell)+2}} := \vec{s}_{\ell+1}$ yields the desired $(\overrightarrow{s'_i})_{i\leqslant f(\ell+1)}$ satisfying (i) and (ii) for all $k\leqslant\ell+1$. To do so, it remains to show that $\vec{s}_\ell < \vec{r} < \vec{s}_{\ell+1}$ and that both pairs $\vec{s}_\ell < \vec{r}$ and $\vec{r} < \vec{s}_{\ell+1}$ are strongly $\mathcal{P}$-relevant.

Let us begin with $\vec{s}_\ell < \vec{r}$. Recall that both $r, s_\ell\in N$ distinguish P_ℓ and $P_{\ell+1}$ with $\vec{r},\vec{s}_\ell\in P_{\ell+1}$. So by Lemma 9.4.2, we have either $\vec{r} < \vec{s}_\ell$ or $\vec{s}_\ell < \vec{r}$; we show that the former cannot occur. The pre-tangle Q_ℓ in G orients r as it orients s_ℓ and $|r| < |s_\ell|$. So if $\vec{r} < \vec{s}_\ell$, then we have $\vec{r}\in Q_\ell$ since $\vec{s}_\ell\in Q_\ell$ and Q_ℓ is consistent. But $\overleftarrow{r}\in P_\ell$, so r distinguishes P_ℓ and Q_ℓ and $|r| < |s_\ell|$ which contradicts that s_ℓ efficiently distinguishes P_ℓ and Q_ℓ by their choice. Thus, we must have $\vec{s}_\ell < \vec{r}$, as desired.

Next, we claim that $\vec{r} < \vec{s}_{\ell+1}$. Recall that r and s_ℓ are in N and hence nested. By the definition of nested, it suffices to show that neither of $\overleftarrow{r} < \vec{s}_{\ell+1}$, $\vec{s}_{\ell+1} < \vec{r}$ and $\vec{s}_{\ell+1} < \overleftarrow{r}$ holds; we address all these cases separately.

We cannot have $\vec{s}_{\ell+1} < \vec{r}$, as we would otherwise have a contradiction to the consistency of the pre-tangle $P_{\ell+1}$ in G which by choice already contains both $\vec{r}$ and $\overleftarrow{s}_{\ell+1}$. For $\overleftarrow{r} < \vec{s}_{\ell+1}$, observe that r cannot distinguish $P_{\ell+1}$ and $Q_{\ell+1}$ as they are efficiently distinguished by $s_{\ell+1}$ and $|r| < |s_{\ell+1}|$. Thus, $\vec{r} \in P_{\ell+1}$ yields $\vec{r} \in Q_{\ell+1}$. But by the choice of $Q_{\ell+1}$ it also contains $\vec{s}_{\ell+1}$ which together with $\overleftarrow{r} < \vec{s}_{\ell+1}$ then contradicts the consistency of $Q_{\ell+1}$. Finally, if $\vec{s}_{\ell+1} < \overleftarrow{r}$, then this implies $\vec{s}_\ell < \overleftarrow{r}$ since $\vec{s}_\ell < \vec{s}_{\ell+1}$. But we also have $\vec{s}_\ell < \vec{r}$ as shown above, a contradiction to s_ℓ being proper.

It remains to show that both pairs $\vec{s}_\ell < \vec{r}$ and $\vec{r} < \vec{s}_{\ell+1}$ are strongly $\mathcal{P}$-relevant. More precisely, we claim that this is witnessed by the pre-tangles $P_\ell, Q_\ell, P_{\ell+1} \in \mathcal{P}$ and the pre-tangles $Q_\ell, P_{\ell+1}, Q_{\ell+1} \in \mathcal{P}$, respectively. To show this, it suffices to prove that r efficiently distinguishes Q_ℓ and $P_{\ell+1}$ with $\overleftarrow{r} \in Q_\ell$ and $\vec{r} \in P_{\ell+1}$.

Since s_ℓ efficiently distinguishes P_ℓ and Q_ℓ and $|r| < s_\ell$, the pre-tangle Q_ℓ in G orients r in the same way as P_ℓ. Hence, $\overleftarrow{r} \in P_\ell$ yields $\overleftarrow{r} \in Q_\ell$. This implies that r distinguishes Q_ℓ and $P_{\ell+1}$ with $\overleftarrow{r} \in Q_\ell$ and $\vec{r} \in P_{\ell+1}$, and we claim that r does so efficiently. Indeed, the pre-tangles Q_ℓ and P_ℓ in G orient every separation of order less than $|s_\ell|$ in the same way, since s_ℓ efficiently distinguishes them. Hence, every separation of order less than $|s_\ell|$ that distinguishes Q_ℓ and $P_{\ell+1}$ also distinguishes P_ℓ and $P_{\ell+1}$. But r efficiently distinguishes P_ℓ and $P_{\ell+1}$ by the choice of r and $|r| < |s_\ell|$, so there is no separation of order less than r distinguishing Q_ℓ and $P_{\ell+1}$, as desired. □

The fourth Lemma 9.4.6 shows that by transition to suitable subsequences, we can maintain both (IM1) and (IM2) while passing from 'the orders cofinitely exceed every integer' to 'strictly increasing orders'.

Lemma 9.4.6. *Let $(\vec{s}_i)_{i\in\mathbb{N}}$ be a strictly increasing sequence of separations of a graph G, and let $(P_i)_{i\in\mathbb{N}}$ be a sequence of pre-tangles in G such that (IM1) and (IM2) hold. If $(|s_i|)_{i\in\mathbb{N}}$ cofinitely exceeds every integer, then there exists a subsequence $(\vec{s}_{j(i)})_{i\in\mathbb{N}}$ of strictly increasing orders such that it still satisfies (IM1) and (IM2) together with $(P_{j(i)})_{i\in\mathbb{N}}$.*

Proof. We inductively define $j(i)$ for all $i \in \mathbb{N}$. To do so, note that since $(|s_\ell|)_{\ell\in\mathbb{N}}$ cofinitely exceeds every integer, there exist only finitely many s_ℓ of some fixed order. For $i = 0$ let K_0 be the minimal order of any s_ℓ, and set $j(0)$ to be the maximal $\ell \in \mathbb{N}$ with $|s_\ell| = K_0$; this maximum exists since there are only finitely many such ℓ, as noted above. For $i \geqslant 1$ we analogously let K_i be the minimal order of any s_ℓ with $\ell > j(i-1)$ and set $j(i)$ to be the maximal $\ell \in \mathbb{N}$ with $|s_\ell| = K_i$. With this construction we ensure that $|s_\ell| > |s_{j(i)}|$ for all $\ell > j(i)$.

First, we verify (IM1) for $(\vec{s}_{j(i)})_{i\in\mathbb{N}}$ and $(P_{j(i)})_{i\in\mathbb{N}}$. The original sequences satisfy (IM1) which immediately yields $\overleftarrow{s}_{j(i)} \in P_{j(i)}$ and $\vec{s}_{j(i+1)-1} \in P_{j(i+1)}$. By construction, $|s_{j(i)}| \leqslant |s_{j(i+1)-1}|$, so the pre-tangle $P_{j(i+1)}$ in G orients $s_{j(i)}$. Since $\vec{s}_{j(i)} \leqslant \vec{s}_{j(i+1)-1} \in P_{j(i+1)}$, the consistency of $P_{j(i+1)}$ yields $\vec{s}_{j(i)} \in P_{j(i+1)}$. This completes the verification of (IM1).

Secondly, we prove (IM2), and we do so by checking (IM2') which suffices by Lemma 9.4.3. Applying (IM2) to $j(i) < j(i+1)$ in the original sequence yields that every minimal-order separation among $\{s_{j(i)}, \dots, s_{j(i+1)-1}\}$ efficiently distinguishes $P_{j(i)}$ and $P_{j(i+1)}$. But by the definition of $j(i)$, $s_{j(i)}$ is the unique such separation. Hence, (IM2') holds. □

With Lemmas 9.4.3 to 9.4.6 at hand, we are now ready to prove Theorem 9.4.1.

Proof of Theorem 9.4.1. Since $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive, Lemma 9.3.2 and Lemma 9.3.5 together imply that $(|A_i, B_i|)_{i\in\mathbb{N}}$ cofinitely exceeds every integer. In particular, $((A_i, B_i))_{i\in\mathbb{N}}$ contains a subsequence whose orders are even strictly increasing. Note that if a sequence of separations is interlaced with a subsequence of $((A_i, B_i))_{i\in\mathbb{N}}$, then it is also interlaced with $((A_i, B_i))_{i\in\mathbb{N}}$ itself. So by transition to a subsequence, we may assume without loss of generality that $(|A_i, B_i|)_{i\in\mathbb{N}}$ is strictly increasing.

Thus, we can apply Lemma 9.4.5 to $((A_i, B_i))_{i\in\mathbb{N}}$ and obtain a strictly increasing sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ in $\overrightarrow{N}$ which is strongly $\mathcal{P}$-relevant and contains $((A_i, B_i))_{i\in\mathbb{N}}$ as a subsequence; in particular, $((A_i, B_i))_{i\in\mathbb{N}}$ and $((A'_i, B'_i))_{i\in\mathbb{N}}$ are interlaced. Lemma 9.4.4 then yields a corresponding sequence $(P_i)_{i\in\mathbb{N}}$ in $\mathcal{P}$ such that (IM1) and (IM2') hold. Lemma 9.4.3 finally implies that (IM2) holds as well.

The 'moreover'-part can be obtained by applying Lemma 9.4.6 to $((A'_i, B'_i))_{i\in\mathbb{N}}$ and $(P_i)_{i\in\mathbb{N}}$. In order to apply Lemma 9.4.6, it remains to check that $(|A'_i, B'_i|)_{i\in\mathbb{N}}$ cofinitely exceeds every integer. By Lemma 9.3.2, both $((A_i, B_i))_{i\in\mathbb{N}}$ and $((A'_i, B'_i))_{i\in\mathbb{N}}$ have the same limit; in particular, $((A'_i, B'_i))_{i\in\mathbb{N}}$ is non-exhaustive as $((A_i, B_i))_{i\in\mathbb{N}}$ is. Then Lemma 9.3.2 and Lemma 9.3.5 imply that $(|A'_i, B'_i|)_{i\in\mathbb{N}}$ cofinitely exceed every integer, as desired. □

We remark that Theorem 9.4.1 and its proof are purely about sequences of separations and do not refer to the underlying graph in any way. Thus, Theorem 9.4.1 can also be transferred to the more general framework of abstract separation systems which are extensively described in [Die18a]. There, we would then consider a universe of separations equipped with a (not necessarily submodular) order function and pre-tangles in this universe which are efficiently distinguished by a regular nested set of separations.

9.5 Proofs of Theorems 38 and 39

In this section, we prove our main results, Theorem 38 and its structural version Theorem 39, from which we will deduce Theorem 38. To prove Theorem 39, we will apply the interlacing theorem, Theorem 9.4.1, to analyse the limit separation of a non-exhaustive strictly increasing sequence of separations in a tree of tangles, and we then exhibit a thick end in the separator of the limit separation.

Let ω be an end of a graph G. Recall that for every finite set X of vertices of G, there exists a unique component of $G - X$ which contains cofinitely many vertices of some (or equivalently every) ray in ω. Thus, the end ω *induces* a pre-tangle P_ω in G: it orients every finite-order

separation of G to the side which contains cofinitely many vertices of some (or equivalently every) ray in ω [Die17].

Recall that an end ω is *thick* if it contains infinitely many disjoint rays, and if ω is not thick, then it is *thin*. The next lemma shows that the pre-tangle P_ω induced by a thin end ω can be particularly well distinguished from all other pre-tangles in G.

Lemma 9.5.1. *Let ω be a thin end of a graph G, and let P_ω be the pre-tangle in G that is induced by ω. Then there exists $K \in \mathbb{N}$ such that every pre-tangle P in G with $P \not\subseteq P_\omega$ is distinguished from P_ω by a separation of order at most K.*

Proof. Since ω is thin, there exists $K \in \mathbb{N}$ together with an exhaustive strictly increasing sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations of order at most K such that $(A_i, B_i) \in P_\omega$ and $G[B_i]$ is connected for all $i \in \mathbb{N}$ [GH22, Corollary 5.5]. Let P be a pre-tangle in G with $P \not\subseteq P_\omega$. Since P_ω orients all finite-order separations of G, there thus is some finite-order separation $\{C, D\}$ of G distinguishing P and P_ω in that $(C, D) \in P_\omega$ and $(D, C) \in P$. If $\{C, D\}$ has order at most K, then it is as desired; so suppose that $\{C, D\}$ has order greater than K.

Since the sequence $((A_i, B_i))_{i\in\mathbb{N}}$ is exhaustive and all the $G[B_i]$ are connected, Lemma 9.3.10 yields $I \in \mathbb{N}$ such that either $(C, D) < (A_I, B_I)$ or $(D, C) < (A_I, B_I)$. We claim that $\{A_I, B_I\}$ is as desired; since $\{A_I, B_I\}$ has order at most K, it remains to show that $\{A_I, B_I\}$ distinguishes P and P_ω. Since P_ω is consistent and both (C, D) and (A_I, B_I) are contained in P_ω, we must have $(C, D) < (A_I, B_I)$. The separation $\{A_I, B_I\}$ is also oriented by the pre-tangle P in G since it has order at most K and hence order less than $(D, C) \in P$. The consistency of P then implies $(B_I, A_I) \in P$ since $(B_I, A_I) < (D, C) \in P$. Thus, $\{A_I, B_I\}$ distinguishes P and P_ω, as desired. □

A *comb* is a union of a ray R with infinitely many 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 star-comb lemma [Die25, Proposition 8.2.2] implies the following simple observation about combs and infinite sets of vertices in locally finite graphs.

Lemma 9.5.2. *If U is an infinite set of vertices of a locally finite graph G, then G contains a comb with teeth in U as a subgraph.* □

Consequently, we say that an end ω of the locally finite G is in the *closure* of a set $U \subseteq V(G)$ if there exists a comb with teeth in U whose spine is in ω.[2] It is immediate from this definition that if ω is in the closure of U and (A, B) is a separation in the pre-tangle P_ω in G induced by ω, then infinitely many vertices of U are contained in B. Similarly, if $\{A, B\}$ is a finite-order separation of G and the end ω is in the closure of B, then $(A, B) \in P_\omega$.

Recall that if a strictly increasing sequence of tight separations of a connected locally finite graph is non-exhaustive, then its limit separation has infinite order by Theorem 9.3.6. Then Lemma 9.5.2 implies that some end of the graph is in the closure of the separator of such a

[2] An end of a locally finite graph G is in the closure of a set U if and only if it is contained in its topological closure $\bar{U}$ in the Freudenthal compactification $|G|$ of G (cf. [Die25, Section 8.6]).

limit separation. The next lemma shows that this end is unique if the sequence is contained in a nested set distinguishing all the pre-tangles induced by ends in the closure of the separator.

Lemma 9.5.3. *Let N be a nested set of separations of a connected locally finite graph G, let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence of tight separations in $\overrightarrow{N}$, and let (A, B) be its limit. If $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive and N distinguishes all pre-tangles in G that are induced by ends in the closure of $A \cap B$, then precisely one end of G is in the closure of $A \cap B$.*

Proof. Since $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive, $A \cap B$ is infinite by Theorem 9.3.6. So at least one end of G is in the closure of $A \cap B$ by Lemma 9.5.2. Suppose for a contradiction that two distinct ends ω and ω' of G are in the closure of $A \cap B$, and let P_ω and $P_{\omega'}$ be the pre-tangles in G that are induced by ω and ω', respectively. By assumption, there exists a finite-order separation $\{C, D\} \in N$ which distinguishes P_ω and $P_{\omega'}$ with $(D, C) \in P_\omega$ and $(C, D) \in P_{\omega'}$. It suffices to show that $\{C, D\}$ must cross $\{A, B\}$. Indeed, Lemma 9.3.9 then implies that $\{C, D\}$ crosses some $\{A_i, B_i\}$ which contradicts the nestedness of N in which they are both contained.

To show that $\{C, D\}$ crosses $\{A, B\}$, observe that there exist infinitely many elements of $A \cap B$ in C, since ω is in the closure of $A \cap B$. This implies $(C \setminus D) \cap (A \cap B) \neq \emptyset$ as $C \cap D$ is finite. Arguing analogously for $(C, D) \in P_{\omega'}$, we obtain $(D \setminus C) \cap (A \cap B) \neq \emptyset$. But then $\{A, B\}$ and $\{C, D\}$ cross by Lemma 9.2.1. □

We now use the interlacing theorem, Theorem 9.4.1, to show Theorem 9.5.4. It asserts that given a non-exhaustive strictly increasing sequence in a nested set distinguishing all pre-tangles of a locally finite graph, the (by Lemma 9.5.3 unique) end which is in the closure of the separator of its limit separation cannot be thin.

Theorem 9.5.4. *Let $\mathcal{P}$ be a set of pre-tangles in a connected locally finite graph G, and let N be a nested set of $\mathcal{P}$-relevant tight separations which efficiently distinguishes $\mathcal{P}$. Let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence in $\overrightarrow{N}$, and let (A, B) be its limit.*

If $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive and $\mathcal{P}$ contains all the pre-tangles in G which are induced by ends in the closure of $A \cap B$, then there is precisely one end ω in the closure of $A \cap B$ and it is thick.

Proof. The assumptions on N and $((A_i, B_i))_{i\in\mathbb{N}}$ allow us to apply Theorem 9.4.1 to obtain a strictly increasing sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ in $\overrightarrow{N}$ and a corresponding sequence $(P_i)_{i\in\mathbb{N}}$ in $\mathcal{P}$ satisfying (IM1) and (IM2). Furthermore, the sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ is interlaced with $((A_i, B_i))_{i\in\mathbb{N}}$, which implies by Lemma 9.3.2 that (A, B) is also the limit of $((A'_i, B'_i))_{i\in\mathbb{N}}$. In particular, this sequence is non-exhaustive. By the 'moreover'-part of Theorem 9.4.1, we may additionally assume that $(|A'_i, B'_i|)_{i\in\mathbb{N}}$ is strictly increasing and then (A'_i, B'_i) efficiently distinguishes P_i from every P_j with $j > i$.

By Lemma 9.5.3, there is a unique end ω in the closure of $A \cap B$. Let P_ω be the pre-tangle in G which is induced by ω. The following claim shows that every separation in our new sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ efficiently distinguishes P_ω from the respective P_i.

Claim 1. *Every $\{A'_i, B'_i\}$ efficiently distinguishes P_i and P_ω.*

Proof. We first show that $\{A'_i, B'_i\}$ distinguishes P_i and P_ω. Indeed, $(B'_i, A'_i) \in P_i$ by (IM1). To see $(A'_i, B'_i) \in P_\omega$, recall that $(A'_i, B'_i) \leqslant (A, B)$. Thus, $A \cap B$ is contained in B'_i. Since ω is in the closure of $A \cap B$, it is thus in the closure of B'_i as well, and this implies $(A'_i, B'_i) \in P_\omega$ by definition.

Suppose that $\{A'_i, B'_i\}$ does not distinguish P_i and P_ω efficiently. The nested set N efficiently distinguishes $\mathcal{P}$, so N contains a separation $\{C, D\}$ which efficiently distinguishes P_i and P_ω in that $(C, D) \in P_\omega$ and $(D, C) \in P_i$. In particular, we have $|C, D| < |A'_i, B'_i|$. We now claim that $(A'_i, B'_i) < (C, D)$ and that there exists some integer $I > i$ with $(C, D) < (A'_I, B'_I)$.

First, let us show that $(A'_i, B'_i) < (C, D)$. The nested separations $\{C, D\}$ and $\{A'_i, B'_i\}$ both distinguish P_i and P_ω with $(C, D), (A'_i, B'_i) \in P_\omega$ as above. Thus, Lemma 9.4.2 yields that either $(A'_i, B'_i) < (C, D)$ or $(C, D) < (A'_i, B'_i)$. Suppose for a contradiction that $(C, D) < (A'_i, B'_i)$. Recall that $(A'_i, B'_i) \in P_{i+1}$ by (IM1). The consistency of the pre-tangle P_{i+1} in G thus yields $(C, D) \in P_{i+1}$, since $|C, D| < |A'_i, B'_i|$. We also have $(D, C) \in P_i$ by the choice of $\{C, D\}$, so $\{C, D\}$ distinguishes P_i and P_{i+1}. But $|C, D| < |A'_i, B'_i|$ which contradicts (IM2) for $i < i + 1$. Thus, we must have $(A'_i, B'_i) < (C, D)$, as desired.

Secondly, we find an integer $I > i$ with $(C, D) < (A'_I, B'_I)$. By Lemma 9.3.11, it suffices to show that $(C, D) \leqslant (A, B)$. Since $\{C, D\}$ and all the $\{A'_j, B'_j\}$ are contained in the nested set N, Lemma 9.3.9 implies that $\{C, D\}$ has to be nested with $\{A, B\}$ as well. By the choice of $\{C, D\}$, we have $(C, D) \in P_\omega$, so D contains infinitely many elements of $A \cap B$, and since $\{C, D\}$ has finite order, we have $(D \smallsetminus C) \cap (A \cap B) \neq \emptyset$. Hence, either $(C, D) \leqslant (A, B)$ or $(C, D) \leqslant (B, A)$ holds by Lemma 9.2.1 because $\{C, D\}$ and $\{A, B\}$ are nested. Suppose for a contradiction that $(C, D) \leqslant (B, A)$. Then $(A'_i, B'_i) < (C, D) \leqslant (B, A)$, but we also have $(A'_i, B'_i) \leqslant (A, B)$ which contradicts the fact that $\{A'_i, B'_i\}$ is tight and hence proper.

We complete the proof of the claim by obtaining a contradiction from $(A'_i, B'_i) < (C, D) < (A'_I, B'_I)$. By (IM1), we have that $(B'_i, A'_i) \in P_i$ and $(A'_I, B'_I) \in P_{I+1}$. Since $|C, D| < |A'_i, B'_i| < |A'_I, B'_I|$, both pre-tangles P_i and P_{I+1} in G orient $\{C, D\}$. The consistency of the pre-tangles P_i and P_{I+1} together with $(A'_i, B'_i) < (C, D) < (A'_I, B'_I)$ then yields $(D, C) \in P_i$ and $(C, D) \in P_{I+1}$; in particular, $\{C, D\}$ distinguishes P_i and P_{I+1}. But $|C, D| < |A'_i, B'_i|$, which contradicts that $\{A'_i, B'_i\}$ efficiently distinguishes P_i and P_{I+1} by the choice of the sequence $((A'_i, B'_i))_{i \in \mathbb{N}}$. ■

Now suppose for a contradiction that ω is thin. Then Lemma 9.5.1 yields some $K \in \mathbb{N}$ such that P_ω is distinguished from every other pre-tangle $P \not\subseteq P_\omega$ in G by a separation of order at most K. But the orders of the $\{A'_i, B'_i\}$ are strictly increasing, so there exists $j \in \mathbb{N}$ such that (A'_j, B'_j) has order greater than K. By the claim, this $\{A'_j, B'_j\}$ efficiently distinguishes P_j and P_ω, which then contradicts the choice of K and hence the thinness of ω. So ω is a thick end of G, as desired. □

Instead of finishing the proof of Theorem 9.5.4 by a contradiction in terms of cardinality via Lemma 9.5.1, one can also perform a structural analysis to obtain that the thickness of ω is even witnessed by an infinite family of disjoint rays in $G[B \smallsetminus A]$ (see Section 9.6).

It remains to deduce our main results, Theorems 38 and 39, by applying Theorem 9.5.4 to a tree of tangles of the graph G. To do so, recall that a *tangle P in G* is a pre-tangle in G such that for every three – not necessarily distinct – separations $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in P$, we have $G[A_1] \cup G[A_2] \cup G[A_3] \neq G$. We remark that every pre-tangle induced by some end is in fact a tangle. Recall further that a *tree of tangles* of G is a nested set of separations which efficiently distinguishes the set $\mathcal{P}$ of all the tangles in G and which contains only $\mathcal{P}$-relevant separations.

The next lemma asserts that every separation in a tree of tangles, and even every $\mathcal{P}$-relevant separation, is tight.

Lemma 9.5.5 ([EKT22], Lemma 6.1 and its proof)**.** *Let $\{A, B\}$ be a separation of a locally finite graph G which efficiently distinguishes a pre-tangle P and a tangle Q in G. Suppose that $(A, B) \in Q$. Then there exists a tight component K of $G - (A \cap B)$ with $V(K) \subseteq B \setminus A$. In particular, if P is also a tangle, then $\{A, B\}$ is tight.*

Proof of Theorem 39. Each separation in the tree of tangles N efficiently distinguishes some two tangles in G and is hence tight by Lemma 9.5.5. Since every tangle in G is also a pre-tangle in G, Theorem 9.5.4 implies that the unique end in the closure of $A \cap B$ is thick, as desired. □

Proof of Theorem 38. By Theorem 9.1.2, there exists a tree of tangles N of G. Since G does not have any thick end, Theorem 39 implies that every strictly increasing sequence in $\vec{N}$ is exhaustive, so N is itself exhaustive. Therefore, N induces a canonical tree-decomposition of G by Lemma 9.3.4, as all the separations in N are tight and hence proper. This tree-decomposition of G then efficiently distinguishes all distinguishable tangles in G by construction and hence is as desired. □

If one considers, instead of tangles, the more general 'profiles' as introduced in [DHL19], then there is a direct generalisation of Theorem 9.1.2: Elbracht, Kneip and Teegen showed in [EKT22, Theorem 6.6] that for a robust set $\mathcal{P}$ of regular profiles in a connected locally finite graph, there is a canonical nested of $\mathcal{P}$-relevant separations which efficiently distinguishes $\mathcal{P}$. Since profiles are also pre-tangles and every tangle, in particular the ones induced by ends, are profiles, one can then obtain direct analogues of Theorem 38 and Theorem 39 for sets $\mathcal{P}$ of profiles with the same proofs.

9.6 Infinitely many rays beyond the limit

In this section we prove a structurally refined version of Theorem 9.5.4: while Theorem 9.5.4 asserts that the (unique) end ω of G in the closure of the separator of the limit separation (A, B) is thick, we here show in Proposition 9.6.1 that there are infinitely many disjoint rays in $G[B \setminus A]$ which are contained in the end ω; in other words, these rays witness the thickness of ω 'beyond' the limit separation (A, B).

Proposition 9.6.1. *Let $\mathcal{P}$ be a set of pre-tangles in a connected locally finite graph G, and let N be a nested set of $\mathcal{P}$-relevant tight separations of G which efficiently distinguishes $\mathcal{P}$. Let $((A_i, B_i))_{i\in\mathbb{N}}$ be a strictly increasing sequence in $\vec{N}$, and let (A, B) be its limit.*

If $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive, then the unique end ω of G which is in the closure of $A \cap B$ contains an infinite family of disjoint rays in $G[B \setminus A]$.

As preparation for the proof, we remark that following the proof of the ubiquity of rays [Hal65, Satz 1] by Halin one easily obtains the following (see also [Die25, Theorem 8.2.5 & Exercise 43]).

Theorem 9.6.2. *An end ω in a graph G contains infinitely many disjoint rays, if it contains k many disjoint rays for every $k \in \mathbb{N}$.* □

Proof of Proposition 9.6.1. We first proceed as in the proof of Theorem 9.5.4 to obtain a strictly increasing sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ in $\vec{N}$ with the same limit (A, B) and a corresponding sequence $(P_i)_{i\in\mathbb{N}}$ in $\mathcal{P}$ satisfying (IM1) and (IM2). Moreover, the (A'_i, B'_i) have strictly increasing order and, by the claim, every (A'_i, B'_i) efficiently distinguishes P_i and P_ω, the pre-tangle in G induced by ω.

We claim that it suffices that there are k disjoint rays in $G[B \setminus A]$ which belong to ω in G for every $k \in \mathbb{N}$. Indeed, let Y be the set of all ends of $G[B \setminus A]$ which contain a ray in ω. If Y is infinite, then one greedily finds the desired infinitely many disjoint rays in $G[B \setminus A]$ in ω. Otherwise Y is finite. Thus, some end $\omega' \in Y$ contains k disjoint rays for every Y by the pigeonhole principle. Hence by Theorem 9.6.2, ω' contains infinitely many disjoint rays. These are also contained in ω by the definition of Y.

So consider some fixed $k \in \mathbb{N}$. Let $X := A \cap B$, and fix any set $Z \subseteq X$ of size k. By Lemma 9.3.3, there exists $I \in \mathbb{N}$ such that $Z \subseteq X_i := A'_i \cap B'_i$ for all $i \geqslant I$. Since every $\{A'_i, B'_i\}$ efficiently distinguishes P_i from P_ω, no X_i–ω separator of G of size less than $|X_i|$ exists. Thus, it follows from [BDS05, Theorem 1.1] that there is a family $\mathscr{R}_i$ of $|X_i|$ many disjoint X_i–ω rays in G for every $i \in \mathbb{N}$. We remark that every vertex in X_i is the starting vertex of some ray in $\mathscr{R}_i$. For every $i \geqslant I$, we fix the subfamily $\mathscr{R}'_i$ of $\mathscr{R}_i$ containing precisely those rays starting in Z.

For every $\ell \in \mathbb{N}$, let $\mathcal{M}^\ell$ be the set consisting of those families of k paths of length ℓ which appear jointly as initial segments of the rays in $\mathscr{R}'_i$ for infinitely many $i \geqslant I$; note that the paths in such families are disjoint and start in Z. Since the graph G is locally finite, there are only finitely many paths of length ℓ in G which start in Z. To see that $\mathcal{M}^\ell$ is non-empty, we can thus apply the pigeonhole principle to the infinite collection of families of k paths of length ℓ starting in Z which are given by the initial segments of the rays in the $\mathscr{R}'_i$.

Given any family $\{R_1, \dots, R_k\} \in \mathcal{M}^\ell$, we have $\{R'_1, \dots, R'_k\} \in \mathcal{M}^{\ell-1}$ where R'_i is the subpath of R_i of length $\ell - 1$ which still starts in Z. Thus, König's infinity lemma [Kön36, VI. Kapitel § 2, Satz 6] (see also [Die25, Lemma 8.1.2]) applied to $(\mathcal{M}^\ell)_{\ell\in\mathbb{N}}$ yields for every $\ell \in \mathbb{N}$ a family $\{R^\ell_1, \dots, R^\ell_k\} \in \mathcal{M}^\ell$ such that $R^{\ell-1}_i$ is the subpath of R^ℓ_i of length $\ell - 1$ starting in Z for $i = 1, \dots, k$. Let R_i be the ray $\bigcup_{\ell\in\mathbb{N}} R^\ell_i$ for $i = 1, \dots, k$. These rays R_i are disjoint as the R^ℓ_i

are disjoint for every $\ell \in \mathbb{N}$ by the definition of $\mathcal{M}^\ell$. We claim that $R_1, \dots, R_k$ are as desired, i.e. a family of k disjoint rays in ω which are contained in $G[B \setminus A]$ (except their startvertices).

First, $R_1, \dots, R_k$ are contained in $G[B \setminus A]$ except their startvertices which are in $Z \subseteq A \cap B$. Indeed, let v be a vertex of some R_i which is not its startvertex. By the definition of R_i and $(\mathcal{M}^\ell)_{\ell \in \mathbb{N}}$, v is contained in a ray in $\mathscr{R}_i$ for infinitely many $i \geqslant I$. All the rays in $\mathscr{R}_i$ avoid X_i except in their startvertices. Hence, every ray in $\mathscr{R}_i$ is contained in $G[B'_i \setminus A'_i]$ except its startvertex, since ω is in the closure of $A \cap B \subseteq B'_i$. Altogether, v is contained in $B_i \setminus A_i$ for infinitely many $i \geqslant I$ and thus in $B \setminus A$ by Lemma 9.3.3.

It remains to show that $R_1, \dots, R_k$ belong to the end ω of G. Consider any $R := R_i$, and let Q be an arbitrary ray in ω. We now inductively exhibit infinitely many disjoint paths in G joining R and Q which then yields that they are in the same end of G and, in particular, that $R \in \omega$. Every vertex in $R \cap Q$ is a trivial path joining R and Q. So if $R \cap Q$ is infinite, then we are done. Suppose that $R \cap Q$ is finite.

Let us assume that we have already found disjoint paths $P_1, \dots, P_j$ joining R and Q for some $j \in \mathbb{N}$. We seek an additional path P_{j+1} joining R and Q which is disjoint from all the previous P_i. Since $P_1, \dots, P_j$ are finite paths, there are finite initial segments R^* and Q^* of R and Q, respectively, such that $R - R^*$ and $Q - Q^*$ are both disjoint from $P_1, \dots, P_j$. By the construction of the R_i, the initial segment of R which consists of R^* together with the following edge is also the initial segment of some ray $Q_{j+1} \in \omega$. Since both Q and Q_{j+1} are in ω, there are infinitely many disjoint Q–Q_{j+1} paths, and we especially find one such path P'_{j+1} which is disjoint from both R^* and Q^* as well as the paths $P_1, \dots, P_j$. We then let P_{j+1} be the (unique) path which first follows $Q_{j+1} - R^*$ and then P'_{j+1}. This path P_{j+1} starts in R and ends in Q by the choice of Q_{j+1} and P'_{j+1}, and it is disjoint from $P_1, \dots, P_j$ by construction, as desired. Iterating this recursive construction yields the desired set of infinitely many disjoint paths starting in R and ending in Q. □

A natural strengthening of Proposition 9.6.1 would assert that the infinitely many rays in $G[B \setminus A]$ not only belong to the same end ω of G, but even to the same end of $G[B \setminus A]$. This, however, does not hold as the following example demonstrates; this example is an adaption of [EKT22, Example 4.9].

Example 9.6.3. We construct a graph G (see Figure 9.2) as follows.

For every $n \in \mathbb{N}$, we pick a copy K_n of the complete graph $K^{2^{n+4}}$ on 2^{n+4} vertices. In this $K^{2^{n+4}}$, we fix 2^n vertices $u^n_1, \dots, u^n_{2^n}$ and 2^{n+1} vertices $v^n_1, \dots, v^n_{2^{n+1}}$ all of which are distinct. We then identify u^{n+1}_i with v^n_i for all $n \in \mathbb{N}$ and $i \leqslant 2^{n+1}$. Furthermore, we pick a ray $R^n = r^n_0 r^n_1 r^n_2 \dots$ for every $n \in \mathbb{N}$ such that these rays are disjoint from each other and from the graph constructed before. Then for every $m \in \mathbb{N}$ and $n \leqslant m$, we add an edge joining r^n_m and $w^m := u^{m+1}_1 = v^m_1$. This completes the construction of the graph G.

Each of the K_n induces a 2^{n+2}-tangle P_n in G. It is easy to check that the tight separation s_n with separator $S_n := \{v^i_1 \mid i < n\} \cup \{v^n_j \mid j \leqslant 2^{n+1}\}$ efficiently distinguishes P_n and P_{n+1}. Moreover, the nested set $N = \{s_n \mid n \in \mathbb{N}\}$ efficiently distinguishes the set $\mathcal{P} = \{P_n \mid n \in \mathbb{N}\}$ of (pre-)tangles in G.

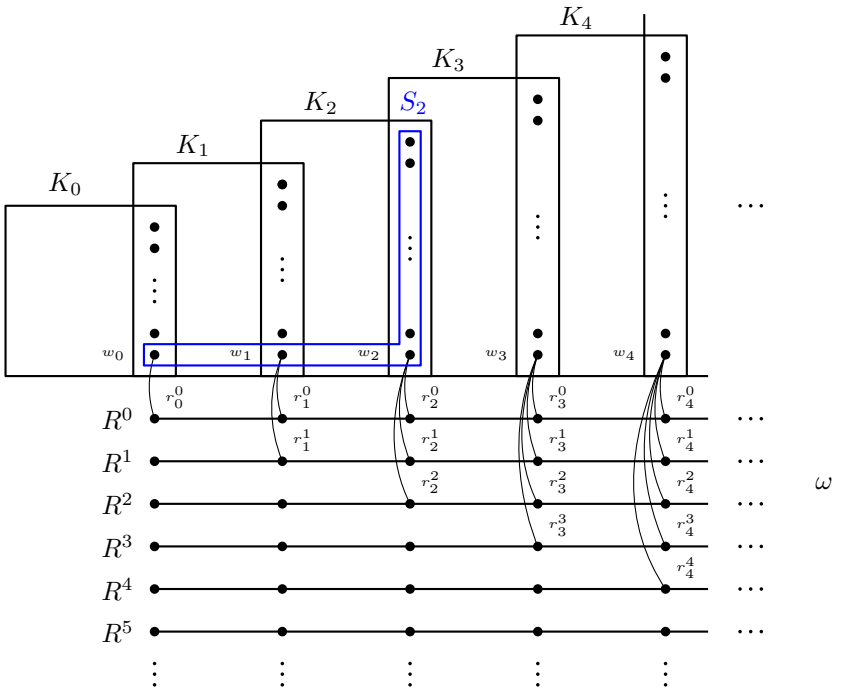

Figure 9.2: The graph G from Example 9.6.3

Now let $\vec{s}_n$ be the orientation of s_n with $\vec{s}_n \in P_{n+1}$. Then $(\vec{s}_n)_{n\in\mathbb{N}}$ is a strictly increasing sequence in $\vec{N}$, and we write $\vec{s} = (A, B)$ for its limit separation. We have $B \setminus A = \bigcup_{n\in\mathbb{N}} V(R^n)$. Hence, the sequence $(\vec{s}_n)_{n\in\mathbb{N}}$ is non-exhaustive, and $G[B \setminus A]$ is the disjoint union of the rays R^n. In particular, $G[B \setminus A]$ does not contain an infinite family of equivalent rays.

The above Example 9.6.3 can be adapted to show that there need not even exist infinitely many rays in $G[B \setminus A]$ which are equivalent in $G[B]$. Instead of joining the rays R_n as described above, we join each R_n to the set $A \cap B = \{v_1^0, v_1^1, v_1^2, \dots\}$ cofinally in such a way that each v_1^i is joined to precisely i of the rays and every two rays R^n and R^m share only finitely many neighbours in $A \cap B$. Now s_n chosen as above is again a tight separation which efficiently distinguishes the tangles P_n and P_{n+1} induced by the respective copies of $K^{2^{n+4}}$ and $K^{2^{(n+1)+4}}$. Now every two disjoint rays in $G[B \setminus A]$ are not equivalent, as desired. We omit the detailed construction here.

CHAPTER 10

On vertex sets inducing tangles

This chapter is based on [ABJKW24], which is joint work with Sandra Albrechtsen, Hanno von Bergen and Raphael W. Jacobs.

10.1 Introduction

10.1.1 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 [RS]. 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 [RS91]: 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 [DHL19, 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 10.1.1. [DHL19] *Is every tangle in a graph G induced by some set $X \subseteq V(G)$?*

In what follows we will often consider Problem 10.1.1 for all tangles of some fixed order $k \in \mathbb{N}$ and then say for short: Problem 10.1.1 for k.

We remark that Problem 10.1.1 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. [RS91, (2.6)]). For $k = 3$, Elbracht [Elb15, Theorem 1.2] proved Problem 10.1.1 directly. Independently, Grohe [Gro16b, 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 10.1.1 for $k = 3$. Similarly, it should be possible to derive Problem 10.1.1 for $k = 4$ from the recent characterisation of 4-tangles in terms of 'internal 4-connectedness' by Carmesin and Kurkofka [CK23b, Theorem 1]. Besides these results for small k, Diestel, Elbracht and Jacobs [DEJ24, Theorem 12] showed that Problem 10.1.1 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 [EKT20, Theorem 2] made substantial progress towards Problem 10.1.1 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 10.1.2. [EKT20] *Every tangle in a graph G is induced by some weight function on $V(G)$.*

Contrary to their positive result, Elbracht, Kneip and Teegen [EKT20, Theorem 10] explicitly construct an example which shows that not only Problem 10.1.1, but also Theorem 10.1.2 fails for tangles in general discrete contexts, such as matroids or data sets (see e.g [Die18a; DHL19; DO19; EKT21]). However, no such example is known for tangles in graphs. Thus, Problem 10.1.1 is open for all $k \geqslant 4$.

10.1.2 Our contributions to Problem 10.1.1

In this chapter, we reduce Problem 10.1.1 for every k to graphs whose size is bounded by a function in k:

Theorem 40. *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 40, we obtain that one may verify the validity of Problem 10.1.1 for any fixed k computationally in explicitly bounded, though impractically long, time by checking Problem 10.1.1 for every k-tangle in every connected graph with fewer than M edges.

Corollary 41. *For $k \geqslant 1$, there exists $M = M(k) \in O(3^{k^{k^5}})$ such that Problem 10.1.1 holds for k if it holds for all k-tangles in connected graphs G with fewer than M edges.*

[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 $X \setminus X'$ is 0.

Our second corollary of Theorem 40 asserts that whenever Problem 10.1.1 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 10.1.1 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 that induces τ, which exists by Theorem 10.1.2, in such way that its *total weight* $w(V(G)) = \sum_{v \in V(G)} f(v)$, which it distributes on $V(G)$, is bounded in k.

Corollary 42. *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 10.1.1 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 40.

We derive Corollary 42 from Theorem 40 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 10.1.2) 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 40 by fixing each such weight function to be an indicator function of some inducing set given by the assumed positive answer to Problem 10.1.1 (see Section 10.8 for details).

10.1.3 An inductive proof method for tangles

Our proof of Theorem 40 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 independent interest.

Theorem 43. *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 $G_0, \dots, G_m$ of graphs and k-tangles τ_i in G_i for every $i \in \{0, \dots, m\}$ 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 43. 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 10.3.6).

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 $A \smallsetminus B$ and $B \smallsetminus A$. 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 10.3.2 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. [RS91, (6.1)] and [CK23b, Lemma 2.1]), we remark that 'surviving' is the reverse notion.

10.1.4 Overview of the proof of Theorem 43

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 10.3.1), and

(ii) if $k = 2$, then τ extends to a k-tangle in $G - e$ for some edge e of G (Lemma 10.3.2).

A rather simple analysis will yield that

(iii) if G is disconnected, then τ induces to a k-tangle in a unique component of G (Lemma 10.3.3),

(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 10.3.4), 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 10.3.5).

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 10.3). 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 44. *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 44 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.

10.1.5 Proof sketch of Theorem 44

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 10.4). 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 10.4.3). 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 10.4.4).

Second, we assume that the graph G contains no tangle of order $> k$ (Section 10.7). 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 10.6). To obtain the desired decomposition, we start with the tree-decomposition obtained from the tangle-tree duality theorem [Die25, 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 10.5.1).

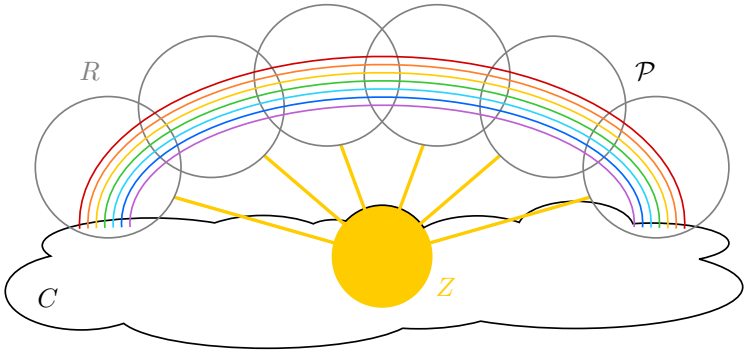

Figure 10.1: Schematic drawing of a rainbow-cloud decomposition

Roughly speaking (see also Figure 10.1), 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 10.5 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 10.7.1).

10.1.6 How this chapter is organised

We recall the relevant terminology concerning tangles in Section 10.2. We then introduce in Section 10.3 the notion of 'survives' and prove Theorem 43 in several special cases: the case $k \leqslant 2$, the case $\delta(G) \leqslant 2$, and when the graph G is disconnected. In Section 10.4 we then prove Theorem 44, the remaining case of Theorem 43 when G contains a tangle of order $k+1$. In Section 10.5, 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 10.6 about the interaction of rainbow-cloud decompositions and separations, we complete the proof of Theorem 44 in Section 10.7. In Section 10.8 we collect all the individual cases to prove Theorem 43, derive Theorem 40, and deduce Corollary 41 and Corollary 42.

10.2 Preliminaries

In this section, we recall the relevant terms and definitions concerning tangles in graphs. For general graph-theoretic concepts and notation, we follow [Die25]. In particular, we denote the set $\{1, \dots, n\}$ with $n \in \mathbb{N}$ by $[n]$. A path is *trivial* if it has length 0. We remark that all graphs in this chapter are finite.

10.2.1 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 sets A and B are the *sides* of the separation $\{A, B\}$, and $A \smallsetminus B$ and $B \smallsetminus A$ are its *strict sides*. The separation $\{A, B\}$ *separates* every two sets $X \subseteq A$ and $Y \subseteq B$. 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. We refer to A as the *small* side of an oriented separation (A, B) and to B as its *big* side. 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 $\vec{S}_k(V)$. We will often also refer to 'oriented separations' simply as 'separations' if the meaning is clear from the context. Similarly, we will transfer definitions for separations verbatim to oriented separations and vice-versa if the transfer is unambiguous.

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 [Bir48, 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.

If two unoriented separations of a set V have orientations which are comparable, then they are called *nested*; if they are not nested, then the two separations *cross*.

By double counting, we obtain that the order of the supremum and infimum of $\{A,B\}$ and $\{C,D\}$ sum up to

$$|(A,B) \wedge (C,D)| + |(A,B) \vee (C,D)| = |A,B| + |C,D|.$$

The inequality '$\leqslant$' implied by the equality above implies that $|\cdot|$ is a *submodular* function on the lattice $(\vec{U}(V), \leqslant)$. In particular, this implies that for two separations $\{A,B\}, \{C,D\}$ of order $< k$, at least one of $\{A \cap C, B \cup D\}$ and $\{A \cup C, B \cap D\}$ has also order $< k$.

10.2.2 Tangles and separations of graphs

Let $G = (V,E)$ be a graph. A *separation* $\{A,B\}$ *of* G is a separation of its vertex set V such that no edge of G joins $A \setminus B$ and $B \setminus A$ (equivalently: $G = G[A] \cup G[B]$). We write $U(G)$ for the set of all unoriented separations of G and $S_k(G) := \{\{A,B\} \in U(G) : |A \cap B| < k\}$; as above, we then define $\vec{U}(G)$ and $\vec{S}_k(G)$ as the respective sets of all oriented separations. Note that $U(G) \subseteq U(V)$ by definition. Moreover, $(\vec{U}(G), \leqslant)$ is a sublattice of $(\vec{U}(V), \leqslant)$, as it is straight forward to check that given two separations (A,B), (C,D) of G, the ordered pairs $(A \cap C, B \cup D)$ and $(A \cup C, B \cap D)$ are again separations of G.

Now let $k \geqslant 1$ be an integer. An *orientation* of $S_k(G)$ is a set $O \subseteq \vec{S}_k(G)$ that contains precisely one orientation of every separation $\{A,B\} \in S_k(G)$. If an orientation O of $S_k(G)$ contains (A,B), we say that O *orients* $\{A,B\}$ *as* (A,B). A *k-tangle* or *tangle of order k* in G is an orientation τ of $S_k(G)$ which contains no element of

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

as a subset. We say the k-tangle *avoids* $\mathcal{T}$. The elements of $\mathcal{T}$ are called *forbidden triples*. A forbidden triple of $\mathcal{T}$ such that two of the (A_i, B_i) are equal is a *forbidden tuple*. Tangles of unspecified order are referred to as *tangles in* G. A k-tangle τ in G *extends* to a k^*-tangle τ^* in G for $k^* \geqslant k$ if $\tau \subseteq \tau^*$. We say that a separation $\{A,B\}$ of G *distinguishes* two tangles τ, τ' in G if τ and τ' contain distinct orientations of $\{A,B\}$.

It is well-known [RS91, (2.6)] that tangles of order 1 and 2 are in a one-to-one correspondence to the components and blocks[2], respectively:

[2]Recall that a *block* of a graph is a $\subseteq$-maximal connected subgraph H such that $H - v$ is connected for every

Proposition 10.2.1. *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.* □

From its definition, tangles have several properties that we will frequently make use of. First, every tangle τ is *consistent*: there are no two separations $(A,B),(C,D) \in \tau$ such that $(D,C) \leqslant (A,B)$. Secondly, if τ has order k, then $(X, V(G)) \in \tau$ for all $X \subseteq V(G)$ of size $< k$; this property is sometimes called the *regularity* of τ. Thirdly, τ has the *profile property*: if $(A,B),(C,D) \in \tau$ and their supremum $(A \cup C, B \cap D)$ again has order $< k$, then it is also in the tangle τ.

Finally, let us make two observations about the interplay of tangles with the partial order $\leqslant$ on oriented separations. First, if an orientation τ of $S_k(G)$ is not a tangle and thus contains a forbidden triple, then τ also contains a forbidden triple whose elements are $\leqslant$-maximal in τ. Secondly, if τ is a tangle in G, then every $\leqslant$-maximal separation $(A,B) \in \tau$ satisfies that $G[B \setminus A]$ is connected.

10.2.3 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.

10.2.4 Lifts of tangles

It is well-known that tangles in minors lift to tangles in the host graph (see [RS91, (6.1)] or [CK23b, 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 $(A \setminus \{v\}, B \setminus \{v\})$ and by C'_i the set $C' \cup \{u_i\}$ for $C' = A', B'$ and $i = 1, 2$.

---

vertex $v \in H$. In particular, as the empty graph is not considered to be connected, every block contains some edge.

(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 $\{i, j\} = \{1, 2\}$, $u_i \in A \setminus B$ and $u_j \in B \setminus A$, 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 is hence an orientation of $S_k(G)$. Indeed, we otherwise have $(A', B'_i), (B', A'_j) \in \tau$, 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 we have $(G'[A'] \cup G'[B']) + u_1u_2 = G'$, and thus (A', B'_i), (B', A'_j) and $(\{u_1, u_2\}, V(G'))$ 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.

10.3 Definition of 'survive' and special cases of Theorem 43

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 10.3.1 and then for graphs G' obtained from G by suppressing a single vertex of degree 2 in Section 10.3.2. Alongside these definitions, we prove several lemmas which all deal with special cases of Theorem 43. Finally, we provide Example 10.3.6 which demonstrates that suppressing vertices of degree 2 needs to be allowed in Theorem 43.

10.3.1 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.

Lemma 10.3.1. *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 10.2.1, the set $X_\tau := \bigcap_{(A,B)\in\tau} B$ is a component of G corresponding to the 1-tangle τ. Let X' be a component of G' whose vertex set is contained in X_τ. By Proposition 10.2.1 the component X' of G' corresponds to a 1-tangle τ' in G' with $X_{\tau'} = \bigcap_{(A,B)\in\tau'} B = V(X')$. 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 10.3.1 the deletion of e may give rise to (at most) two components X' in $G[X_\tau]$. Thus, τ does not necessarily extend to a unique 1-tangle in G'.

Lemma 10.3.2. *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 10.2.1, $G[X_\tau]$ 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 10.2.1, 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 10.3.1 also in the proof of Lemma 10.3.2 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'.

Lemma 10.3.3. *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 10.2.1, $G[X]$ with $X := X_{\tau_1} = \bigcap_{(A,B)\in\tau_1} B$ 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 $A := A' \cup (V(G) \setminus X)$ 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 $((V(G) \setminus X), X)$ of order 0 is in $\tau_1 \subseteq \tau$. Thus, the forbidden triple

$$\{(A, B), (C, D)), ((V(G) \setminus X), X)\} \text{ or } \{(B, A), (C, D)), ((V(G) \setminus X), X)\})$$

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, $(A_i, B_i) := (A'_i \cup (V(G) \setminus X), B'_i)$ 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. □

Lemma 10.3.4. *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. □

10.3.2 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 10.3.5 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 10.3.5. *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. □

10.3.3 Suppressing vertices of degree 2 in Theorem 43

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 43 it is necessary to allow the suppression of a vertex.

Example 10.3.6. 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 $H_0 := H \supseteq H_1 \supseteq \ldots \supseteq H_n$ 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 10.3.3, 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 tangle would induce a k-tangle in $G' - e'$ by Lemmas 10.3.4 and 10.3.5 where $e' := e$ if $e \in E(G)$ or $e' := uv$ otherwise. □

10.4 Proof of Theorem 44 if G has a higher-order tangle

In this section we prove Theorem 44 for graphs which contain a tangle of order $> k$.

Theorem 10.4.1. *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 show that if τ itself extends to a $(k+1)$-tangle in G, then τ extends to a k-tangle in $G - e$ for every edge $e \in E(G)$ (Lemma 10.4.3). 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 10.4.4).

For both cases, we will make use of the following observation.

Lemma 10.4.2. *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 $< k$, and so τ orients both of them. The consistency of 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 10.4.3. *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 10.4.2. 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 10.4.3 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 10.4.2. 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 10.4.4. *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 $C \smallsetminus D$ and the other one in $D \smallsetminus C$. 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 10.4.2, 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 $\tilde{\tau} \cap \vec{S}_k$, we then have $(A \cap C, B \cup D) \in \tilde{\tau}$. Since $\{(B, A), (A \cap C, B \cup D), (D, C)\}$ is 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 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, $e \in G[A \smallsetminus B] \subseteq G[C_1 \smallsetminus D_1]$. 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 $G[A \smallsetminus B]$. 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 $e \in G[A \smallsetminus B]$, $\{B \cap C, A \cup D\}$ 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 $G[A \smallsetminus B]$ by its choice. Additionally, it has one endvertex in $C \smallsetminus D$ and the other one in $D \smallsetminus C$ 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 10.4.1. This follows immediately from Lemma 10.4.3 and Lemma 10.4.4. □

Our proof of Theorem 10.4.1 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 10.4.5. *Let τ be a tangle in G of order $k \geqslant 3$. Suppose further that there exists another k-tangle τ^* in G with $\tau^* \not\subseteq \tau$ 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 $G - e$?*

10.5 Rainbow-cloud-decompositions in the absence of high-order tangles

Recall that Theorem 10.4.1 in Section 10.4 immediately yields Theorem 44 if the graph has a $(k+1)$-tangle. So from now on we work towards a proof of Theorem 44 for graphs without tangles of high order. 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 [KNTW12]. 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) $\bigcup_{i=0}^{M} G[W_i] = G$,

[3]Decompositions satisfying (L1) and (L2) are often known as *path-decompositions* (cf. [Die25, §12.6]). In [KNTW12] 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).

(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:

(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 10.2 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.

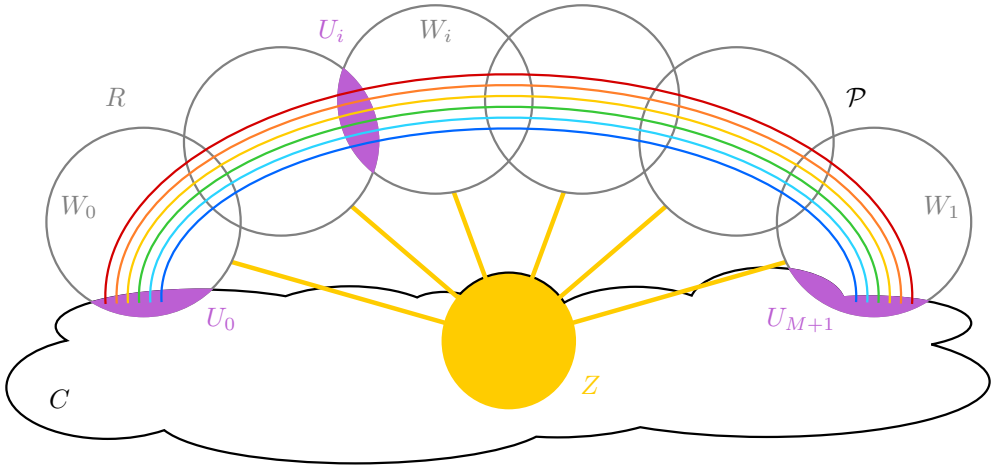

Figure 10.2: A rainbow-cloud decomposition $(R, \mathcal{W}, Z, C)$: the rainbow R together with its decomposition $\mathcal{W}$ is indicated in grey, the foundational linkage $\mathcal{P}$ of $\mathcal{W}$ is depicted in rainbow colours, and the cloud C is depicted in black. Further, indicated in yellow, 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.

With the definition of rainbow-cloud-decompositions at hand, we can state the main result of this section:

Theorem 10.5.1. *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 10.5.1, 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 10.5.3 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 for every two successive separations, there exists a linkage between its separators. By Menger's theorem (e.g. [Die25, 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 10.5.2. *Let $n, m, p \geqslant 1$ be integers with $p \geqslant n^m$. Then every sequence $(a_1, \dots, a_p)$ of integers of length p with $a_i \in \{0, \dots, m-1\}$ for every $i \in [p]$ has a subsequence $(a_{i_1}, \dots, a_{i_n})$ 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 $i_1, \dots, i_{n'}$ be the respective indices. Now consider the $n' + 1$ subsequences

$$(a_1, \dots, a_{i_1 - 1}), (a_{i_1 + 1}, \dots, a_{i_2 - 1}), \dots, (a_{i_{n'} + 1}, \dots, a_p) \tag{10.1}$$

of consecutive $a_1, \dots, a_p$; 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 (10.1), 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. □

Given a graph G, a sequence $((A_i, B_i))_{i \in [p]}$ of (oriented) separations is *strictly increasing* if $(A_i, B_i) < (A_j, B_j)$ for every two elements $i < j$ of $[p]$.

Lemma 10.5.3. *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) *$\ell := |A_1, B_1| = \cdots = |A_n, B_n|$, 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 10.5.2 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 10.5.2. 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

$$(A_1 \cap A, B_1 \cup B), \dots, (A_q \cap A, B_q \cup B), (A, B), (A_1 \cup A, B_1 \cap B), \dots, (A_r \cup A, B_r \cap B).$$

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). □

To proceed to the proof of Theorem 10.5.1, we recall the tangle-tree duality theorem and all necessary definition: For a tree T and two nodes or edges x, y of T, we denote by xTy the (unique) $\subseteq$-minimal path in T which contains x and y. A *tree-decomposition* of a graph G is a pair $(T, \mathcal{V})$ of a tree T and a family $\mathcal{V}$ of subsets V_t of $V(G)$ indexed by the nodes of T such that

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

(T2) for every three nodes $r, s, t \in T$ with $s \in rTt$ we have $V_r \cap V_t \subseteq V_s$.

The maximum of the sizes of the V_t minus 1 is the *width* of $(T, \mathcal{V})$. The *adhesion set* V_e corresponding to an edge $e = t_1t_2$ of T is $V_{t_1} \cap V_{t_2}$. The maximum of the sizes of the V_e is the *adhesion* of $(T, \mathcal{V})$. Every orientation (t_1, t_2) of an edge e of T *induces* a separation $(\bigcup_{t\in T_1} V_t, \bigcup_{t\in T_2} V_t)$ of G with separator V_e where $T_1 \ni t_1$ and $T_2 \ni t_2$ are the two components of $T - t_1t_2$ (cf. [Die25, Lemma 12.3.1]). It is immediate from the definition that the set of separations induced by a tree-decomposition is nested. Conversely, it is well-known (e.g. proof of [EKT22, Lemma 2.7] or [Die18b; DO19]) that every nested set N of separations of a finite graph *induces* a tree-decomposition whose induced separations are in bijection with N. We

remark that tree-decompositions $(T, \mathcal{V})$ where T is a path of length M correspond precisely to the families $\mathcal{W} = (W_0, \dots, W_M)$ which satisfy (L1) and (L2). We now recall the tangle-tree duality theorem, rephrased here in the version which we need later:

Theorem 10.5.4 (e.g. [Die25], 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 induced separations 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 to obtain a tree-decomposition of small width and then sort out a suitable linear decomposition by Lemma 10.5.3.

Lemma 10.5.5. *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 will show that $N_1(k, M) = 3k \cdot 3^{(M+2)^{k+1}}$ suffices.

Proof of Lemma 10.5.5. 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 10.5.3, 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. [Die25, 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

$$(A_i, B_i) = (\bigcup_{j<i} W_j, \bigcup_{j \geqslant i} W_j).$$

Thus, $(A_i, B_i) < (C, D) < (A_{i+1}, B_{i+1})$ contradicting our choice of the (A_j, B_j) via Lemma 10.5.3.

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 10.5.5 into the desired rainbow-cloud-decomposition. As an intermediate step, let us find a linear decomposition $\mathcal{W} = (W_0, W_1, \dots, W_M)$ 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 10.5.5.

Lemma 10.5.6 ([KNTW12], Lemma 3.5)**.** *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 $\mathcal{W} = (W_0, \dots, W_{M_1})$ 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).*

[4](FL1) and (FL2) are (L7) and (L8) in [KNTW12], respectively.

Proof. This follows from the proof of [KNTW12, Lemma 3.5], as (L6) is only assumed to guarantee that the resulting decomposition also satisfies (L6). □

We remark that the proof of Lemma 10.5.6 in [KNTW12] shows that $M_1(\ell, M) = (M^{\binom{\ell}{2}+1} \cdot \binom{\ell}{2}!)^{\ell+1} \cdot l!$ suffices.

We now turn to the proof of Theorem 10.5.1. 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 10.5.1. Suppose that G has no $(k+1)$-tangle. Then we set $N(k, M)$ to be $N_1(k, M_1(k, M+2))$, where N_1 and M_1 are as in Lemma 10.5.5 and Lemma 10.5.6, respectively. Then Lemma 10.5.5 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 10.5.6 ,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 $Z' := \bigcap_{i\in\{1,\dots,M'\}} U_i''$. 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 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

$$R := G[\bigcup_{i\in[M'-1]} W_i]$$

and

$$C := G[Z' \cup \bigcup_{i\in[M'-1]} (W_i' \setminus W_i) \cup W_0' \cup W_{M'}'].$$

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. □

10.6 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 44. These tools will allow us to control the new low-order separations that arise when we delete an edge deep inside the rainbow.

10.6.1 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 10.6.1. *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 10.6.2. *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 10.6.1 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. □

10.6.2 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 $\{0, \dots, M\}$ and adding the remaining bags to C as follows. For $0 \leqslant i \leqslant j \leqslant M$, we set

$$\begin{aligned} R_{i,j} &:= G\left[\bigcup_{h=i}^{j} W_h\right] = \bigcup_{h=i}^{j} G[W_h], \\ \mathcal{W}_{i,j} &:= (W_i, \dots, W_j), \\ C_{i,j} &:= G\left[V(C) \cup \left(\bigcup_{h=0}^{i-1} W_h\right) \cup \left(\bigcup_{h=j+1}^{M} W_h\right)\right], \text{ and} \\ (R, \mathcal{W}, Z, C)_{i,j} &:= (R_{i,j}, \mathcal{W}_{i,j}, Z, C_{i,j}). \end{aligned}$$

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 10.6.3. *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 10.6.4. *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 $2\ell + |Z|$.* □

10.6.3 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 $i \in \{0, \dots, 2k\}$ 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 10.3): More precisely, let $(R, \mathcal{W}, Z, C)$ be an RC-decomposition of a

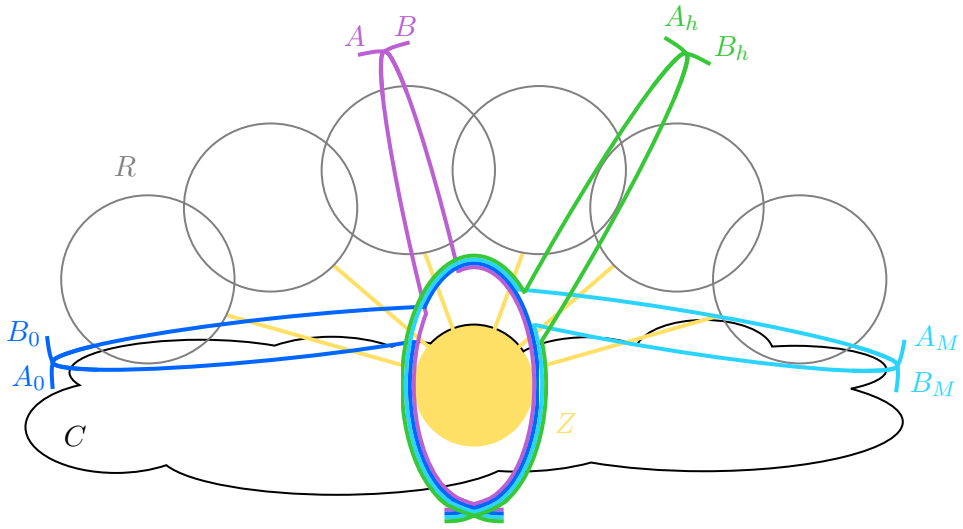

Figure 10.3: A rainbow-crossing separation (A, B) and the arising separations (A^0, B^0), (A^h, B^h) and (A^M, B^M).

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 10.6.5. *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 10.6.4, 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 paths. 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 $\{A^h, B^h\}$:

$$\begin{aligned}
&|A^h \cap B^h| = |((A \cap V(C_{i,j})) \cup V(R_{i,h-1})) \cap (V(R_{h,j}) \cup (B \cap V(C_{i,j})))| \\
&\leqslant |(A \cap B) \cap V(C_{i,j})| + |A \cap (V(C_{i,j}) \cap V(R_{h,j}))| + |B \cap (V(C_{i,j}) \cap V(R_{i,h-1}))| + |V(R_{i,h-1}) \cap V(R_{h,j})| \\
&\leqslant (k - \ell) + |A \cap W_j| + |B \cap W_i| + |U_h| = (k - \ell) + 0 + 0 + \ell = k,
\end{aligned}$$

where the penultimate equation holds because $W_i \subseteq A \setminus B$ and $W_j \subseteq B \setminus A$. This completes the proof. □

10.6.4 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 10.6.6. *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 10.6.7. *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 10.6.8. *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 10.6.1, 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. □

10.7 Proof of Theorem 44 if G has no higher-order tangle

In this section, we complete the proof of Theorem 44. Based on the results in Sections 10.3–10.5, it suffices to consider graphs that admit an RC-decomposition with certain properties. More precisely, the technical main result of this section, which will allow us to finish the proof of Theorem 44, reads as follows:

Theorem 10.7.1. *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 10.7.1. In the very end we combine the previous results to prove Theorem 44.

10.7.1 Living in the rainbow

In the proof of Theorem 10.7.1, we will delete an edge e deep inside the rainbow R of the given long RC-decomposition. With this, we aim to 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 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 10.7.1, 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 $h \in \{i_{A,B} + 1, \dots, j_{A,B}\}$ 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 10.7.2. *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 10.6.4.

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 $B \setminus A \subseteq V(R) \setminus V(C)$ is $\leqslant$-maximal in τ with that property; in particular, $G[B \setminus A]$ is connected. Then $B \setminus A$ cannot contain a bag of $\mathcal{W}$ by Lemma 10.6.6, since $B \setminus A \subseteq V(R) \setminus V(C)$ 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 $B \setminus A$ also meets the separator $A \cap B$. Thus, Lemma 10.6.1 ensures that at most $2k - 2$ bags of $\mathcal{W}$ meet $B \setminus A$. Further, since $G[B \setminus A] \subseteq R$ is connected and each bag W_i separates R, the bags which are met by $B \setminus A$ 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 $B \setminus A$. ■

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 10.6.3, $(R, \mathcal{W}, Z, C)_{i,j}$

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 therefore $Y \smallsetminus X$ 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 10.6.3, $(R, \mathcal{W}, Z, C)_{i,j}$ 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. □

10.7.2 Deleting an edge and orienting the new separations

This subsection is dedicated to the proof of Theorem 10.7.1. We remark that this section is the only part of the proof of Theorem 44 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 10.7.1:

Setting 10.7.3. 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 10.7.4. *If we assume Setting 10.7.3 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 10.7.2 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 10.6.3, 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 10.6.4. 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 10.7.4 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 10.7.1, 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 10.4.4). 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 (Lemma 10.7.6 and Lemma 10.7.7, respectively). This then allows us to show that an even broader class of separations are forced by τ (Lemma 10.7.8). For all these subsequent three lemmas, we assume the following setting:

Setting 10.7.5. Assume Setting 10.7.3 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 10.7.4 yields. In particular, $M \geqslant 8k$.

Lemma 10.7.6. *Assume Setting 10.7.5. 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 10.6.5, $(A^{i+1}, B^{i+1}), (A^j, B^j)$ 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, we have either (A^{i+1}, B^{i+1}), (A^j, B^j) in τ or (B^{i+1}, A^{i+1}), (B^j, A^j) in τ. 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 10.7.7. *Assume Setting 10.7.5. If $\{A, B\} \in S_k(G')$ slices the rainbow, then there is $(X, Y) \in \tau$ with $(A, B) \leqslant (X, Y)$ or $(B, A) \leqslant (X, Y)$ such that $V(R) \subseteq X$. 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 10.6.2 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 10.6.7 yields $2\ell + |Z| \leqslant |A, B| < k$ and therefore $|V(R) \cup Z, V(C)| = 2\ell + |Z| < k$. 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 10.6.7. 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 10.7.8. *Assume Setting 10.7.5. 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 10.6.2 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 and 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 pf $(R, \mathcal{W}, Z, C)$, 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 10.6.2 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 10.6.8 ensures that $\{A, B\}$ either crosses or slices the rainbow in G'. Thus, the previous Lemmas 10.7.6 and 10.7.7 guarantee taht τ forces an orientation of $\{A, B\}$. □

With these tools at hand, we are now ready to prove Theorem 10.7.1:

Proof of Theorem 10.7.1. We may assume Setting 10.7.5, as Setting 10.7.3 with $M_0 \geqslant 18k$ is precisely the premise of Theorem 10.7.1. 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 10.7.8 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 $G[A'_i \smallsetminus B'_i]$ 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 $A'_i \smallsetminus B'_i$ nor in $B'_i \smallsetminus A'_i$ by Lemma 10.6.1. 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 $A'_j \smallsetminus B'_j$ of some (A'_j, B'_j); by renaming the (A'_i, B'_i), we may assume that $j = 1$. Note that as $W_h \subseteq A'_1 \smallsetminus B'_1$ 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 $B'_1 \smallsetminus A'_1$ also contains some bag W_s. By Lemma 10.6.8, (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 10.7.6 and 10.7.7. 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 10.7.6 and 10.7.7 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 $B'_1 \smallsetminus A'_1$. By Lemma 10.6.1, the strict side $A'_1 \smallsetminus B'_1$ 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 did thus 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 10.6.2 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 10.6.2 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 10.6.2 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 10.6.1. 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 10.6.8. Thus, τ forces an orientation of $\{A, B\}$ by Lemmas 10.7.6 and 10.7.7, 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 10.6.1 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 44. Choose $M(k)$ to be $N(k, 18k)$ as in Theorem 10.5.1. 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 10.3.1 and 10.3.2. If there exists a $(k+1)$-tangle in G, then we are done by Theorem 10.4.1. 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 10.5.1. If G has a vertex of degree at most 2, then we are done by Lemmas 10.3.4 and 10.3.5. Thus, Theorem 10.7.1 concludes the proof. □

We remark that one may calculate that $M(k) \in O(3^{k^{k^5}})$.

10.8 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 43. Next, we deduce from Theorem 43 our reduction of Problem 10.1.1 to small graphs, Theorem 40, and then derive Corollaries 41 and 42 from it. Finally, we present a further application of Theorem 43 in Corollary 10.8.4, 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 43. *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 44. Suppose that $G_0, \dots, G_{i-1}$ and $\tau_0, \dots, \tau_{i-1}$ are already defined. If G_{i-1} is disconnected, then we apply Lemma 10.3.3 to obtain G_i and τ_i. If $k \leqslant 2$, then we apply Lemma 10.3.1 or Lemma 10.3.2. If $k \geqslant 3$, but G_{i-1} has a vertex of degree $\leqslant 2$, then we apply Lemma 10.3.4 or Lemma 10.3.5. 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 44, if G_{i-1} has at least $M(k)$ edges. Otherwise, we set $m := i - 1$, completing the proof. □

10.8.1 Application I: sets and functions inducing tangles

In this section, we address Problem 10.1.1.

Problem 10.1.1. *Is every tangle in a graph G induced by some set $X \subseteq V(G)$?*

We will use Theorem 43 to prove Theorem 40, our reduction of Problem 10.1.1 for k-tangles to graphs of size bounded in k. Let us briefly recall the relevant definitions, for which we mostly follow [DEJ24].

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 40 reads as follows:

Theorem 40. *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 40 via Theorem 43, 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 10.8.1. *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 10.8.2. *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 $(A', B') := (A \setminus \{v\}, B)$ is a separation of G', and $(A, B) \in \tau$ witnesses $(A', B') \in \tau'$, as desired. Similarly, if $x, y \in B$ and $v \notin A$, we have $(A', B') := (A, B \setminus \{v\}) \in \tau'$, as desired.

Let us now assume that $x, y \in A$ and $v \in B$. Then $(A, B \setminus \{v\})$ is a separation of G. Since the only neighbours x, y of v are in A, the separations (A, B) and $(B \setminus \{v\}, A)$ form a forbidden tuple in G. Hence, $(A, B \setminus \{v\})$ is in the tangle τ. Now the above described case yields $(A', B') := (A \setminus \{v\}, B \setminus \{v\}) \in \tau'$, as desired. Similarly, if $u, w \in B$ and $v \in A$, then $(A', B') := (A \setminus \{v\}, B \setminus \{v\}) \in \tau'$.

So to conclude the proof, we may assume that $x \in A \setminus B$ and $y \in B \setminus A$ by possibly renaming x, y; in particular $v \in A \cap B$, as $xv, vy \in E(G)$. Hence, $(C, D) := (A \cup \{y\}, B \setminus \{v\})$ 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 $(A', B') := ((A \setminus \{v\}) \cup \{y\}, B \setminus \{v\}) = (C \setminus \{v\}, D) \in \tau'$, 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 10.8.3. *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 40. Let τ_m be the k-tangle in the graph G_m with less than $M(k)$ edges as described in Theorem 43. Let w' be a weight function on $V(G')$ which induces the tangle τ_m. Iteratively applying Lemmas 10.8.1 to 10.8.3 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 40, we now deduce Corollary 41 and Corollary 42:

Corollary 41. *For $k \geqslant 1$, there exists $M = M(k) \in O(3^{k^{k^5}})$ such that Problem 10.1.1 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 40. □

Corollary 42. *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 10.1.1 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 40.

Proof. We enumerate all the finitely many non-isomorphic connected graphs $G_1, \dots, G_m$ 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 10.1.2 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 40 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 40 by choosing the weight functions as indicator functions of the inducing sets given by the assumed positive answer to Problem 10.1.1. □

We remark that the proof of Corollary 42 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 40.

10.8.2 Application II: subgraphs witnessing a tangle

In this section, we demonstrate another application of our inductive proof method Theorem 43 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 [GS15b, Lemma 3.1][5] 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 43, we obtain a new bound which is significantly better for sufficiently large k:

[5]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.

Corollary 10.8.4. *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 10.8.4. Let $M(k)$ be given by Theorem 43, and set $M'(k) := 2M(k)$. Let τ be a k-tangle in some graph G. Theorem 43 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 43.

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 43 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. □

CHAPTER 11

A grid theorem for strong immersions of walls

This chapter is based on [DJKW23], which is joint work with Reinhard Diestel, Raphael W. Jacobs and Paul Wollan.

11.1 Introduction

In their Graph Minors Project [RS], Robertson and Seymour investigated the structure of graphs not containing a fixed graph as a minor. An important example of their structure theorems as well as a result central to their project is the *grid theorem* [RS86b] which asserts the equivalence of large tree-width and the existence of large grid minors. More precisely, every graph of large enough tree-width contains a large grid as a minor, and large grid minors form an obstruction to small tree-width in that large grids have large tree-width and every graph has tree-width at least the tree-width of each of its minors. An equivalent formulation of the grid theorem is that a graph has large tree-width if and only if it contains a large wall as a topological minor [RS95].

Grid theorem-like results, which provide the equivalence of large width with respect to some width-measure and a large (often grid-like) substructure in terms of a corresponding containment relation, have since been proven in various other settings. For example, Kreutzer and Kawarabayashi proved a directed grid theorem [KK15], which establishes an equivalence of large directed tree-width and the existence of butterfly minors of large directed grids, and Geelen, Kwon, McCarty and Wollan showed in [GKMW23] that a graph has large rank-width if and only if it contains a vertex-minor isomorphic to a large comparability-grid.

In this chapter we are concerned with the structure of graphs not containing a large wall as a strong immersion minor. As the minor relation, immersion minors generalise the notion of topological minors, but they are unrelated to minors in that neither the existence of an immersion minor implies the existence of a minor nor vice-versa. Immersion minors have attracted attention from various algorithmic and structural viewpoints in recent years, often finding results analogous to those for minors [AL03; DDFMMS14; DK14b; DW16; FGTW08; GKMW11; KK24; RS10; Wol15].

There are two versions of the immersion relation, a weak one and a strong one. They are defined as follows. A *weak immersion* of a graph H in a graph G is a map α with domain $V(H) \cup E(H)$ that embeds $V(H)$ into $V(G)$ and maps every edge $uv \in H$ to an $\alpha(u)$–$\alpha(v)$ path in G which is edge-disjoint from every other such path. If, additionally, these paths have no internal vertices in $\alpha(V(H))$, then α is a *strong immersion* of H in G. The vertices in $\alpha(V(H))$ are the *branch vertices* of this immersion. We say that H is *weakly/strongly immersed* in G, or that H is a *weak/strong immersion minor* of G, if there is a weak/strong immersion of H in G.

While the (topological) minor relation behaves well with tree-decompositions, the immersion relations do so with *tree-cut decompositions*, which form the edge-cut analogue of tree-decompositions. Formally, a pair $(T, \mathcal{X})$ is a *tree-cut decomposition* of a graph G if T is a tree and $\mathcal{X} = (X_t)_{t \in T}$ a near-partition of $V(G)$, that is, the X_t are disjoint and their union is $V(G)$ but, in contrast to a partition, the X_t may be empty. The vertex sets X_t are the *parts* of $(T, \mathcal{X})$. The *torso of* $(T, \mathcal{X})$ *at a node* $t \in T$ arises from G by identifying for every component T' of $T - t$ the vertices in $\bigcup_{t' \in T'} X_{t'}$ to a single vertex, keeping parallel edges as they arise but omitting loops. These new identification vertices are the *peripheral vertices* of the torso, while the ones in X_t are its *core vertices*. Every edge $t_1t_2 \in T$ induces its *adhesion set* $E_G(\bigcup_{t \in T_1} X_t, \bigcup_{t \in T_2} X_t)$ where T_1 and T_2 are the two components of $T - t_1t_2$ with $t_1 \in T_1$ and $t_2 \in T_2$ and, for $X, Y \subseteq V(G)$, $E_G(X, Y)$ denotes the set of edges xy of G with $x \in X$ and $y \in Y$. The *adhesion* of a tree-cut decomposition then is the maximum size of its adhesion sets.

DeVos, McDonald, Mohar and Scheide [DMMS13] and Wollan [Wol15] independently showed that if a graph does not admit K^ℓ, the complete graph of size ℓ, as a weak immersion, then it admits a tree-cut decomposition each of whose torsos contains at most ℓ vertices of degree at least ℓ^2. As a qualitative converse, they proved that a tree-cut decomposition in which each torso contains at most ℓ vertices of degree at least ℓ precludes the existence of a weak immersion of $K^{\ell+1}$. Introducing the concept of tree-cut width, Wollan [Wol15] also derived a grid theorem for weak immersions, which establishes the equivalence of large tree-cut width and the existence of weak immersions of large walls. His proof made use of the classical grid theorem for excluded wall minors by Robertson and Seymour [RS86b].

For strong immersions Dvořák and Wollan proved the following general structure theorem describing the structure of graphs not containing a fixed graph as a strong immersion minor.

Theorem 11.1.1 ([DW16], Theorem 4)**.** *For every graph F, there exists an integer $\alpha = \alpha(F) > 0$ such that if a graph G does not contain F as a strong immersion minor, then there exists a tree-cut decomposition of G of adhesion less than α such that each of its torsos is α-basic.*

Here, the α-basicness of a graph means, roughly speaking, that it has a path-like decomposition in which we can accommodate vertices of high degree.

When we restrict F to complete graphs in Theorem 11.1.1, then there is also a qualitative converse of Theorem 11.1.1 in [DW16]: For every integer $\alpha > 0$, there exists an integer $n = n(\alpha) > 0$ such that any graph with a tree-cut decomposition which has adhesion less than α and whose torsos are α-basic does not contain K^n as a strong immersion minor.

If we consider strong immersions of walls, as it is the subject of this chapter, and we thus look at Theorem 11.1.1 with F restricted to walls, then Theorem 11.1.1 does not have such a qualitative converse. Indeed, for every integer $\ell > 0$, the *wall* W_ℓ *of size* ℓ is formally defined as the graph arising from the $\ell \times 2\ell$ grid by deleting all edges $(i, j)(i', j')$ with $j = j'$, $i = i' + 1$ and $j \equiv i \ (mod\ 2)$ and then removing the two resulting vertices of degree 1 (see Figure 11.1). By definition the wall W_ℓ has $2\ell^2 - 2$ vertices and maximum degree 3. Thus, W_ℓ does not contain K^5 as a strong immersion minor. Theorem 11.1.1 hence yields an integer $\alpha > 0$ such

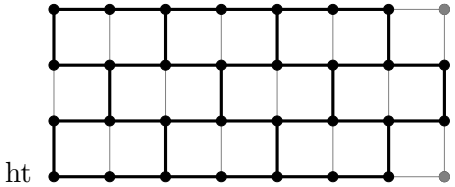

Figure 11.1: The wall W_4 of size 4 with the underlying 4×8 grid.

that every wall, no matter how large its size, has a tree-cut decomposition of adhesion less than α such that each torso is α-basic.

In this chapter, we prove a grid theorem-like result for the strong immersion relation which says that a graph contains a large wall as a strong immersion minor if and only if the graph has a tree-cut decomposition of a specific type. These tree-cut decompositions are inspired by those in Theorem 11.1.1 in that their torsos still admit a path-like structure, but path-likeness is defined differently, as follows.

A graph G is *α-thin* for an integer $\alpha > 0$ if there exists an enumeration $v_1, \dots, v_n$ of its vertices such that there are at most α edges in G joining $\{v_1, \dots, v_{i-1}\}$ and $\{v_{i+1}, \dots, v_n\}$ for $i = 1, \dots, n$, and we call G *almost α-thin* if it becomes α-thin after removing up to α vertices each of which has at most α neighbours in G. In our specific tree-cut decompositions, the torsos then have almost α-thin '3-centres'. Following [Wol15], the 3-*centre* of a torso of a tree-cut decomposition arises from the torso by a maximal sequence of deleting peripheral vertices of degree at most 1 and suppressing peripheral vertices of degree 2, removing arising loops. We will show that if a torso of a tree-cut decomposition is itself (almost) α-thin, then so is its 3-centre, while the converse does not hold. For a further discussion of these notions, we refer the reader to Section 11.2.

Our first main result, a version of Theorem 11.1.1 specifically tailored to walls, then reads as follows:

Theorem 45. *For every integer $\ell > 0$, there exists an integer $\alpha = \alpha(\ell) > 0$ such that if a graph G does not contain the wall W_ℓ as a strong immersion minor, then G has a tree-cut decomposition of adhesion at most α such that the 3-centres of its torsos are almost α-thin.*

We will see that 3-centres are indeed necessary in Theorem 45 in that we cannot consider the torsos themselves instead of their 3-centres (see Example 11.2.6).

In contrast to Theorem 11.1.1, the notion of path-like torsos as in Theorem 45 now gives the desired equivalence of the existence of strong immersions of large walls and these tree-cut decompositions. More precisely, our second main result provides a qualitative converse of Theorem 45:

Theorem 46. *For every integer $\alpha > 0$, there exists an integer $\ell = \ell(\alpha) > 0$ such that if a graph G has a tree-cut decomposition of adhesion at most α such that the 3-centres of its torsos are almost α-thin, then G does not contain the wall W_ℓ as a strong immersion minor.*

While the proof of Theorem 46 is self-contained and does not build on any previous results, the proof of Theorem 45 draws on the grid theorem for excluded wall minors as well as on

several previously established methods for immersions and tree-cut decompositions [DW16; RS86b; Wol15]. We emphasise that we do not make use of any previous structure theorem for weak or strong immersions, but rather use some of their ideas and combine them in a novel way to fit our specific problem.

This chapter is organised as follows. We first discuss the notions of almost α-thin graphs and 3-centres in Section 11.2. Then we prove Theorem 46 in Section 11.3 and Theorem 45 in Section 11.4.

For basic graph-theoretic terms, we follow [Die25]. In this chapter, we only consider graphs without loops. But unless called *simple*, they may contain parallel edges. As a consequence, the *degree* of a vertex is the number of its incident edges which may in general differ from the number of its neighbours. Moreover, we diverge from [Die25] in that we define the *components* of a graph not as its maximal connected subgraphs, but as their vertex sets.

11.2 Almost α-thin graphs and 3-centres

In this section we take a closer look into our notion of path-likeness, (almost) α-thin graphs, and its relation to the concept of 3-centres. Let us first recall the definition of (almost) α-thinness from the introduction.

Definition 11.2.1. Let $\alpha > 0$ be an integer. A graph G is *α-thin* if there exists an enumeration $v_1, \dots, v_n$ of its vertices satisfying that $|E_G(\{v_1, \dots, v_{i-1}\}, \{v_{i+1}, \dots, v_n\})| \leqslant \alpha$ for $i = 1, \dots, n$,[1] and the graph G is *almost α-thin* if there exists a set X of at most α vertices of G such that each vertex in X has at most α neighbours in G and $G - X$ is α-thin.

We remark that if a graph is (almost) α-thin, then so are all its subgraphs.

Even though adding vertices of degree 1 to a graph does not change whether it contains a wall as a strong immersion minor, such vertices affect whether a graph is (almost) α-thin.

Example 11.2.2. *For every integer $\alpha > 0$, the star $K_{1,n}$ with $n = 3\alpha+3$ leaves is not almost α-thin, but becomes 0-thin after repeatedly removing vertices of degree 1.*

Proof. Repeatedly removing its vertices of degree at most 1, the star $G := K_{1,n}$ becomes the graph on one vertex which is trivially 0-thin. So it remains to argue that G is not almost α-thin, and we suppose for a contradiction that there exists a set X of at most α vertices of G, each with at most α neighbours in G, such that $G - X$ is α-thin.

The set X cannot contain the centre c of G since it has $n > \alpha$ many neighbours in G. Thus, X consists of up to α leaves of G, and $G - X$ is again a star with centre c, now with at least $n' \geqslant 2\alpha + 3$ leaves. Now take an arbitrary enumeration $v_1, \dots, v_{n'+1}$ of $V(G - X)$,

[1] Note that the definition of α-thin differs from the similar notion of *cutwidth* at most α, since cutwidth also takes the edges joining v_i and $\{v_{i+1}, \dots, v_n\}$ into account. While cutwidth bounds thinness from above, there are graphs with unbounded cutwidth which are 0-thin, as paths with many parallel edges witness.

and let $i \in \{1, \dots, n'+1\}$ with $v_i = c$. If $i \geqslant \alpha + 3$, then $\{v_1, \dots, v_{i-2}\}$ contains at least $\alpha + 1$ leaves of the star, while if $i \leqslant (n'+1) - (\alpha + 3)$, then $\{v_{i+2}, \dots, v_{n'+1}\}$ contains at least $\alpha + 1$ leaves. So for $j = i - 1$ in the first case and $j = i + 1$ in the second case, there are at least $\alpha + 1$ edges joining $\{v_1, \dots, v_{j-1}\}$ and $\{v_{j+1}, \dots, v_{n'+1}\}$, which yields the desired contradiction. □

Similar to the addition of vertices of degree 1, subdividing edges has an impact on (almost) α-thinness.

Example 11.2.3. *For every integer $\alpha > 0$, the graph G that arises from two disjoint stars, each with $n = 3\alpha + 3$ leaves, by identifying their leaves is not almost α-thin, but becomes 0-thin after repeatedly suppressing vertices of degree 2.*

Proof. If we repeatedly suppress vertices of degree 2, then G turns into the graph consisting of two vertices which are joined by n edges. This graph is trivially 0-thin. Since G contains $K_{1,n}$ as a subgraph and $K_{1,n}$ is not almost α-thin by Example 11.2.2, G cannot be almost α-thin itself. □

Just as a graph stays (almost) α-thin when we delete some of its vertices, suppressing vertices of degree 2 also does not impact the (almost) α-thinness of a graph, as the next lemma demonstrates:

Lemma 11.2.4. *Let G be an (almost) α-thin graph for some integer $\alpha > 0$. Then suppressing a vertex of G of degree 2 (and deleting a potentially arising loop) results in an (almost) α-thin graph.*

Proof. It is enough to prove the lemma for α-thinness, since the case of almost α-thinness then follows immediately. So let G be an α-thin graph, let $v \in G$ be a vertex of degree 2 and let u and w be the other incident vertices of the two edges of G incident to v. Now let G' arise from G by suppressing v, omitting a potentially arising loop.

If $u = w$, then G' is a subgraph of G which is hence α-thin; so we may assume that u and w are distinct. Given an enumeration $v_1, \dots, v_{n+1}$ of $V(G)$ witnessing that G is α-thin, let $k \in \{1, \dots, n+1\}$ with $v_k = v$. We may assume that v appears in this enumeration in between u and w; otherwise, moving v in between them only reduces the number of edges joining $\{v_1, \dots, v_{i-1}\}$ and $\{v_{i+1}, \dots, v_{n+1}\}$ for $i = 1, \dots, n+1$. We then define the enumeration $v'_1, \dots, v'_n$ of $V(G')$ as follows:

$$v'_i := \begin{cases} v_i & \text{if } i < k; \\ v_{i+1} & \text{if } i > k. \end{cases}$$

This definition guarantees that for $i = 1, \dots, n$, the edges of G' joining the sets $\{v'_1, \dots, v'_{i-1}\}$ and $\{v'_{i+1}, \dots, v'_n\}$ coincide with the edges of G joining $\{v_1, \dots, v_{j-1}\}$ and $\{v_{j+1}, \dots, v_{n+1}\}$, where $j := i$ if $i < k$ and $j := i + 1$ if $i \geqslant k$, except that the edge of G' that arose by suppressing v is replaced by one of the two edges incident to v in G. Thus, G' is α-thin as witnessed by the enumeration $v'_1, \dots, v'_n$ of $V(G')$. □

As we will see below, we need to ignore peripheral vertices of degree at most 2 in torsos to establish Theorem 45 (see Example 11.2.6). Formally, we achieve this with the notion of 3-centres which was introduced in [Wol15] and is well-defined by [Wol15, Lemma 9].

Definition 11.2.5. Let X be a set of vertices of a graph G. The 3-*centre of* (G, X) arises from G by a maximal sequence of deleting vertices of degree at most 1 and suppressing vertices of degree 2 (omitting any resulting loops) where all these vertices are not in X. If $(T, \mathcal{X})$ is a tree-cut decomposition of G, then the 3-*centre of the torso* H_t *of* $(T, \mathcal{X})$ *at the node* $t \in T$ is defined as the 3-centre of (H_t, X_t).

We note that the 3-centre of a pair (G, X) cannot be formed by first repeatedly deleting vertices of degree 1 and then suppressing all vertices of degree 2. Indeed, the deletion of resulting loops may create new vertices of degree 1 which have to be deleted afterwards; consider for example the graph G which consists of two triangles joined by an edge and let X be the set of vertices of one of the triangles.

If a graph G is (almost) α-thin, then the 3-centre of (G, X) is also (almost) α-thin since deleting vertices and suppressing vertices of degree 2 does not affect the (almost) α-thinness of a graph, see Lemma 11.2.4. In particular, it is weaker to assume in Theorem 46 that the 3-centres of the torsos are almost α-thin than to assume that the torsos themselves have this property.

As we have seen in Example 11.2.2 and Example 11.2.3, the converse statement is not true: If the 3-centre of (G, X) is α-thin, then G does not even have to be almost α-thin itself. In fact, the converse holds in an even stronger form in that we cannot remove 3-centres from Theorem 45 and consider the torsos themselves instead, as the following example demonstrates.

Example 11.2.6. *For every two integers* $\ell \geqslant 2$ *and* $\alpha > 0$*, there exists a graph which does not contain the wall* W_ℓ *as a strong immersion minor and also does not admit a tree-cut decomposition of adhesion at most* α *whose torsos are almost* α*-thin.*

Proof. Let $G = K_{1,n}$ be the star with $n = \alpha(3\alpha + 3)$ leaves and centre $c \in V(G)$. Clearly, G does not contain W_ℓ as a strong immersion minor since G contains only a single vertex of degree more than 1 while W_ℓ has several of those because of $\ell \geqslant 2$. Thus, it remains to consider an arbitrary tree-cut decomposition $(T, \mathcal{X})$ of G of adhesion at most α and to show that at least one of the torsos of $(T, \mathcal{X})$ is not almost α-thin.

Let s be the (unique) node of T such that X_s contains the centre c of the star G. For each neighbour t of s in T, let $Y_t := \bigcup_{t' \in T'} X_{t'}$ where T' is the component of $T - st$ containing t. Since $(T, \mathcal{X})$ has adhesion at most α and $c \in X_s$, the set Y_t contains at most α leaves of G. But G has $\alpha(3\alpha + 3)$ leaves, so it yields a star G' with at least $3\alpha + 3$ leaves in the torso H_s of $(T, \mathcal{X})$ at s. This star G', however, is not almost α-thin by Example 11.2.2, and hence H_s is not almost α-thin either. □

If we adapt Example 11.2.6 by doubling each edge in G, then this modified version also demonstrates that just deleting peripheral vertices of degree 1 in each torso instead of taking its 3-centre does not suffice either.

11.3 Strongly immersed walls are obstructions

In this section, we prove Theorem 46 which we recall here:

Theorem 46. *For every integer $\alpha > 0$, there exists an integer $\ell = \ell(\alpha) > 0$ such that if a graph G has a tree-cut decomposition of adhesion at most α such that all the 3-centres of its torsos are almost α-thin, then G does not contain the wall W_ℓ as a strong immersion minor.*

As a key ingredient to our proof of Theorem 46, we need Lemma 11.3.1, which asserts that the wall W_ℓ contains a vertex set Z of size $\ell - 2$ which is *well-linked* in W_ℓ, that is, for every two disjoint $A, B \subseteq Z$ with $k := |A| = |B|$, there exist k vertex-disjoint A–B paths in W_ℓ. The existence of such a set Z of size ℓ follows from the facts that the tree-width of a wall of size ℓ is at least ℓ and the largest size of a well-linked set in a graph is lower bounded by its tree-width (see for example the survey [HW17]). We include a direct and constructive proof in the case of walls here which yields a further structural property of the well-linked set.

Lemma 11.3.1. *The wall W_ℓ of size $\ell \geqslant 3$ contains a set Z of vertices of size $\ell - 2$ such that Z is well-linked in W_ℓ and such that every two vertices in Z are joined by three internally vertex-disjoint paths in W_ℓ.*

Proof. Let Z be the set of all vertices of W_ℓ which have the form (i, ℓ) and degree 3. Note that Z has size $\ell - 2$ by the definition of W_ℓ. It is easy to see that every two vertices of degree 3 are joined by three internally vertex-disjoint paths in W_ℓ as $\ell \geqslant 3$. Thus, it remains to show that for every two disjoint subsets A, B of Z with $k := |A| = |B|$, there exist k vertex-disjoint A–B paths in W_ℓ. By Menger's theorem [Men27], it suffices to prove that for every set $X \subseteq V(G)$ of size at most $k - 1$, there exists an A–B path in W_ℓ which avoids X.

For every $j = 1, \ldots, \ell$ let V_j be the j-th *vertical path* of the wall W_ℓ, that is, V_j is the induced subgraph of W_ℓ on the set of all vertices of the form $(i, 2j)$ or $(i, 2j - 1)$. Note that these vertical paths are pairwise vertex-disjoint, that each vertical path meets Z in at most one vertex, and that each vertex in Z is contained in some vertical path. Thus, the pigeonhole principle yields a vertical path V_{j_A} which meets A, but avoids X, since $|A| = k > k - 1 \geqslant |X|$; analogously, we obtain a vertical path V_{j_B} for B.

Next, for every $i = 1, \ldots, \ell$ let H_j be the i-th *horizontal path* of the wall W_ℓ, that is, the subgraph of W_ℓ induced on the set of all vertices of the form (i, j). Note that these horizontal paths are pairwise vertex-disjoint. Thus, the pigeonhole principle yields an horizontal path H_{i_0} which avoids X, since $\ell > |A| > |X|$.

By definition, every horizontal path meets every vertical path. Hence, $V_{j_A} + H_{i_0} + V_{j_B}$ is a connected subgraph of W_ℓ which meets both A and B, but avoids X. In particular, there exists an A–B path in W_ℓ which avoids X, as desired. □

Proof of Theorem 46. We set

$$\ell := \ell(\alpha) := [(\alpha^2 + 1)(2(\alpha + 1) + 4) + \alpha^2 + \alpha] + [\alpha((\alpha^2 + 1)(2(\alpha + 1) + \alpha + 2) + \alpha^2 + \alpha)] + 2,$$

and claim that this ℓ is as desired. So let G be a graph, and let $(T, \mathcal{X})$ be a tree-cut decomposition of G of adhesion at most α such that the 3-centres of its torsos are almost α-thin. Suppose for a contradiction that the wall W_ℓ is strongly immersed in G, and let U be the set of branch vertices of this strong immersion.

By Lemma 11.3.1, W_ℓ contains a vertex set Z of size at least $\ell - 2$ such that Z is well-linked in W_ℓ and every two vertices in Z are joined by three internally vertex-disjoint paths in W_ℓ. Let $Z_U \subseteq U$ be the set of branch vertices corresponding to $Z \subseteq V(W_\ell)$ in G via the strong immersion. Since Z is well-linked in W_ℓ, the definition of strong immersions yields that this set Z_U has the following property:

> *For every two disjoint sets $A, B \subseteq Z_U$ with $k := |A| = |B|$, there exist k edge-disjoint A–B paths in G and these k paths meet each vertex of U at most once.* (∗)

Moreover, since every two vertices in Z are joined by three internally vertex-disjoint paths in W_ℓ, every two vertices in Z_U are joined by three edge-disjoint paths in G.

Consider an edge $t_1t_2 \in T$ and let $Y_1 := \bigcup_{t \in T_1} X_t$ where T_1 is the component of $T - t_1t_2$ containing t_1, and define Y_2 analogously. Since $(T, \mathcal{X})$ has adhesion at most α, the adhesion set $E_G(Y_1, Y_2)$ of G induced by t_1t_2 has size at most α. Therefore, precisely one Y_i, say Y_2, contains at least $\alpha + 1$ vertices from Z_U by (∗) and since $\ell - 2 \geqslant 2(\alpha + 1)$. We then orient the edge t_1t_2 from t_1 to t_2. In this way, Z_U induces an orientation of the edges of T such that for each node $t \in T$ at most one incident edge is oriented away from t.

Hence, there is a unique sink s of this orientation, that is, a node $s \in T$ such that all incident edges are oriented towards s. Let H be the torso of $(T, \mathcal{X})$ at s. For each neighbour t of s in T, let $Y_t := \bigcup_{t' \in T'} X_{t'}$ where T' is the component of $T - st$ containing t, and let v_t be the peripheral vertex of H corresponding to the identification of the vertices in Y_t.

Consider the set Z_H of vertices of H corresponding to the vertices in Z_U, that is, the set consisting of the core vertices $Z_U \cap X_s$ together with those peripheral vertices v_t with $Z_U \cap Y_t \neq \emptyset$. This set Z_H is then contained in the 3-centre $\bar{H}$ of the torso H. Indeed, for every peripheral vertex $v_t \in Z_H$, the set Z_U contains a vertex in Y_t by the definition of Z_H. Since st is oriented towards s, the set Z_U also contains a vertex which is not in Y_t. But these two vertices, just as any two vertices in Z_U, are joined by three edge-disjoint paths in G. Thus, v_t is never a candidate for deletion or suppression in the construction of $\bar{H}$ from H.

For every peripheral vertex v_t in Z_H, we have $|Z_U \cap Y_t| \leqslant \alpha$, since st was oriented towards s. This implies that $|Z_H| \geqslant |Z_U \cap X_s| + |Z_U \setminus X_s|/\alpha$. In particular, Z_H contains either many core vertices or many peripheral vertices of H.

We now aim to use the set Z_H to derive a contradiction to the 3-centre $\bar{H}$ of the torso H being almost α-thin. To do so, consider a set X of at most α vertices of $\bar{H}$ each of which has at most α neighbours in $\bar{H}$ and an enumeration $v_1, \dots, v_n$ of $V(\bar{H} - X)$ which witnesses that $\bar{H} - X$ is α-thin. Observe that if we remove the at most α^2 neighbours of the set X among the v_i, then this partitions the enumeration $v_1, \dots, v_n$ into a set $\mathcal{I}$ of at most $\alpha^2 + 1$ intervals.

If Z_H contains many core vertices of H in that $|Z_H \cap X_s| \geqslant (\alpha^2 + 1)(2(\alpha + 1) + 4) + \alpha^2 + \alpha$,

then $\{v_1, \dots, v_n\}$ contains at least $(\alpha^2+1)(2(\alpha+1)+4)$ core vertices in Z_H which are not neighbours of X. Thus, the pigeonhole principle yields an interval $v_j, \dots, v_k$ in $\mathcal{I}$ containing at least $2(\alpha+1)+4$ elements of $Z_H \cap X_s$; shortening the interval if necessary, we may assume additionally that it contains precisely $2(\alpha+1)+4$ elements of $Z_H \cap X_s$ and that its endpoints v_j and v_k are among those.

Let A_H be the set consisting of the $2(\alpha+1)+2$ elements of $Z_H \cap X_s$ in the interior $\{v_{j+1}, \dots, v_{k-1}\}$ of our fixed interval, and let B_H be a subset of $Z_H \cap X_s$ of the same size as A_H and which is disjoint from $\{v_j, \dots, v_k\}$; such a set B_H exists since $Z_H \cap X_s$ is large enough. Now A_H and B_H are subsets of Z_H consisting of core vertices of H, and hence $A_H, B_H \subseteq Z_U$. So we may apply $(*)$ to $A_H, B_H \subseteq Z_U$ and obtain $2(\alpha+1)+2$ edge-disjoint A_H–B_H paths in G such that at most two of these paths meet v_j or v_k, since v_j and v_k are contained in $Z_H \cap X_s \subseteq Z_U \subseteq U$.

Since the torso H arises from G by identifications of vertices keeping parallel edges as they arise, these $2(\alpha+1)+2$ paths induce $2(\alpha+1)+2$ edge-disjoint A_H–B_H walks in H which we can shorten to $2(\alpha+1)+2$ edge-disjoint A_H–B_H paths in H. By the definition of 3-centres, these $2(\alpha+1)+2$ paths in the torso H in turn induce $2(\alpha+1)+2$ edge-disjoint A_H–B_H paths in the 3-centre $\bar{H}$ of H. We remark that throughout this process we maintain the property that at least $2(\alpha+1)$ of these paths avoid v_j and v_k.

Since none of the vertices $v_j, \dots, v_k$ is a neighbour of X, these $2(\alpha+1)$ edge-disjoint paths joining $\{v_{j+1}, \dots, v_{k-1}\}$ and $\{v_1, \dots, v_{j-1}\} \cup \{v_{k+1}, \dots, v_n\}$ indeed contain at least $2(\alpha+1)$ edges in $\bar{H} - X$ joining $\{v_{j+1}, \dots, v_{k-1}\}$ and $\{v_1, \dots, v_{j-1}\} \cup \{v_{k+1}, \dots, v_n\}$. In particular, there are either at least $\alpha+1$ edges joining $\{v_1, \dots, v_{j-1}\}$ and $\{v_{j+1}, \dots, v_n\}$ or at least $\alpha+1$ edges joining $\{v_1, \dots, v_{k-1}\}$ and $\{v_{k+1}, \dots, v_n\}$. In both cases we obtain a contradiction to the choice of the enumeration $v_1, \dots, v_n$ to witness the α-thinness of $\bar{H} - X$.

If Z_H does not contain many core vertices as in the above case, then $|Z_H| \geqslant |Z_U \cap X_s| + |Z_U \setminus X_s|/\alpha$ implies that Z_H contains many peripheral vertices in that $|Z_H \setminus X_s| \geqslant (\alpha^2+1)(2(\alpha+1)+\alpha+2)+\alpha^2+\alpha$ by our choice of ℓ. We now proceed analogously to the above case of many core vertices in Z_H.

By the lower bound on the size of $Z_H \setminus X_s$, there are at least $(\alpha^2+1)(2(\alpha+1)+\alpha+2)$ peripheral vertices in Z_H which are neither neighbours of X nor contained in X. So the pigeonhole principle yields an interval $v_j, \dots, v_k$ in $\mathcal{I}$ which contains at least $2(\alpha+1)+\alpha+2$ elements of $Z_H \setminus X_s$. We may assume additionally, by shortening the interval if necessary, that the interval contains precisely $2(\alpha+1)+\alpha+2$ elements of $Z_H \setminus X_s$ and that its endpoints v_j and v_k are among those.

Let A_H be the set of the $2(\alpha+1)+\alpha$ elements of $Z_H \setminus X_s$ in the interior $\{v_{j+1}, \dots, v_{k-1}\}$ of our fixed interval, and let B_H be a set of $2(\alpha+1)+\alpha$ elements of $Z_H \setminus X_s$ that are not in $\{v_j, \dots, v_k\}$; such a set B_H exists since $Z_H \setminus X_s$ is large enough. Since each peripheral vertex v_t in Z_H satisfies $Z_U \cap Y_t \neq \emptyset$, we may replace each element v_t of A_H and B_H by an arbitrary vertex in $Z_U \cap Y_t$ to obtain subsets A_U and B_U of Z_U, respectively, which are disjoint and of the same size as A_H and B_H by the disjointness of the Y_t. Applying $(*)$ to $A_U, B_U \subseteq Z_U$

then yields $2(\alpha+1)+\alpha$ edge-disjoint A_U–B_U paths in G.

The construction of A_U and B_U now guarantees that these $2(\alpha+1)+\alpha$ many A_U–B_U paths in G induce $2(\alpha+1)+\alpha$ edge-disjoint A_H–B_H paths in $\bar{H}$ as above. Since $(T,\mathcal{X})$ has adhesion at most α and $v_j, v_k \in Z_H \setminus X_s$ are neither contained in A_H nor in B_H, each of the peripheral vertices v_j and v_k lies on at most $\alpha/2$ of these paths. Thus, there are at least $2(\alpha+1)+\alpha-2\cdot\alpha/2 = 2(\alpha+1)$ edge-disjoint paths in $\bar{H}$ joining $\{v_{j+1},\dots,v_{k-1}\}$ and $\{v_1,\dots,v_{j-1}\}\cup\{v_{k+1},\dots,v_n\}$ which avoid v_j and v_k. Since $\{v_j,\dots,v_k\}$ contains no neighbours of X, these paths contain at least $2(\alpha+1)$ edges in $\bar{H}-X$ joining $\{v_{j+1},\dots,v_{k-1}\}$ and $\{v_1,\dots,v_{j-1}\}\cup\{v_{k+1},\dots,v_n\}$. Thus, there are either at least $\alpha+1$ edges joining $\{v_{j+1},\dots,v_{k-1}\}$ and $\{v_1,\dots,v_{j-1}\}$ or at least $\alpha+1$ joining $\{v_{j+1},\dots,v_{k-1}\}$ and $\{v_{k+1},\dots,v_n\}$. In particular, there are either at least $\alpha+1$ edges joining $\{v_{j+1},\dots,v_n\}$ and $\{v_1,\dots,v_{j-1}\}$ or at least $\alpha+1$ edges joining $\{v_1,\dots,v_{k-1}\}$ and $\{v_{k+1},\dots,v_n\}$, a contradiction. □

11.4 Excluding strongly immersed walls

This section is devoted to the proof of Theorem 45, recalled here.

Theorem 45. *For every integer $\ell > 0$, there exists some integer $\alpha = \alpha(\ell) > 0$ such that if a graph G does not contain the wall W_ℓ as a strong immersion minor, then G has a tree-cut decomposition of adhesion at most α such that the 3-centres of its torsos are almost α-thin.*

For the proof of Theorem 45, we first reduce the problem to 3-edge-connected graphs in Section 11.4.1. In Section 11.4.2 we collect some tools from previous research dealing with immersions and tree-cut decompositions. Then we finally prove Theorem 45 in Section 11.4.3.

11.4.1 Reduction to 3-edge-connected graphs

We begin by reducing the proof of Theorem 45 to 3-edge-connected graphs. This reduction will be done separately for each integer $\ell > 0$; so throughout this section, let $\ell > 0$ be an arbitrary, but fixed integer. We first reduce our attention to graphs with minimum degree 3 which will then facilitate the reduction to 3-edge-connected graphs.

Lemma 11.4.1. *If Theorem 45 holds for all graphs with minimum degree 3, then it holds for all graphs.*

Proof. Let $\alpha := \alpha(\ell) \geqslant 2$ be an integer such that Theorem 45 holds for all graphs with minimum degree 3. We now show that Theorem 45 holds with the same α for all graphs G, and we will do so by induction on $|G|$. The case $|G| = 1$ is trivial; so suppose $|G| > 1$, and assume that G does not contain the wall W_ℓ as a strong immersion minor.

Let $v \in G$ be a vertex with degree $d_G(v)$ at most 2 in G, and let G' arise from G by deleting v if $d_G(v) \leqslant 1$ and suppressing v (omitting any loops which arise) if $d_G(v) = 2$. Clearly, G' does not contain the wall W_ℓ as a strong immersion minor, as G does not do so. Since $|G'| < |G|$, we can apply the induction hypothesis to G' to obtain a tree-cut decomposition $(T',\mathcal{X}')$ of G' of adhesion at most α such that the 3-centres of its torsos are almost α-thin.

If $d_G(v) \geqslant 1$, let u be a neighbour of v in G, and otherwise, let u be an arbitrary vertex in G other than v. Then there exists a (unique) node $t_u \in T'$ with $u \in X'_{t_u}$. We then obtain a tree-cut decomposition $(T, \mathcal{X})$ of G from $(T', \mathcal{X}')$ in that T arises from T' by adding a vertex t_v adjacent to t_u and in that we set $X_{t_v} := \{v\}$ while all other parts remain unchanged. We claim that $(T, \mathcal{X})$ is as desired.

We first prove that $(T, \mathcal{X})$ has adhesion at most α. The edge $t_u t_v \in T$ induces an adhesion set of size at most $2 \leqslant \alpha$. Every other edge e of T is also an edge of T', and the adhesion set of $(T, \mathcal{X})$ induced by e contains the same edges as the one of $(T', \mathcal{X}')$ induced by e – except that if v has a second neighbour $w \neq u$ in G, then the edge of G' that arose by suppressing v is replaced by the edge joining v and w in G. This replacement, however, leaves the size of the adhesion set unchanged. Hence, the adhesion of $(T, \mathcal{X})$ is at most α since $(T', \mathcal{X}')$ has adhesion at most α.

It remains to show that the 3-centres of the torsos of $(T, \mathcal{X})$ are almost α-thin. Since v has degree at most 2 in G, the 3-centre of the torso of $(T, \mathcal{X})$ at t_v consists only of the vertex v and is hence almost α-thin. For all other nodes $t \in T$, the 3-centre of the torso of $(T, \mathcal{X})$ at t equals the 3-centre of the torso of $(T', \mathcal{X}')$ at t, since the new peripheral vertex v of the torso of $(T, \mathcal{X})$ at $t = t_u$ has degree at most 2 and is hence deleted or suppressed in the construction of the 3-centre from the torso. □

Lemma 11.4.2. *If Theorem 45 holds for all 3-edge-connected graphs, then it holds for all graphs.*

Proof. Let $\alpha := \alpha(\ell) \geqslant 2$ be an integer such that Theorem 45 holds for all 3-edge-connected graphs. By Lemma 11.4.1, it is enough to prove the statement for all graphs G of minimum degree 3, and we will do so by induction on $|G|$. The case $|G| = 1$ is trivial; so suppose that $|G| > 1$, that G has minimum degree 3, that G is not 3-edge-connected, and that G does not contain the wall W_ℓ as a strong immersion minor.

Let $\{A, B\}$ be a bipartition of $V(G)$ such that the size of the cut $E_G(A, B)$ of G is minimal among all bipartitions of $V(G)$. Since G is not 3-edge-connected, the cut $E_G(A, B)$ has size at most 2. Hence, both A and B have size at least 2, since G has minimum degree 3. Let G^A arise from G by contracting B to a single vertex b, keeping parallel edges as they arise, but omitting loops; analogously define G^B by contracting A to a. Then both G^A and G^B are strictly smaller than G, and they do not contain W_ℓ as a strong immersion minor, since G does not do so. By applying the induction hypothesis, we obtain tree-cut decompositions $(T^A, \mathcal{X}^A)$ of G^A and $(T^B, \mathcal{X}^B)$ of G^B which both have adhesion at most α and such that the 3-centres of their torsos are almost α-thin.

Using $(T^A, \mathcal{X}^A)$ and $(T^B, \mathcal{X}^B)$ we now construct a tree-cut decomposition $(T, \mathcal{X})$ of G which, as we will show afterwards, has the desired properties.[2] Let $t^A \in T^A$ and $t^B \in T^B$ be the (unique) nodes with $b \in X^A_{t^A}$ and $a \in X^B_{t^B}$. We then define the tree T as the disjoint union of T^A and T^B together with an edge joining t^A and t^B. We set $X_{t^A} := X^A_{t^A} \setminus \{b\}$ and $X_{t^B} := X^B_{t^B} \setminus \{a\}$, and take X_t as the corresponding X^A_t or X^B_t for all other nodes $t \in T$.

[2]This construction follows the one in terms of edge-sums given in [Wol15, Lemma 5].

Let us first check that $(T, \mathcal{X})$ has adhesion at most α. By construction, every adhesion set of $(T, \mathcal{X})$ that is induced by an edge of T other than $t^A t^B$ has the same size as the adhesion set of $(T^A, \mathcal{X}^A)$ or $(T^B, \mathcal{X}^B)$ induced by the corresponding edge of T^A or T^B. The adhesion set of $(T, \mathcal{X})$ induced by $t^A t^B$ is precisely the cut $E_G(A, B)$ and hence of size at most $2 \leqslant \alpha$.

Next, we verify that the 3-centres of the torsos of $(T, \mathcal{X})$ are almost α-thin. For any node $t \in T$ with $t \neq t^A, t^B$, the torso of $(T, \mathcal{X})$ at t equals the corresponding torso of $(T^A, \mathcal{X}^A)$ or $(T^B, \mathcal{X}^B)$ and hence, its 3-centre is almost α-thin. The torso H of $(T, \mathcal{X})$ at t^A differs from the torso H^A of $(T^A, \mathcal{X}^A)$ at t^A only in that b is not a core vertex any more, but a peripheral vertex now. So the 3-centre of H arises from the one of H^A by deleting b if $d(b) \leqslant 1$ and suppressing b if $d(b) = 2$. But then the 3-centre of the torso H is almost α-thin since the 3-centre of H^A is, because a graph remains almost α-thin after the deletion of a vertex by definition and the suppression of a vertex of degree 2 does not affect the almost α-thinness by Lemma 11.2.4. The analysis of the torso at t^B is symmetrical. □

11.4.2 Immersions and tree-cut decompositions

In this section we collect some tools and auxiliary results for the proof of Theorem 45.

The first lemma asserts that the absence of a large wall as a strong immersion minor implies a bound on the number of neighbours of every vertex. This is the lemma for whose application in the proof of Theorem 45 we reduced to 3-edge-connected graphs in Section 11.4.1.

Lemma 11.4.3 ([DK14b], Corollary 1.3)**.** *For every integer $\ell > 0$, there is an integer $g(\ell) > 0$ such that if a 3-edge-connected graph G contains a vertex with at least $g(\ell)$ neighbours, then the wall W_ℓ is a strong immersion minor of G.*

The next two lemmas again exclude specific substructures in the absence of a large wall as a strong immersion minor. First, for every two integers $k, n > 0$, let $S_{k,n}$ be the graph on $n + 1$ vertices $x, v_1, \ldots, v_n$ with k parallel edges joining x and v_i for $i = 1, \ldots, n$. For every graph G on n vertices which has maximum degree k, one can greedily construct a strong immersion of G in $S_{k,n}$ (cf. [Wol15, Observation 1]).

Lemma 11.4.4. *For every two integers $k, n > 0$, the graph $S_{k,n}$ contains every graph on n vertices with maximum degree k as a strong immersion minor. In particular, $S_{3,2\ell^2}$ contains the wall W_ℓ as a strong immersion minor for every integer $\ell > 0$.* □

Secondly, let $P_n * v$ denote the graph arising from the path P_n of length n by adding a vertex v that is adjacent to every vertex of P_n.

Lemma 11.4.5. *For every integer $k > 0$, the graph $P_{3k-1} * v$ contains a subdivision of $S_{3,k}$. In particular, the wall W_ℓ is a strong immersion minor of $P_{6\ell^2-1} * v$ for every integer $\ell > 0$.*

Proof. Deleting every third edge on the path P_{3k-1}, the graph $P_{3k-1} * v$ becomes a subdivision of $S_{3,k}$. The observation on W_ℓ follows immediately from Lemma 11.4.4. □

The following results together show that the absence of a subdivision of the wall W_ℓ, which is implied by the absence of W_ℓ as a strong immersion, implies the existence of a tree-cut decomposition whose adhesion and torso size are both bounded in terms of the maximum degree and the size ℓ of the wall.

Theorem 11.4.6 (Grid Theorem, [RS86b], (2.1))**.** *For every integer $\ell > 0$, there exists an integer $w(\ell) > 0$ such that if a graph G does not contain a subdivision of the wall W_ℓ, then G has tree-width at most $w(\ell)$.*

Lemma 11.4.7 ([Wol15], Lemma 12)**.** *Let $w, d > 0$ be integers, and let G be a graph with maximum degree at most d and tree-width at most w. Then there exists a tree-cut decomposition of G of adhesion at most $(2w + 2)d$ such that every torso has at most $(d + 1)(w + 1)$ vertices.*

The next lemma says that a large enough set U of vertices in a connected graph is either reflected in a subdivision of a star with many leaves in U or a comb with many teeth in U. Here, a *comb* is a union of a path P with disjoint, possibly trivial, paths which have precisely their first vertex on P, and we call the last vertices of these paths their *teeth*. Note that if a graph has small maximum degree, then the graph does not contain any large subdivided star.

Lemma 11.4.8 ([BK22], Lemma 6.2)**.** *For every two integers $s, t > 0$, there exists an integer $h(s, t) > 0$ such that the following holds: Whenever a set U of vertices of a connected graph G has size at least $h(s, t)$, then there is either some subdivision of a star in G with all its s leaves in U or a comb in G with all its t teeth in U.*

As a final ingredient to our proof of Theorem 45, the following lemma asserts that if a simple graph does not contain a star as a minor, then it is 'close' to a union of vertex-disjoint paths.

Lemma 11.4.9 ([MW14], Lemma 20)**.** *Let $k > 0$ be an integer. If a simple connected graph G does not contain $K_{1,k}$ as a minor, then $G - X$ is a vertex-disjoint union of paths for some set X of at most $4k$ vertices of G.*

11.4.3 Proof of Theorem 45

Proof of Theorem 45. We will define the integer $\alpha(\ell) > 0$ explicitly in terms of ℓ in what follows. Let G be a graph in which the wall W_ℓ is not strongly immersed. By Lemma 11.4.2, we may assume that G is 3-edge-connected.

To construct the desired tree-cut decomposition $(T, \mathcal{X})$ of G, we will first apply Lemma 11.4.7 to an appropriate minor G' of G in order to obtain an auxiliary tree-cut decomposition $(T, \mathcal{X}')$ of G' from which we then construct $(T, \mathcal{X})$. This minor G' of G is defined as follows: Let A be the simple graph on $V(G)$ with an edge joining two vertices if there are more than $p(\ell) := 6\ell^2$ edges joining them in G. The graph G' then arises from G by contracting the components of A, keeping parallel edges as they arise (omitting potentially arising loops).

Now G' inherits two key properties of G, which we will need in what follows. First, G' is still 3-edge-connected as it is a minor of the 3-edge-connected graph G. Secondly, G' does still

not contain the wall W_ℓ as a strong immersion minor. Indeed, G does not contain W_ℓ as a strong immersion minor, and since $p(\ell) \geqslant 3$ and W_ℓ has maximum degree 3, one could greedily construct a strong immersion of W_ℓ in G from a strong immersion of W_ℓ in G'.

In order to apply Lemma 11.4.7 to G', we next bound its tree-width and its maximum degree in terms of ℓ.

Claim 1. *The tree-width of G' is at most $w(\ell)$, the integer given by Theorem 11.4.6.*

Proof. Since G' does not contain W_ℓ as a strong immersion minor, in particular it does not contain a subdivision of W_ℓ. Therefore, Theorem 11.4.6 yields the desired upper bound $w(\ell)$ on the tree-width of G'. ■

Claim 2. *There exists an integer $d(\ell)$ such that G' has maximum degree $d(\ell)$.*

Proof. We prove that G' has maximum degree at most $d(\ell) := g(\ell) \cdot p(\ell) \cdot h(2\ell^2, 6\ell^2)^2$, where $h(2\ell^2, 6\ell^2)$ and $g(\ell)$ are the integers given by Lemma 11.4.8 and Lemma 11.4.3, respectively. Since G' is 3-edge-connected and does not contain W_ℓ as strong immersion minor, Lemma 11.4.3 implies that every vertex of G' has less than $g(\ell)$ neighbours in G'. Thus, it is enough to show that there are at most $d(\ell)/g(\ell)$ many parallel edges joining any fixed pair of vertices. By the construction of G', this is the case if there are at most $d(\ell)/g(\ell)$ edges in G joining any two components of A.

Suppose for a contradiction that there are two components C and C' of A such that there are more than $d(\ell)/g(\ell) = p(\ell) \cdot h(2\ell^2, 6\ell^2)^2$ edges of G joining them. By the definition of A, each vertex in C is joined to each vertex in C' by at most $p(\ell)$ parallel edges in G. Thus, at least one of $N_G(C) \cap C'$ and $N_G(C') \cap C$ is of size at least $h(2\ell^2, 6\ell^2)$, and we may assume without loss of generality that it is $N_G(C') \cap C$.

Now the simple graph A has maximum degree less than $2\ell^2$. Indeed, since $p(\ell) \geqslant 3$, a vertex of degree at least $2\ell^2$ in A would immediately imply the existence of a strong immersion of W_ℓ in G by Lemma 11.4.4. Thus, by Lemma 11.4.8, $A[C]$ contains a comb with at least $6\ell^2$ teeth in $N_G(C') \cap C$.

Hence, the wall W_ℓ is strongly immersed in G by Lemma 11.4.5, since G contains $P_{6\ell^2-1} * v$ as a strong immersion minor: Fix a vertex v in C'. Since $p(\ell) \geqslant 6\ell^2$, we may pick greedily edge-disjoint paths from v to each tooth in $N_G(C') \cap C$ such that all edges of each path except its last edge are contained in $G[C']$. These paths together with the comb form a strong immersion of $P_{6\ell^2-1} * v$ in G.

All in all, the maximum degree of G' is at most $d(\ell)$, as desired. ■

By Claim 1 and Claim 2, we can now apply Lemma 11.4.7 to G' and obtain a tree-cut decomposition $(T, \mathcal{X}')$ of G' of adhesion at most $a(\ell) := (2w(\ell) + 2)d(\ell)$ such that every torso of $(T, \mathcal{X}')$ has size at most $k(\ell) := (d(\ell)+1)(w(\ell)+1)$. From the tree-cut decomposition $(T, \mathcal{X}')$ of the minor G' of G, we then construct the tree-cut decomposition $(T, \mathcal{X})$ of G by defining, for each node $t \in T$, its corresponding part X_t to be the union of the branch sets of the vertices in X'_t.

Now let

$$\alpha = \alpha(\ell) := \max\{k(\ell)(8\ell^2{+}1{+}6\ell^2{\cdot}16\ell^4), p(\ell){\cdot}(6\ell^2)^2{+}\binom{16\ell^4}{2}{\cdot}p(\ell){\cdot}(6\ell^2)^2{+}k(\ell){\cdot}d(\ell)/2{+}a(\ell){\cdot}k(\ell)\}.$$

We show that $(T, \mathcal{X})$ has the desired properties in that it has adhesion at most α and the 3-centres of its torsos are almost α-thin. In doing so, the first term in the maximum defining α will be used to bound the size of the set of vertices that we will delete from the 3-centre of a torso to make it α-thin as well as the number of its neighbours, while the second term in the maximum will be used to show that the remainder of the 3-centre is α-thin.

Claim 3. *The tree-cut decomposition $(T, \mathcal{X})$ has adhesion at most $a(\ell) \leqslant \alpha$.*

Proof. By the construction of $(T, \mathcal{X})$ from $(T, \mathcal{X}')$, the adhesion sets of $(T, \mathcal{X})$ are the same as those of $(T, \mathcal{X}')$. Thus, $(T, \mathcal{X})$ has the same adhesion as $(T, \mathcal{X}')$ and, in particular, adhesion at most $a(\ell)$. ■

Claim 4. *The 3-centres of the torsos of $(T, \mathcal{X})$ are almost α-thin.*

Proof. We first observe that the torsos of $(T, \mathcal{X})$ equal their 3-centres: Every adhesion set of $(T, \mathcal{X})$ has size at least 3, since G is 3-edge-connected. Therefore, all peripheral vertices of torsos of $(T, \mathcal{X})$ have degree at least 3 in the respective torso. This implies that the 3-centre of a torso of $(T, \mathcal{X})$ equals the torso itself.

Given any node $t \in T$, let H_t be the torso of $(T, \mathcal{X})$ at t. We first find a set X of vertices of H_t such that $A[X_t] - X$ is a disjoint union of paths, and we show that both the size of X and the size of its neighbourhood in H_t are bounded by our suitably chosen $\alpha = \alpha(\ell)$. In a next step, we then use the structure of $A[X_t] - X$ to exhibit a natural enumeration of $V(H_t - X)$, and we finally show that this enumeration witnesses the α-thinness of $H_t - X$.

For every component C of $A[X_t]$, and more generally for every component C of A, the graph $A[C]$ does not contain the star $K_{1,2\ell^2}$ as a minor. To prove this, we show that $G[C]$ otherwise contains W_ℓ as a strong immersion minor, contradicting our assumptions on G. By Lemma 11.4.4, it is enough to show that $S_{3,2\ell^2}$ is strongly immersed in $G[C]$. So let $U \subseteq C$ be a set of vertices containing precisely one vertex of each branch set of the $K_{1,2\ell^2}$-minor of $A[C]$. We then map the centre of $S_{3,2\ell^2}$ to the vertex of U in the branch set of the centre of $K_{1,2\ell^2}$, and map the rest of the vertices of $S_{3,2\ell^2}$ bijectively to the remaining ones in U. Since $p(\ell) \geqslant 6\ell^2$, we can now greedily embed the edges of $S_{3,2\ell^2}$ as edge-disjoint paths in $G[C]$ joining the vertices in U along the $K_{1,2\ell^2}$-minor. Hence, we have found our desired strong immersion of $S_{3,2\ell^2}$ in $G[C]$.

We are now ready to find the desired set X of vertices of the torso H_t of $(T, \mathcal{X})$ at t. For every component C of $A[X_t]$, the graph $A[C]$ does not contain $K_{1,2\ell^2}$ as a minor. So Lemma 11.4.9 yields a set $X_C \subseteq C$ of size at most $8\ell^2$ such that $A[C] - X_C$ is a disjoint union of paths. We then let the set X be the union of these X_C. The construction of X then implies that $A[X_t] - X$ is indeed a disjoint union of paths.

We first show that the size of X is bounded by α. Note that since A has maximum degree at most $2\ell^2$ (as in the proof of Claim 2), the number of components of $A[C] - X_C$ is at most $|X_C| \cdot 2\ell^2 \leqslant 16\ell^4$. Since X'_t has size at most $k(\ell)$, the graph $A[X_t]$ has at most $k(\ell)$ components and thus $|X| \leqslant 8\ell^2 \cdot k(\ell) \leqslant \alpha$.

Next, we claim that each vertex x in X has at most $8\ell^2 \cdot k(\ell) + k(\ell) + k(\ell) \cdot 6\ell^2 \cdot 16\ell^4 \leqslant \alpha$ neighbours in H_t. Indeed, suppose for a contradiction that x has more than $k(\ell)+k(\ell)\cdot 6\ell^2 \cdot 16\ell^4$ neighbours in $H_t - X$. We first note that x is adjacent to at most $k(\ell)$ peripheral vertices of H_t, since there are at most $k(\ell)$ many peripheral vertices of H_t: They are in one-to-one correspondence to the peripheral vertices of the torso of $(T, \mathcal{X}')$ at t which has size at most $k(\ell)$. Thus, x has more than $k(\ell) \cdot 6\ell^2 \cdot 16\ell^4$ many neighbours which are core vertices of H_t and not contained in X. By the pigeonhole principle, there then exists one component C of the at most $k(\ell)$ components of $A[X_t]$ such that x has at least $6\ell^2 \cdot 16\ell^4$ neighbours in $C \setminus X_C$. Applying the pigeonhole principle a second time, we find one component P of the at most $16\ell^4$ components of $A[C] - X_C$ that contains at least $6\ell^2$ neighbours of x. Since $A[P]$ is a path by the choice of X_C, we have hence found a subdivision and thus a strong immersion of $P_{6\ell^2-1} * v$ in H_t. As all the vertices in C as well as x are core vertices of the torso H_t, this also is a strong immersion of $P_{6\ell^2-1} * v$ in G. So by Lemma 11.4.5, the wall W_ℓ is contained in G as a strong immersion minor – a contradiction.

Since X has size at most α and each element of X has at most α neighbours in H_t, it remains to prove that $H_t - X$ is α-thin. To do so, we now give an enumeration of $V(H_t - X)$ which we will then show to witness the α-thinness of $H_t - X$.

Let $c_1, \dots, c_n$ be an arbitrary enumeration of the part X'_t of the tree-cut decomposition $(T, \mathcal{X}')$ of G', let C_i be the component of A corresponding to the vertex c_i of G', and write $X_i := X_{C_i}$. Then take an enumeration $v_1, \dots, v_N$ of $X_t \setminus X$ such that for $i = 1, \dots, n$, the set $C_i - X_i$ forms an interval of this enumeration and within such an interval, every component of $A[C_i] - X_i$ forms an interval with the same linear order which is given by the path that it induces in A by the choice of X_i. For the peripheral vertices of H_t, we finally pick an arbitrary enumeration $v_{N+1}, \dots, v_M$.

To show that this enumeration $v_1, \dots, v_M$ of $V(H_t - X)$ is as desired, first pick an arbitrary $i \in \{1, \dots, N\}$. Then there is a component C_j of $A[X_t]$ with $v_i \in C_j$, and a component of $A[C_j] - X_j$ containing v_i; this component of $A[C_j] - X_j$ then induces a path P in A by the choice of X_j (which might have length 0). Every edge in H_t joining $\{v_1, \dots, v_{i-1}\}$ and $\{v_{i+1}, \dots, v_M\}$ belongs to one of the following classes:

(1) edges joining vertices of the path P,

(2) edges joining the distinct components in $A[C_j] - X_j$,

(3) edges joining distinct components of $A[X'_t]$, and

(4) edges incident to peripheral vertices of H_t.

We will now separately bound the size of each of these classes (1) to (4) in terms of ℓ. Note

that only (1) depends on the choice of the enumeration; in particular, we will indeed bound the absolute number of all such edges in H_t in terms of ℓ for (2) to (4).

For (1), the edges joining vertices of P, note that $V(P)$ is a component of $A[C_j] - X_j$. So there are no edges in A joining the vertices of P except the ones contained in P itself. Hence, each two non-adjacent vertices of P are joined by at most $p(\ell)$ parallel edges. Let P_1 and P_2 be the two components of $P - v_i$, and note that both P_1 and P_2 consist of vertices of G since all vertices on P are in X_t. If $|E_G(P_1, P_2)| > p(\ell) \cdot (6\ell^2)^2$, then one of P_i, say P_1, has more than $6\ell^2$ neighbours in P_{2-i}. Fix any $v \in P_1$. Since $p(\ell) \geqslant 6\ell^2$, we can greedily find edge-disjoint paths in $G[P_1 \cup P_2]$ joining v and $6\ell^2$ of these neighbours in P_2, and we thus find a strong immersion of $P_{6\ell^2-1} * v$ in $G[P_1 \cup P_2]$. By Lemma 11.4.5, this graph contains the wall W_ℓ as a strong immersion minor which yields the desired contradiction.

With the same type of argument, one can show that there are at most $p(\ell) \cdot (6\ell^2)^2$ edges in G joining any two components of $A[C_j] - X_j$. Recall that the number of components of $A[C_j] - X_j$ is at most $16\ell^4$. Thus, there are at most $\binom{16\ell^4}{2} \cdot p(\ell) \cdot (6\ell^2)^2$ edges in G as in (2), i.e. edges joining components of $A[C_j] - X_j$.

Next, we turn to (3) and bound the number of edges in G joining distinct components of $A[X'_t]$. Since G' arises from G by contracting the components of A, these edges are in one-to-one correspondence with the edges in $G'[X'_t]$. Since X'_t has size at most $k(\ell)$ and G' has maximum degree at most $d(\ell)$, there are at most $k(\ell) \cdot d(\ell)/2$ edges in $G'[X'_t]$ and hence joining distinct components of $A[X'_t]$.

Along similar lines, we can bound the number of edges in (4), i.e. the edges incident to peripheral vertices of H_t, in terms of ℓ. Indeed, the edges incident to peripheral vertices of H_t are in one-to-one correspondence to the edges incident to the peripheral vertices of the torso of $(T, \mathcal{X}')$ at t. Since $(T, \mathcal{X}')$ has adhesion at most $a(\ell)$, each peripheral vertex has degree at most $a(\ell)$. Moreover, since X'_t has size at most $k(\ell)$, there are at most $k(\ell)$ peripheral vertices of H_t. Altogether, there are at most $a(\ell) \cdot k(\ell)$ edges incident to peripheral vertices of H_t.

All the above bounds show that by the choice of $\alpha = \alpha(\ell)$, there are at most α edges joining $\{v_1, \dots, v_{i-1}\}$ and $\{v_{i+1}, \dots, v_M\}$ in H_t for $i = 1, \dots, N$. By the choice of our enumeration, the vertices $v_{N+1}, \dots, v_M$ are precisely the peripheral vertices of H_t. So for $i = N+1, \dots, M$, all edges joining $\{v_1, \dots, v_{i-1}\}$ and $\{v_{i+1}, \dots, v_M\}$ in H_t are incident to peripheral vertices of H_t. As shown above there are at most $a(\ell) \cdot k(\ell) \leqslant \alpha(\ell)$ such edges. This completes the proof of this claim. ■

By Claim 3 and Claim 4, our constructed tree-cut decomposition $(T, \mathcal{X})$ of G now is as desired, completing the proof of Theorem 45. □

CHAPTER 12

The Lovász-Cherkassky theorem for locally finite graphs with ends

This chapter is based on [JJKKM23], which is joint work with Raphael W. Jacobs, Attila Joó, Jan Kurkofka, Ruben Melcher.

12.1 Introduction

A non-trivial path P is a *T-path* for a set T of vertices if P has its endvertices but no inner vertex in T. For disjoint vertex sets X and Y in a graph G, we write $\lambda_G(X, Y)$ for the size of a smallest cut in G that separates X and Y.

Now let T be any set of vertices in a finite graph G. In a set $\mathcal{P}$ of edge-disjoint T-paths, there are at most $\lambda_G(t, T \setminus \{t\})$ many paths that link a vertex $t \in T$ to $T \setminus \{t\}$. It follows that $|\mathcal{P}| \leqslant \frac{1}{2} \sum_{t \in T} \lambda_G(t, T \setminus \{t\})$. The question about the sharpness of this upper bound can be formulated in the following structural way. For every vertex $t \in T$, let $\mathcal{P}_t$ be a set of $\lambda_G(t, T \setminus \{t\})$ many edge-disjoint t–$(T \setminus \{t\})$ paths.

Question. Can we choose the paths in the sets $\mathcal{P}_t$ for each vertex $t \in T$ in such a way that the union of all the sets $\mathcal{P}_t$ is an edge-disjoint path-system?

Clearly, the answer is no: the three leaves of a star $K_{1,3}$ form a set T where each set $\mathcal{P}_t$ must consist of a single path, but the union $\bigcup_{t \in T} \mathcal{P}_t$ always contains two distinct paths that share an edge, no matter how we choose the paths in each set $\mathcal{P}_t$. Lovász and Cherkassky independently showed that, perhaps surprisingly, the answer is yes under the additional assumption that the graph G is *inner-Eulerian* for T in that every vertex of G which is not in T has even degree in G.

Theorem 12.1.1 (Lovász-Cherkassky Theorem [Che77; Lov76])**.** *Let G be any finite graph, and let $T \subseteq V(G)$ such that G is inner-Eulerian for T. Then the maximum number of pairwise edge-disjoint T-paths in G is equal to*

$$\frac{1}{2} \sum_{t \in T} \lambda_G(t, T \setminus \{t\}).$$

Diestel asked whether Theorem 12.1.1 extends to all locally-finite infinite graphs and their ends [Die20]. In this note, we answer his question in the affirmative. For this, we employ the Freudenthal compactification in the spirit of [Die10; Die11; DS12], as customary in the study of locally-finite infinite graphs [BDS05; BS07; Die03; DK04a; DK04b; DK04c]. For infinite graphs and their ends, we follow and assume familiarity with the terminology in [Die25, §8], in particular in §8.6.

We allow a graph to have parallel edges, but we do not allow any loops; if a graph has no parallel edges, we call it *simple*. The *degree* of a vertex v in a graph G is the number $d_G(v) \in \mathbb{N} \cup \{\infty\}$ of edges of G incident with v. If all the vertices of G have finite degree, then we say that G is *locally finite*. Note that in a locally finite graph, there can be only finitely many parallel edges between any two vertices. We write $\hat{V}(G)$ for the union of the vertex set $V(G)$ of G and the set $\Omega(G)$ of all ends of G. An arc A in the end compactification $|G|$ of a locally-finite connected graph G is a *T-arc* for a set $T \subseteq \hat{V}(G)$ if A has its endpoints but no inner points in T. An arc A is an *X–Y arc* between two sets X and Y if A intersects X precisely in one endpoint and Y precisely in the other. We call an arc *graphic* if it is defined by a finite graph-theoretic path, a ray, or a double ray with its tails in distinct ends. Two arcs in $|G|$ are *edge-disjoint* if they do not meet in inner points of edges.

A subset $X \subseteq \hat{V}(G)$ *lives in* a subgraph $C \subseteq G$ or a vertex set $C \subseteq V(G)$ if all the vertices of X lie in C and all the rays of ends in X have tails in C or $G[C]$, respectively. A finite cut F of a graph G is an *X–Y cut* for two sets $X, Y \subseteq \hat{V}(G)$ if X and Y live in distinct sides of F. If $Y \subseteq \hat{V}(G)$ is a set of vertices and ends of a locally finite graph G, and $x \in \hat{V}(G)$ is not contained in the closure of Y in $|G|$, then G admits a (finite) x–Y cut, and we denote the least size of such a cut by $\lambda_G(x, Y)$.

Our main result reads as follows:

Theorem 47. *Let G be any locally-finite graph, and let $T \subseteq \hat{V}(G)$ be discrete in $|G|$. If every finite cut of G such that T lives on one of its sides is even, then $|G|$ contains a set $\mathcal{A}$ of pairwise edge-disjoint graphic T-arcs such that for every $t \in T$, the number of t–$(T \setminus \{t\})$ arcs in $\mathcal{A}$ is equal to $\lambda_G(t, T \setminus \{t\})$.*

On first sight, one might wonder why the notion of 'inner-Eulerian' from the premise of Theorem 12.1.1 has been replaced in the premise of Theorem 47 with a new condition on the vertex sets X. In short, in finite graphs 'inner-Eulerian' implies the new premise, so Theorem 47 is more general; see Section 12.3 for a discussion of the necessity of the new condition for end-compactifications of locally finite graphs.

The assumption that T is a discrete subset of $|G|$ naturally arises here as it is equivalent to asking that there exists a t–$(T \setminus \{t\})$ cut in G for each end $t \in T$, which precisely ensures that $\lambda_G(t, T \setminus \{t\})$ is defined for all $t \in T$. We remark that the assumption that T is discrete is also motivated by the work of Bruhn, Diestel, and Stein [BDS05], which is a generalisation of the Erdős-Menger theorem by Aharoni and Berger [AB09] from infinite graphs to infinite graphs and their ends, under a similar assumption on the ends which implies discreteness in our setting. (Diestel discusses this assumption in detail in [Die03, §3].)

We conclude the introduction with an example that discusses why a natural weakening of the discreteness-assumption in Theorem 47 cannot be made.

Example 12.1.2. We claim that it is not possible to drop in Theorem 47 the requirement that T is discrete in $|G|$ and, instead, replace $\lambda_G(t, T \setminus \{t\})$ in the wording of the theorem with the maximum number $\mu_G(t, T \setminus \{t\})$ of pairwise edge-disjoint graphic t–$(T \setminus \{t\})$ arcs in $|G|$.

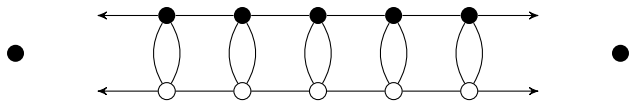

Figure 12.1: T consists of the black vertices and ends

Indeed, let G be obtained from the double ladder by duplicating each rung, and let T consist of both ends of G together with the vertices of one of the two main double rays; see Figure 12.1. Then

$$\mu_G(t, T \setminus \{t\}) = \begin{cases} 4 & \text{if } t \in T \cap V(G), \\ 1 & \text{if } t \in T \cap \Omega(G). \end{cases}$$

But any t–$(T \setminus \{t\})$ arc for some $t \in T \cap \Omega(G)$ does already preclude the existence of four t'–$(T \setminus \{t'\})$ arcs for some of the $t' \in T \cap V(G)$, and hence there is no desired arc-system for T.

12.1.1 The Lovász-Cherkassky Theorem for infinite graphs without ends

Joó has extended Theorem 12.1.1 to infinite – not necessarily locally finite – graphs where the set T is a countably infinite set of vertices, but does not contain ends. A cut F is said to lie *on* a set $\mathcal{P}$ of edge-disjoint paths in G if F consists of a choice of exactly one edge from each path in $\mathcal{P}$. For a vertex set $X \subseteq V(G)$, we denote by $d_G(X)$ the number of edges of G between X and its complement $V(G) \setminus X$. Note that this notation is consistent with the above definition of the degree $d_G(v)$ of a vertex v of G in that $d_G(v) = d_G(\{v\})$ for every vertex $v \in G$.

Theorem 12.1.3 ([Joó23], Theorem 1.3)**.** *Let G be any graph, and let $T \subseteq V(G)$ be a countable vertex set such that there is no $X \subseteq V(G) \setminus T$ for which $d_G(X)$ is an odd natural number. Then G contains a set $\mathcal{P}$ of edge-disjoint T-paths such that for each vertex $t \in T$, the graph G contains a t–$(T \setminus \{t\})$ cut on the set of t–$(T \setminus \{t\})$ paths in $\mathcal{P}$.*

We will use Theorem 12.1.3 in the proof of Theorem 47.

Theorem 12.1.3 compares to Theorem 47 as follows. On the one hand, Theorem 12.1.3 can be applied in the setting of countable graphs that are not locally finite, where Theorem 47 cannot be applied. In the setting of locally finite graphs, on the other hand, Theorem 47 is more general than Theorem 12.1.3, as it allows T to consist of vertices and ends alike. Indeed, the locally-finite version of Theorem 12.1.3 can easily be re-obtained from Theorem 47 since every set of vertices of a locally finite graph G is discrete in $|G|$ and does not contain any other vertices in its closure.

12.2 Proof of the main result of this chapter

We need one auxiliary result for the proof Theorem 47, and to state this lemma we make the following definition. If S is a set of vertices of a graph G and ω is an end of G, then by an S–ω *ray* we mean a ray which has precisely its first vertex in S and belongs to ω.

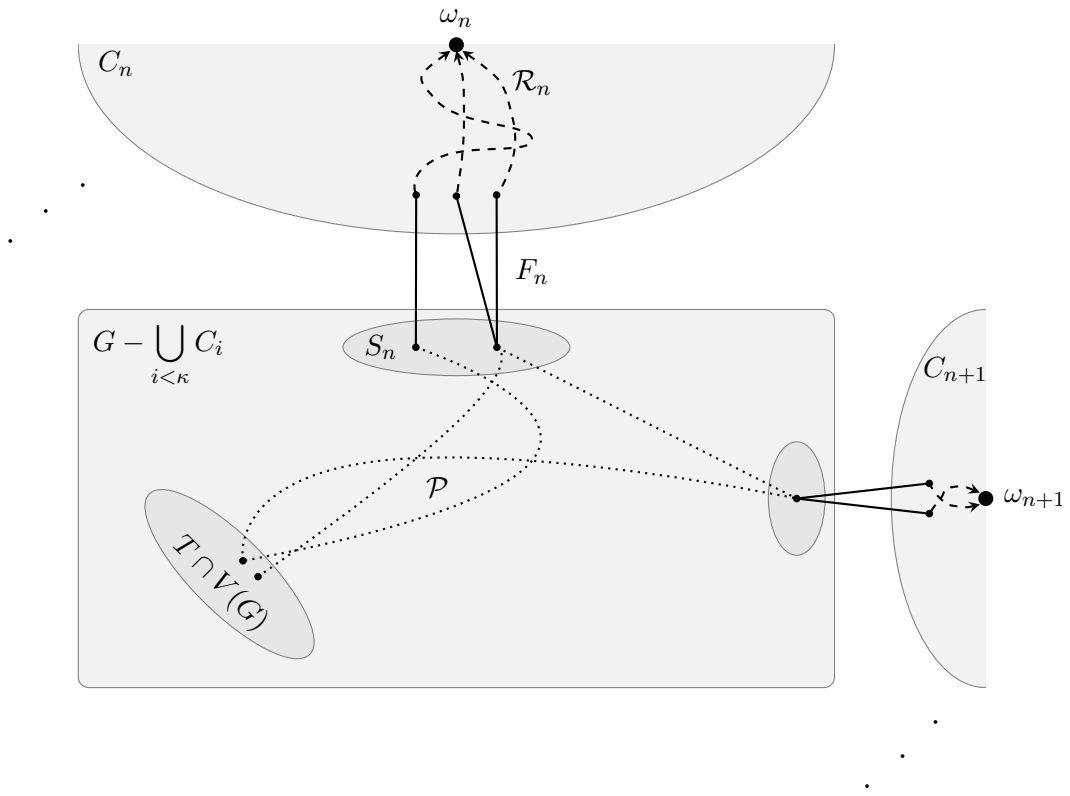

Figure 12.2: Situation in the proof of the main result

Lemma 12.2.1 ([BS07], Lemma 10)**.** *Let G be a locally finite connected graph, let ω be an end of G, and let S be a finite set of vertices in G. Then the maximum number of edge-disjoint S–ω rays is equal to the minimum size of a cut that separates S and ω.*

Originally, Lemma 12.2.1 has been proved only for simple graphs. However, it extends to graphs with parallel edges: Given a graph with parallel edges, just subdivide each edge once, apply the original result, and then suppress all subdividing vertices.

Proof of Theorem 47. Without loss of generality, we may assume that G is connected.

Let us first show that T is countable. Since G is locally finite, $|G|$ is second-countable, meaning that the topology on $|G|$ has some countable base $\mathcal{U}$. Recall that by assumption, T is discrete in $|G|$. Therefore, we find for each $t \in T$ a basic neighbourhood $U_t \in \mathcal{U}$ so that U_t and $T \setminus \{t\}$ are disjoint; in particular, $U_t \neq U_{t'}$ for distinct $t, t' \in T$. Then T is countable because $\mathcal{U} \supseteq \{U_t \colon t \in T\}$ is countable.

Since $T \cap \Omega(G)$ is countable, we may fix an enumeration $(\omega_n \colon n < \kappa)$ of $T \cap \Omega(G)$ where $\kappa := |T \cap \Omega(G)| \leqslant \aleph_0$. Next, we recursively find for each $n < \kappa$ an ω_n–$(T \setminus \{\omega_n\})$ cut F_n of G such that the component C_n of $G - F_n$ in which ω_n lives is disjoint from the component C_m of $G - F_m$ in which ω_m lives for all $m < \kappa$ other than n (see also Figure 12.2 for a visualisation of the whole proof).

Given any $n < \kappa$, assume that we have already found suitable finite cuts F_i for all $i < n$. Let G_n be the graph obtained from G by contracting each component C_i for $i < n$ to a single

vertex v_i, keeping all the parallel edges that may arise. Since all cuts F_i for $i < n$ are finite, the contraction minor G_n is again locally finite. Let

$$T_n := (T \setminus \{\omega_i \colon i < n\}) \cup \{v_i \colon i < n\} \subseteq \hat{V}(G_n).$$

Note that $|G_n| = |G| \,/\, \left\{\overline{C_i} \colon i < n\right\}$, since all the F_i are finite.

Since T is discrete in $|G|$ and the components C_i for $i < n$ are disjoint, T_n is discrete in $|G_n|$. Therefore, the end ω_n is not contained in $\overline{T_n \setminus \{\omega_n\}}$ where the closure is taken in $|G_n|$. Thus, there is a smallest ω_n–$(T_n \setminus \{\omega_n\})$ cut F_n^* in G_n. This finite cut F_n^* in G_n defines a finite cut F_n of G of the same size.

Finally, we observe that the component $C_n \subseteq G_n$ of $G_n - F_n^*$ in which ω_n lives does not contain any v_i with $i < n$, since F_n^* separates ω_n and $T_n \setminus \{\omega_n\} \supseteq \{v_i \colon i < n\}$. Thus, C_n is also a component of $G - F_n$ and disjoint from each previous C_i. Altogether, this shows that the cut F_n is as desired.

Next, we simultaneously contract each component C_n for $n < \kappa$ to a single vertex v_n, again keeping all the parallel edges that may arise, and obtain a contraction minor G_κ of G. As before, this contraction minor G_κ is locally finite since all the cuts F_n are finite. Let

$$T_\kappa := (T \setminus \Omega(G)) \cup \{v_n \colon n < \kappa\},$$

and note that $T_\kappa \subseteq V(G_\kappa)$. Moreover, T_κ is countable because it has the same size as T.

We show that there is no vertex set $X \subseteq V(G)$ whose closure in $|G|$ is disjoint from T and for which $d_G(X)$ is an odd natural number: Let X be a set of vertices of G such that its closure in $|G|$ is disjoint from T and $d_G(X)$ is finite. It follows from the Jumping Arc Lemma [Die25, Proposition 8.6.3 (i)] that T lives in $V(G) \setminus X$. Hence, $d_G(X) = |E_G(X, V(G) \setminus X)|$ is even by assumption.

This yields that there is no vertex set $X \subseteq V(G_\kappa)$ whose closure in $|G_\kappa|$ is disjoint from T_κ and for which $d_{G_\kappa}(X)$ is an odd natural number. Since T_κ contains no ends, it further follows that there is no vertex set $X \subseteq V(G_\kappa) \setminus T_\kappa$ for which $d_{G_\kappa}(X)$ is an odd natural number. Therefore, we can apply Theorem 12.1.3 in G_κ to T_κ to obtain a set $\mathcal{P}$ of edge-disjoint T_κ-paths in G_κ with the following property: For every vertex $t' \in T_\kappa$, there is a cut $F'_{t'}$ of G on the set of t'–$(T_\kappa \setminus \{t'\})$ paths in $\mathcal{P}$. Note that all cuts $F'_{t'}$ are finite, because G_κ is locally finite and the paths in $\mathcal{P}$ are edge-disjoint. It remains to translate the set $\mathcal{P}$ of edge-disjoint T_κ-paths in G_κ into the desired set $\mathcal{A}$ of pairwise edge-disjoint graphic T-arcs in $|G|$.

For every $n < \kappa$, the finite ω_n–$(T \setminus \{\omega_n\})$ cut F_n has smallest size. Let S_n be the set of those endvertices of F_n in $G - C_n$. By definition, F_n is a minimal S_n–ω_n cut in G. So we can apply Lemma 12.2.1 to S_n and ω_n to find a set $\mathscr{R}_n$ of $|F_n|$ many edge-disjoint S_n–ω_n rays. Note that the rays in $\mathscr{R}_n$ use only vertices of $V(C_n) \cup S_n$ since F_n is a cut. In particular, there is for each edge in F_n precisely one ray in $\mathscr{R}_n$ which starts in this edge.

Now every T_κ-path $P \in \mathcal{P}$ in G_κ uniquely defines a graphic T-arc in $|G|$ as follows: If the first or last edge e of P runs between a contraction vertex v_n and another vertex u, then we replace it with the unique u–ω_n ray in $\mathscr{R}_n$ which traverses the inner points of e and add the

end ω_n to it. Here we allow one exception: if P consists of just one edge e between two vertices v_n and v_m in T_κ, then we replace e with the double ray which arises from the two rays in $\mathscr{R}_n$ and $\mathscr{R}_m$ that start in e, and add the ends ω_n and ω_m to it. In either case, let us write $A(P)$ for the graphic arc defined by P in this way. We claim that $\mathcal{A} := \{A(P) \mid P \in \mathcal{P}\}$ is the desired set of graphic T-arcs.

The arcs in $\mathcal{A}$ are edge-disjoint because the paths in $\mathcal{P}$ are edge-disjoint, the components C_n are disjoint, and the rays in each set $\mathscr{R}_n$ are edge-disjoint. It remains to show that for each $t \in T$, the number of arcs in $\mathcal{A}$ that link t to $T \setminus \{t\}$ is equal to $\lambda_G(t, T \setminus \{t\})$. Given any $t \in T$, let $t' \in T_\kappa$ be equal to t if t is a vertex, and let $t' := v_n$ if t is an end ω_n. The finite t'–$(T_\kappa \setminus \{t'\})$ cut $F'_{t'}$ of G_κ witnesses that the set of t'–$(T_\kappa \setminus \{t'\})$ paths in $\mathcal{P}$ form a maximal sized edge-disjoint t'–$(T_\kappa \setminus \{t'\})$ path system in G_κ. Since the t–$(T \setminus \{t\})$ arc system defined by $\mathcal{A}$ arises from this path system by replacing each path P with the arc $A(P)$, the number of arcs in this system is equal to the size of the finite cut $F'_{t'}$. As the t'–$(T_\kappa \setminus \{t'\})$ cut $F'_{t'}$ of G_κ induces a t–$(T \setminus \{t\})$ cut of G of the same size, we have $\lambda_G(t, T \setminus \{t\}) \leqslant |F'_{t'}|$, which completes the proof. □

12.3 A degree condition for vertices and ends

We recall that an end ω of a locally finite graph G is *even*, or *has even degree*, if there is a finite vertex set $S \subseteq V(G)$ such that for every finite set $S' \supseteq S$ of vertices, the maximum number of edge-disjoint S'–ω rays is even; otherwise, ω is *odd* or *has odd degree*. We refer to [BS07, §3] for a discussion of the parity of ends. With the parity of ends at hand, it seems natural to extend the notion of *inner-Eulerian* from finite to locally-finite infinite graphs and their ends, as follows. A locally finite graph G is *inner-Eulerian* for a set $T \subseteq \hat{V}(G)$ if every vertex and every end in $\hat{V}(G) \setminus T$ have even degree in G.

This raises the following three questions:

(1) Is it necessary to consider ends in this generalisation of 'inner-Eulerian'?

(2) Why does the premise of Theorem 47 not use this notion of 'inner-Eulerian'?

(3) How does the the notion of 'inner-Eulerian' compare to the premise of Theorem 47?

The first question is answered by the following example.

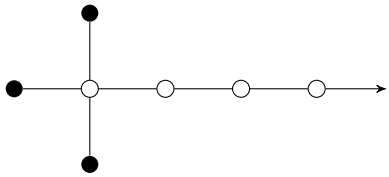

Figure 12.3: T consists of the black vertices

Example 12.3.1. Theorem 47 becomes wrong if we replace its premise on the vertex sets X with an 'inner-Eulerian' condition that ignores end degrees. Indeed, let us consider the tree G

in Figure 12.3, and let T consist of its leaves. All the non-leaves of G have even degree, but the unique end of G has degree one. Any set $\mathcal{A}$ of pairwise edge-disjoint T-arcs has size at most one, so $\mathcal{A}$ contains no t–$(T \setminus \{t\})$ arc for at least one $t \in T$, even though $\lambda_G(t, T \setminus \{t\}) = 1$.

We answer the second question twofold: On the one hand, with the premise of Theorem 47 it is clear that Theorem 47 implies Theorem 12.1.3 for locally-finite graphs. On the other hand, the premise of G being inner-Eulerian for $T \subseteq \hat{V}(G)$ is more specific than the premise on the finite cuts in Theorem 47, see Lemma 12.3.3 below. The latter also answers the third question.

Lemma 12.3.2 (Infinite Handshaking Lemma [BS07, Proposition 15])**.** *The number of odd vertices and ends in a locally finite graph is even or infinite.*

Lemma 12.3.3. *If G is a locally finite graph which is inner-Eulerian for $T \subseteq \hat{V}(G)$, then every finite cut of G such that T lives on one of its sides is even.*

Proof. Assume for a contradiction that there is an odd finite cut $E_G(A, B)$ such that T lives in A. Consider the graph H that arises from G by contracting $G - B$ to a single vertex v, keeping parallel edges. Then v has odd degree in H, but no other vertex or end of H has odd degree, which contradicts Lemma 12.3.2. □

The converse of Lemma 12.3.3 fails if G is a ray and $T = V(G)$. Hence Theorem 47 is more general than its corollary below (which follows with Lemma 12.3.3):

Corollary 12.3.4. *Let G be any locally-finite graph, and let $T \subseteq \hat{V}(G)$ be discrete in $|G|$ such that G is inner-Eulerian for T. Then $|G|$ contains a set $\mathcal{A}$ of pairwise edge-disjoint graphic T-arcs such that for every $t \in T$, the number of t–$(T \setminus \{t\})$ arcs in $\mathcal{A}$ is equal to $\lambda_G(t, T \setminus \{t\})$.* □

Bibliography

[ADK] T. Abrishami, R. Diestel, and P. Knappe. “A new width parameter that exhibits local-global structure”. In preparation.

[AJKK25] T. Abrishami, R. W. Jacobs, P. Knappe, and J. Kobler. “Locally chordal graphs”. *arXiv preprint* (2025). arXiv: 2501.17320.

[AL03] F. N. Abu-Khzam and M. A. Langston. “Graph coloring and the immersion order”. *International Computing and Combinatorics Conference*. 2003, pp. 394–403.

[AB09] R. Aharoni and E. Berger. “Menger’s theorem for infinite graphs”. *Inventiones mathematicae* 176.1 (2009), pp. 1–62. arXiv: 0509397.

[Alb23] S. Albrechtsen. “Optimal trees of tangles: refining the essential parts”. *arXiv preprint* (2023). arXiv: 2304.12078.

[Alb24a] S. Albrechtsen. “Refining tree-decompositions so that they display the k-Blocks”. *arXiv preprint* (2024). arXiv: 2403.19585.

[Alb24b] S. Albrechtsen. “Refining trees of tangles in abstract separation systems: inessential parts”. *Combinatorial Theory* 4.1 (2024), #17. arXiv: 2302.01808.

[Alb24c] S. Albrechtsen. “Tangle-tree duality in infinite graphs”. *arXiv preprint* (2024). arXiv: 2405.06756.

[ABJKW24] S. Albrechtsen, H. v. Bergen, R. W. Jacobs, P. Knappe, and P. Wollan. “On vertex sets inducing tangles”. *arXiv preprint* (2024). arXiv: 2411.13656.

[ADEFJKW23] 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). arXiv: 2307.08497.

[AHJKW24] 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* 38.2 (2024), pp. 1438–1450. arXiv: 2305.04721.

[AJKP24a] S. Albrechtsen, R. W. Jacobs, P. Knappe, and M. Pitz. “Linked tree-decompositions into finite parts”. *arXiv preprint* (2024). arXiv: 2405.06753.

[AJKP24b] S. Albrechtsen, R. W. Jacobs, P. Knappe, and M. Pitz. “Counterexamples regarding linked and lean tree-decompositions of infinite graphs”. *arXiv preprint* (2024). arXiv: 2405.06755.

[AJKW24] 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). arXiv: 2408.15335.

[AL02] A. Amit and N. Linial. “Random graph coverings I: General theory and connectivity”. *Combinatorica* 22 (2002), pp. 1–18.

[BD08] G. Bell and A. Dranishnikov. “Asymptotic dimension”. *Topology and its Applications* 155 (2008), pp. 1265–1296. arXiv: `math/0703766`.

[BD02] P. Bellenbaum and R. Diestel. “Two short proofs concerning tree-decompositions”. *Combinatorics, Probability and Computing* 11.6 (2002), pp. 541–547.

[BFGR17] R. Belmonte, F. V. Fomin, P. A. Golovach, and M. Ramanujan. “Metric dimension of bounded tree-length graphs”. *SIAM Journal on Discrete Mathematics* 31.2 (2017), pp. 1217–1243. arXiv: `1602.02610`.

[BR22] O. Bendele and D. Rautenbach. “Additive tree $O(\varrho \log n)$-spanners from tree breadth ϱ”. *Theoretical Computer Science* 914 (2022), pp. 39–46. arXiv: `2002.12103`.

[BS24] E. Berger and P. D. Seymour. “Bounded-diameter tree-decompositions”. *Combinatorica* 44.3 (2024), pp. 659–674. arXiv: `2306.13282`.

[Bie91] D. Bienstock. “On the complexity of testing for odd holes and induced odd paths”. *Discrete Mathematics* 90.1 (1991), pp. 85–92.

[Bie92] D. Bienstock. “Corrigendum: “On the complexity of testing for odd holes and induced odd paths””. *Discrete Mathematics* 102.1 (1992), p. 109.

[Bir48] G. Birkhoff. *Lattice theory*. revised edition. Colloquium publications. American Mathematical Society, 1948.

[BP93] J. R. Blair and B. Peyton. “An introduction to chordal graphs and clique trees”. *Graph theory and sparse matrix computation*. 1993, pp. 1–29.

[BJKKP] D. Blankenhagen, R. W. Jacobs, P. Knappe, J. Kurkofka, and N. Pich Preuss. “The local 2-separator theorem via coverings”. In preparation.

[Bol78] B. Bollobas. *Extremal Graph Theory*. Academic Press, 1978.

[BBEGLPS23] 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* 26.10 (2023), pp. 3739–3791. arXiv: `2012.02435`.

[BDEW24] M. Bonamy, J. Davies, L. Esperet, and A. Wesolek. *Personal communication*. 2024.

[Bow21] N. Bowler. *Personal communication*. 2021.

[BGHKMR23] N. Bowler, F. Gut, M. Hatzel, K.-i. Kawarabayashi, I. Muzi, and F. Reich. “Decomposition of (Infinite) Digraphs along Directed 1-Separations”. *arXiv preprint* (2023). arXiv: `2305.09192`.

[BDS05] H. Bruhn, R. Diestel, and M. Stein. “Menger’s theorem for infinite graphs with ends”. *Journal of Graph Theory* 50.3 (2005), pp. 199–211.

[BS07] H. Bruhn and M. Stein. “On end degrees and infinite cycles in locally finite graphs”. *Combinatorica* 27.3 (2007), p. 269.

[BFT24] B. M. Bumpus, J. Fairbanks, and W. J. Turner. “Pushing Tree Decompositions Forward along Graph Homomorphisms”. *arXiv preprint* (2024). arXiv: 2408.15184.

[BKM22] B. M. Bumpus, Z. A. Kocsis, and J. E. Master. “Structured Decompositions: Structural and Algorithmic Compositionality”. *arXiv preprint* (2022). arXiv: 2207.06091.

[BK22] C. Bürger and J. Kurkofka. “Duality theorems for stars and combs IV: Undominating stars”. *Journal of Graph Theory* 100.1 (2022), pp. 140–162. arXiv: 2004.00591.

[Bur65] J. P. Burling. “On coloring problems of families of prototypes”. PhD thesis. University of Colorado, 1965.

[CGS17] Y. Cao, L. N. Grippo, and M. D. Safe. “Forbidden induced subgraphs of normal Helly circular-arc graphs: Characterization and detection”. *Discrete Applied Mathematics* 216 (2017), pp. 67–83. arXiv: 1405.0329.

[Car16] J. Carmesin. “A short proof that every finite graph has a tree-decomposition displaying its tangles”. *European Journal of Combinatorics* 58 (2016). arXiv: 1511.02734.

[Car19] J. Carmesin. “All graphs have tree-decompositions displaying their topological ends”. *Combinatorica* 39.3 (2019), pp. 545–596. arXiv: 1409.6640.

[Car21] J. Carmesin. “Local k-separators”. Manuscript. 2021.

[Car22] J. Carmesin. “Local 2-separators”. *Journal of Combinatorial Theory, Series B* 156 (2022), pp. 101–144. arXiv: 2008.03032.

[CDHH14] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. “k-Blocks: a connectivity invariant for graphs”. *SIAM Journal on Discrete Mathematics* 28 (2014), pp. 1876–1891. arXiv: 1305.4557.

[CDHH16a] 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), pp. 1–24. arXiv: 1406.3797.

[CDHH16b] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. “Canonical tree-decompositions of finite Graphs II. Essential parts”. *Journal of Combinatorial Theory, Series B* 118 (2016), pp. 268–283. arXiv: 1305.4909.

[CDHS14] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. “Connectivity and tree structure in finite graphs”. *Combinatorica* 34.1 (2014), pp. 1–35. arXiv: 1105.1611.

[CF23] J. Carmesin and S. Frenkel. “How to apply tree decomposition ideas in large networks”. *arXiv preprint* (2023). arXiv: 2312.12354.

[CG17] 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), pp. 1–20. arXiv: 1506.02904.

[CHM22] J. Carmesin, M. Hamann, and B. Miraftab. “Canonical trees of tree-decompositions”. *Journal of Combinatorial Theory, Series B* 152 (2022), pp. 1–26. arXiv: 2002.12030.

[CJKK25] 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). arXiv: 2501.16170.

[CKKT24] J. Carmesin, G. Kontogeorgiou, J. Kurkofka, and W. J. Turner. “Towards a Stallings-type theorem for finite groups”. *arXiv preprint* (2024). arXiv: 2403.07776.

[CK23a] J. Carmesin and J. Kurkofka. “Canonical Decompositions of 3-Connected Graphs”. *IEEE 64th Annual Symposium on Foundations of Computer Science*. 2023, pp. 1887–1920. arXiv: 2304.00945.

[CK23b] J. Carmesin and J. Kurkofka. “Characterising 4-tangles through a connectivity property”. *arXiv preprint* (2023). arXiv: 2309.00902.

[CK24] J. Carmesin and J. Kurkofka. “Entanglements”. *Journal of Combinatorial Theory, Series B* 164 (2024), pp. 17–28. arXiv: 2205.11488.

[CDEHV08] V. Chepoi, F. Dragan, B. Estellon, M. Habib, and Y. Vaxès. “Diameters, centers, and approximating trees of delta-hyperbolicgeodesic spaces and graphs”. *Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry*. 2008, pp. 59–68.

[CDNRV12] V. Chepoi, F. 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), pp. 187–214. arXiv: 1007.0489.

[Che77] B. V. Cherkassky. “A solution of a problem on multicommodity flows in a network”. *Ekonomika i Matematicheskie Metody* 13.1 (1977). In Russian, pp. 143–151.

[Cou90] B. Courcelle. “The monadic second-order logic of graphs. I. Recognizable sets of finite graphs”. *Information and Computation* 85.1 (1990), pp. 12–75.

[DMŠ24] C. Dallard, M. Milanič, and K. Štorgel. “Treewidth versus clique number. II. Tree-independence number”. *Journal of Combinatorial Theory, Series B* 164 (2024), pp. 404–442. arXiv: 2206.15092.

[DDFMMS14] M. DeVos, Z. Dvořák, J. Fox, J. McDonald, B. Mohar, and D. Scheide. "A minimum degree condition forcing complete graph immersion". *Combinatorica* 34.3 (2014), pp. 279–298. arXiv: 1101.2630.

[DMMS13] M. DeVos, J. McDonald, B. Mohar, and D. Scheide. "A note on forbidding clique immersions". *The Electronic Journal of Combinatorics* 20.3 (2013), #P55. arXiv: 1207.2117.

[Die88] R. Diestel. "Tree-decompositions, tree-representability and chordal graphs". *Discrete Mathematics* 71.2 (1988), pp. 181–184.

[Die94] R. Diestel. "The depth-first search tree structure of $TK^{\aleph_0}$-free graphs". *Journal of Combinatorial Theory, Series B* 61.2 (1994), pp. 260–262.

[Die03] R. Diestel. "The countable Erdős–Menger conjecture with ends". *Journal of Combinatorial Theory, Series B* 87.1 (2003), pp. 145–161.

[Die10] R. Diestel. "Locally finite graphs with ends: A topological approach, II. Applications". *Discrete Mathematics* 310.20 (2010), pp. 2750–2765. arXiv: 0912.4213.

[Die11] R. Diestel. "Locally finite graphs with ends: A topological approach, I. Basic theory". *Discrete Mathematics* 311.15 (2011), pp. 1423–1447. arXiv: 0912.4213.

[Die17] R. Diestel. "Ends and tangles". *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 87, Special issue in memory of Rudolf Halin (2017), pp. 223–244. arXiv: 1510.04050.

[Die18a] R. Diestel. "Abstract separation systems". *Order* 35 (2018), pp. 157–170. arXiv: 1406.3797.

[Die18b] R. Diestel. "Tree sets". *Order* 35.1 (2018), pp. 171–192. arXiv: 1512.03781.

[Die20] R. Diestel. *Personal communication*. 2020.

[Die24] R. Diestel. *Tangles: a structural approach to artificial intelligence in the empirical sciences*. Cambridge University Press, 2024.

[Die25] R. Diestel. *Graph Theory*. 6th ed. Vol. 173. Graduate Texts in Mathematics. Electronic edition available at www.diestel-graph-theory.com. Springer, 2025.

[DEJ24] R. Diestel, C. Elbracht, and R. W. Jacobs. "Point sets and functions inducing tangles of set separations". *Journal of Combinatorics* 15.3 (2024), pp. 283–306. arXiv: 2107.01087.

[DEW19] R. Diestel, J. Erde, and D. Weißauer. "Structural submodularity and tangles in abstract separation systems". *Journal of Combinatorial Theory, Series A* 167 (2019), pp. 155–180. arXiv: 1805.01439.

[DHL19] R. Diestel, F. Hundertmark, and S. Lemanczyk. “Profiles of separations: in graphs, matroids, and beyond”. *Combinatorica* 39.1 (2019), pp. 37–75. arXiv: 1110.6207.

[DJKK22] R. Diestel, R. W. Jacobs, P. Knappe, and J. Kurkofka. “Canonical graph decompositions via coverings”. *arXiv preprint* (2022). arXiv: 2207.04855.

[DJKW23] R. Diestel, R. W. Jacobs, P. Knappe, and P. Wollan. “A grid theorem for strong immersions of walls”. *arXiv preprint* (2023). arXiv: 2301.05134.

[DK03] R. Diestel and D. Kühn. “Graph-theoretical versus topological ends of graphs”. *Journal of Combinatorial Theory, Series B* 87 (2003), pp. 197–206.

[DK04a] R. Diestel and D. Kühn. “On infinite cycles I”. *Combinatorica* 24.1 (2004), pp. 69–89.

[DK04b] R. Diestel and D. Kühn. “On infinite cycles II”. *Combinatorica* 24.1 (2004), pp. 91–116.

[DK04c] R. Diestel and D. Kühn. “Topological paths, cycles and spanning trees in infinite graphs”. *European Journal of Combinatorics* 25.6 (2004), pp. 835–862.

[DK05] R. Diestel and D. Kühn. “Graph minor hierarchies”. *Discrete Applied Mathematics* 145 (2005), pp. 167–182.

[DL01] R. Diestel and I. Leader. “Normal spanning trees, Aronszajn trees and excluded minors”. *Journal of the London Mathematical Society* 63.1 (2001), pp. 16–32.

[DO19] R. Diestel and S.-i. Oum. “Tangle-tree duality: in graphs, matroids and beyond”. *Combinatorica* 39.4 (2019), pp. 879–910. arXiv: 1701.02651.

[DS12] R. Diestel and P. Sprüssel. “Locally finite graphs with ends: A topological approach, III. Fundamental group and homology”. *Discrete Mathematics* 312.1 (2012), pp. 21–29. arXiv: 1004.0110.

[Dir61] G. A. Dirac. “On rigid circuit graphs”. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 25.1 (1961), pp. 71–76.

[Djo74] D. Djoković. “Automorphisms of graphs and coverings”. *Journal of Combinatorial Theory, Series B* 16.3 (1974), pp. 243–247.

[DG07] Y. Dourisboure and C. Gavoille. “Tree-decompositions with bags of small diameter”. *Discrete Mathematics* 307.16 (2007), pp. 2008–2029.

[DK14a] 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), pp. 884–905.

[Dun85] M. Dunwoody. “The accessibility of finitely presented groups”. *Inventiones Mathematicae* 81.3 (1985), pp. 449–457.

[DK14b] Z. Dvořák and T. Klimosova. "Strong immersions and maximum degree". *SIAM Journal on Discrete Mathematics* 28.1 (2014), pp. 177–187. arXiv: 1304.0728.

[DW16] Z. Dvořák and P. Wollan. "A structure theorem for strong immersions". *Journal of Graph Theory* 83.2 (2016), pp. 152–163. arXiv: 1411.0522.

[Elb15] C. Elbracht. "Tangles determined by majority vote". Electronic version available at www.math.uni-hamburg.de/spag/dm/papers/Elbracht_MSc.pdf. Master's thesis. University of Hamburg, 2015.

[EFKKRTL23] C. Elbracht, D. Fioravanti, S. Klepper, J. Kneip, L. Rendsburg, M. Teegen, and U. von Luxburg. "Clustering with tangles: Algorithmic framework and theoretical guarantees". *Journal of Machine Learning Research* 24.190 (2023), pp. 1–56. arXiv: 2006.14444.

[EKT20] C. Elbracht, J. Kneip, and M. Teegen. "Tangles are decided by weighted vertex sets". *Advances in Combinatorics* (2020). arXiv: 1811.06821.

[EKT21] C. Elbracht, J. Kneip, and M. Teegen. "Trees of tangles in abstract separation systems". *Journal of Combinatorial Theory, Series A* 180 (2021), p. 105425. arXiv: 1909.09030.

[EKT22] C. Elbracht, J. Kneip, and M. Teegen. "Trees of tangles in infinite separation systems". *Mathematical Proceedings of the Cambridge Philosophical Society* 173.2 (2022), pp. 297–327. arXiv: 2005.12122.

[EK22] A.-K. Elm and J. Kurkofka. "A tree-of-tangles theorem for infinite tangles". *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 92.2 (2022), pp. 139–178. arXiv: 2003.02535.

[Erd17] J. Erde. "Refining a tree-decomposition which distinguishes tangles". *SIAM Journal on Discrete Mathematics* 31.3 (2017), pp. 1529–1551. arXiv: 1512.02499.

[Erd18] J. Erde. "A unified treatment of linked and lean tree-decompositions". *Journal of Combinatorial Theory, Series B* 130 (2018), pp. 114–143. arXiv: 1703.03756.

[EGL24] L. Esperet, U. Giocanti, and C. Legrand-Duchesne. "The structure of quasi-transitive graphs avoiding a minor with applications to the domino problem". *Journal of Combinatorial Theory, Series B* 169 (2024), pp. 561–613. arXiv: 2304.01823.

[FGTW08] M. Ferrara, R. J. Gould, G. Tansey, and T. Whalen. "On H-immersions". *Journal of Graph Theory* 57.3 (2008), pp. 245–254.

[FV10] F. V. Fomin and Y. Villanger. “Finding Induced Subgraphs via Minimal Triangulations”. *27th International Symposium on Theoretical Aspects of Computer Science*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik. 2010, pp. 383–394. arXiv: 0909.5278.

[FP23] K. Fujiwara and P. Papasoglu. “A coarse-geometry characterization of cacti”. *arXiv preprint* (2023). arXiv: 2305.08512.

[GKL23] P. Gartland, T. Korhonen, and D. Lokshtanov. “On Induced Versions of Menger’s Theorem on Sparse Graphs”. *arXiv preprint* (2023). arXiv: 2309.08169.

[Gav72] F. Gavril. “Algorithms for minimum coloring, maximum clique, minimum covering by cliques and maximum independent set of a chordal graph”. *SIAM Journal of Computing* 1 (1972), pp. 180–187.

[Gav74] F. Gavril. “The intersection graphs of subtrees in trees are exactly the chordal graphs”. *Journal of Combinatorial Theory, Series B* 16.1 (1974), pp. 47–56.

[Gav87] F. Gavril. “Generating the maximum spanning trees of a weighted graph”. *Journal of Algorithms* 8.4 (1987), pp. 592–597.

[Gav96] F. Gavril. “Intersection graphs of Helly families of subtrees”. *Discrete Applied Mathematics* 66 (1996), pp. 45–56.

[GKMW23] J. Geelen, O.-J. Kwon, R. McCarty, and P. Wollan. “The grid theorem for vertex minors”. *Journal of Combinatorial Theory, Series B* 158 (2023), pp. 93–116. arXiv: 1909.08113.

[Geo17] A. Georgakopoulos. “On covers of graphs by Cayley graphs”. *European Journal of Combinatorics* 64 (2017), pp. 57–65. arXiv: 1504.00173.

[GP23] A. Georgakopoulos and P. Papasoglu. “Graph minors and metric spaces”. *arXiv preprint* (2023). arXiv: 2305.07456.

[GHW25] J. P. Gollin, M. Hatzel, and S. Wiederrecht. “On graphs with a simple structure of maximal cliques”. *arXiv preprint* (2025). arXiv: 2504.16863.

[GH22] J. P. Gollin and K. Heuer. “Characterising k-connected sets in infinite graphs”. *Journal of Combinatorial Theory, Series B* 157 (2022), pp. 451–499. arXiv: 1811.06411.

[GK21] J. P. Gollin and J. Kneip. “Representations of infinite tree sets”. *Order* 38.1 (2021), pp. 79–96. arXiv: 1908.10327.

[Gro16a] M. Grohe. “Quasi-4-Connected Components”. *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. 2016, pp. 8–1. arXiv: 1602.04505.

[Gro16b] M. Grohe. “Tangles and connectivity in graphs”. *Language and Automata Theory and Applications* (2016), pp. 24–41. arXiv: 1602.04727.

[GKMW11] M. Grohe, K.-i. Kawarabayashi, D. Marx, and P. Wollan. "Finding topological subgraphs is fixed-parameter tractable". *Proceedings of the Forty-Third annual ACM symposium on Theory of Computing*. 2011, pp. 479–488. arXiv: 1011.1827.

[GS15a] M. Grohe and P. Schweitzer. "Computing with Tangles". *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*. 2015, pp. 683–692. arXiv: 1503.00190. Also published in: *SIAM Journal on Discrete Mathematics* 30.2 (2016), pp. 1213–1247.

[GS15b] M. Grohe and P. Schweitzer. "Isomorphism testing for graphs of bounded rank width". *IEEE 56th Annual Symposium on Foundations of Computer Science* (2015), pp. 1010–1029. arXiv: 1505.03737.

[Gro93] M. Gromov. *Geometric group theory. Volume 2: Asymptotic invariants of infinite groups. Proceedings of the symposium held at the Sussex University, Brighton, July 14-19, 1991*. Vol. 182. London Mathematical Society Lecture Note Series. Cambridge University Press, 1993.

[GL70] A. Gyárfás and J. Lehel. "A Helly-type problem in trees". *Combinatorial theory and its applications* 4 (1970), pp. 571–584.

[Hal65] R. Halin. "Über die Maximalzahl fremder unendlicher Wege in Graphen". *Mathematische Nachrichten* 30 (1965), pp. 63–85.

[Hal75] R. Halin. "Some path problems in graph theory". *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 44.1 (1975), pp. 175–186.

[Hal77] R. Halin. "Systeme disjunkter unendlicher Wege in Graphen". *Numerische Methoden bei Optimierungsaufgaben Band 3: Optimierung bei graphentheoretischen und ganzzahligen Problemen* (1977), pp. 55–67.

[Hal78] R. Halin. "Simplicial decompositions of infinite graphs". *Annals of Discrete Mathematics* 3 (1978), pp. 93–109.

[Hal91a] R. Halin. "Lattices of Cuts in Graphs". *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 61.1 (1991), pp. 217–230.

[Hal91b] R. Halin. "Tree-partitions of infinite graphs". *Discrete Mathematics* 97.1-3 (1991), pp. 203–217.

[Ham18] M. Hamann. "Accessibility in transitive graphs". *Combinatorica* 38.4 (2018), pp. 847–859. arXiv: 1404.7677.

[HLMR22] 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), pp. 40–69. arXiv: 1812.06312.

[HW17] D. J. Harvey and D. R. Wood. "Parameters tied to treewidth". *Journal of Graph Theory* 84.4 (2017), pp. 364–385. arXiv: 1312.3401.

[Hås99] J. Håstad. "Clique is Hard to Approximate within $n1-\epsilon$". *Acta Mathematica* 182 (1999), pp. 105–142.

[Hat02] A. Hatcher. *Algebraic Topology*. Electronic edition available at `https://pi.math.cornell.edu/~hatcher/AT/ATpage.html`. Cambrigde University Press, 2002.

[HNST23] K. Hendrey, S. Norin, R. Steiner, and J. Turcotte. "On an induced version of Menger's theorem". *The Electronic Journal of Combinatorics* (2023), #P4–28. arXiv: `2309.07905`.

[HL18] C. Hofer-Temmel and F. Lehner. "Clique Trees of Infinite Locally Finite Chordal Graphs". *The Electronic Journal of Combinatorics* 25.2 (2018), #P2.9. arXiv: `1311.7001`.

[JK24] R. W. Jacobs and P. Knappe. "Efficiently distinguishing all tangles in locally finite graphs". *Journal of Combinatorial Theory, Series B* 167 (2024), pp. 189–214. arXiv: `2303.09332`.

[JJKKM23] R. W. Jacobs, A. Joó, P. Knappe, J. Kurkofka, and R. Melcher. "The Lovász-Cherkassky theorem for locally finite graphs with ends". *Discrete Mathematics* 346.12 (2023), #113586. arXiv: `2109.11260`.

[JRFR24] James Davies, Robert Hickingbotham, Freddie Illingworth, and Rose McCarty. "Fat minors cannot be thinned (by quasi-isometries)". *arXiv preprint* (2024). arXiv: `2405.09383`.

[Jän97] K. Jänich. *Topology*. Undergraduate Texts in Mathematics. Springer, 1997.

[Joó23] A. Joó. "The Lovász-Cherkassky Theorem in countable graphs". *Journal of Combinatorial Theory, Series B* 159 (2023), pp. 1–19. arXiv: `2102.04203`.

[Jun69] H. A. Jung. "Wurzelbäume und unendliche Wege in Graphen". *Mathematische Nachrichten* 41 (1969), pp. 1–22.

[kab24] kabel abel. *Cycle space and chordal graphs*. Mathematics Stack Exchange: `https://math.stackexchange.com/q/4854624`. 2024.

[Kar72] R. M. Karp. "Reducibility among combinatorial problems". *In Complexity of Computer Computations* (1972), pp. 85–103.

[KK15] K.-i. Kawarabayashi and S. Kreutzer. "The directed grid theorem". *Proceedings of the Forty-Seventh annual ACM symposium on Theory of Computing*. 2015, pp. 655–664. arXiv: `1411.5681`.

[KNTW12] 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). arXiv: `1203.2171`.

[KK24] P. Knappe and J. Kurkofka. "The immersion-minimal infinitely edge-connected graph". *Journal of Combinatorial Theory, Series B* 164 (2024), pp. 492–516. arXiv: 2207.08459.

[KT] P. Knappe and T. T. M. Trinh. "Fundamental groups of planar graphs". In preparation.

[KKP23] M. Koloschin, T. Krill, and M. Pitz. "End spaces and tree-decompositions". *Journal of Combinatorial Theory, Series B* 161 (2023), pp. 147–179. arXiv: 2205.09865.

[Kom95] P. Komjáth. "A note on minors of uncountable graphs". *Mathematical Proceedings of the Cambridge Philosophical Society* 117.1 (1995), pp. 7–9.

[Kön36] D. König. *Theorie der endlichen und unendlichen Graphen*. Akademische Verlagsgesellschaft, 1936.

[Kor24] T. Korhonen. "Linear-Time Algorithms for k-Edge-Connected Components, k-Lean Tree Decompositions, and More". *arXiv preprint* (2024). arXiv: 2411.02658.

[KLMPW03] D. Kratsch, H.-O. Le, H. Müller, E. Prisner, and D. Wagner. "Additive Tree Spanners". *SIAM Journal on Discrete Mathematics* 17.2 (2003), pp. 332–340.

[KT91] I. Kříž and R. Thomas. "The Menger-like Property of the Tree-Width of Infinite Graphs". *Journal of Combinatorial Theory, Series B* 52.1 (1991), pp. 86–91.

[Krö10] B. Krön. "Cutting up Graphs Revisited–a Short Proof of Stallings' Structure Theorem". *Groups Complexity Cryptology* 2.2 (2010), pp. 213–221. arXiv: 1003.1096.

[Kru56] J. B. Kruskal. "On the shortest spanning subtree of a graph and the traveling salesman problem". *Proceedings of the American Mathematical Society* 7 (1956), pp. 48–50.

[KMP22] J. Kurkofka, R. Melcher, and M. Pitz. "A strengthening of Halin's grid theorem". *Mathematika* 68.4 (2022), pp. 1009–1013. arXiv: 2104.10672.

[Lei17] A. Leitert. "Tree-Breadth of Graphs with Variants and Applications". PhD thesis. Kent State University, 2017.

[LB62] C. Lekkeikerker and J. Boland. "Representation of a finite graph by a set of intervals on the real line". *Fundamenta Mathematicae* 51.1 (1962), pp. 45–64.

[Löh17] C. Löh. *Geometric Group Theory*. Universitext. Springer, 2017.

[LVV14] D. Lokshantov, M. Vatshelle, and Y. Villanger. "Independent set in P 5-free graphs in polynomial time". *Proceedings of the Twenty-Fifth annual ACM-SIAM Symposium on Discrete Algorithms*. 2014, pp. 570–581.

[Lov76] L. Lovász. “On some connectivity properties of Eulerian graphs”. *Acta Mathematica Academiae Scientiarum Hungarica* 28.1–2 (1976), pp. 129–138.

[MKS04] W. Magnus, A. Karrass, and D. Solitar. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*. Courier Corporation, 2004.

[Man05] J. F. Manning. “Geometry of pseudocharacters”. *Geometry & Topology* 9.2 (2005), pp. 1147–1185. arXiv: `math/0303380`.

[MW14] D. Marx and P. Wollan. “Immersions in highly edge connected graphs”. *SIAM Journal on Discrete Mathematics* 28.1 (2014), pp. 503–520. arXiv: `1305.1331`.

[Men27] K. Menger. “Zur allgemeinen Kurventheorie”. *Fundamenta mathematicae* 10.1 (1927), pp. 96–115.

[Mil68] J. Milnor. “A note on curvature and fundamental group”. *Journal of Differential Geometry* 2.1 (1968), pp. 1–7.

[Mur92] O. J. Murphy. “Computing independent sets in graphs with large girth”. *Discrete Applied Mathematics* 35 (1992), pp. 167–170.

[NSS24] T. Nguyen, A. Scott, and P. D. Seymour. “A counterexample to the coarse Menger conjecture”. *arXiv preprint* (2024). arXiv: `2401.06685`.

[NSS25] T. Nguyen, A. Scott, and P. D. Seymour. “Coarse tree-width”. *arXiv preprint* (2025). arXiv: `2501.09839`.

[Pit20] M. Pitz. “A unified existence theorem for normal spanning trees”. *Journal of Combinatorial Theory, Series B* 145 (2020), pp. 466–469. arXiv: `2003.11575`.

[Pit21] M. Pitz. “Proof of Halin’s normal spanning tree conjecture”. *Israel Journal of Mathematics* 246.1 (2021), pp. 353–370. arXiv: `2005.02833`.

[Pit23] M. Pitz. “A note on minor antichains of uncountable graphs”. *Journal of Graph Theory* 102.3 (2023), pp. 552–555. arXiv: `2005.05816`.

[Pla24] T. Planken. “Colouring the 1-Skeleton of d-Dimensional Triangulations”. *arXiv preprint* (2024). arXiv: `2409.11762`.

[Pol96] N. Polat. “Ends and multi-endings I”. *Journal of Combinatorial Theory, Series B* 67 (1996), pp. 86–110.

[Pol74] S. Poljak. “A note on stable sets and colorings of graphs”. *Commentationes Mathematicae Universitatis Carolinae* 15 (1974), pp. 307–309.

[PT23] P. Pournajafi and N. Trotignon. “Burling graphs revisited, part I: New characterizations”. *European Journal of Combinatorics* 110 (2023), #103686. arXiv: `2104.07001`.

[PT24] P. Pournajafi and N. Trotignon. “Burling graphs revisited, part II: Structure”. *European Journal of Combinatorics* 116 (2024), #103849. arXiv: `2106.16089`.

[RS86a] N. Robertson and P. D. Seymour. "Graph minors. II. Algorithmic aspects of tree-width". *Journal of Algorithms* 7.3 (1986), pp. 309–322.

[RS86b] N. Robertson and P. D. Seymour. "Graph minors. V. Excluding a planar graph". *Journal of Combinatorial Theory, Series B* 41 (1986), pp. 92–114.

[RS90a] N. Robertson and P. D. Seymour. "Graph Minors. IV. Tree-Width and Well-Quasi-Ordering". *Journal of Combinatorial Theory, Series B* 48 (1990), pp. 227–254.

[RS90b] N. Robertson and P. D. Seymour. "Graph minors. VIII. A Kuratowski theorem for general surfaces". *Journal of Combinatorial Theory, Series B* 48.2 (1990), pp. 255–288.

[RS91] N. Robertson and P. D. Seymour. "Graph minors. X. Obstructions to tree-decomposition". *Journal of Combinatorial Theory, Series B* 52 (1991), pp. 153–190.

[RS95] N. Robertson and P. D. Seymour. "Graph minors. XIII. The disjoint paths problem". *Journal of Combinatorial Theory, Series B* 63 (1995), pp. 65–110.

[RS03] N. Robertson and P. D. Seymour. "Graph minors. XVI. Excluding a non-planar graph". *Journal of Combinatorial Theory, Series B* 89 (2003), pp. 43–76.

[RS04] N. Robertson and P. D. Seymour. "Graph minors. XX. Wagner's conjecture". *Journal of Combinatorial Theory, Series B* 92.2 (2004), pp. 325–357.

[RS10] N. Robertson and P. D. Seymour. "Graph minors. XXIII. Nash-Williams' immersion conjecture". *Journal of Combinatorial Theory, Series B* 100 (2010), pp. 181–205.

[RS] N. Robertson and P. D. Seymour. "Graph minors. I–XX". *Journal of Combinatorial Theory, Series B* (1983–2004).

[RST95] N. Robertson, P. D. Seymour, and R. Thomas. "Excluding infinite clique minors". *Memoirs of the American Mathematical Society* 118.566 (1995), pp. vi+103.

[Ros70] D. J. Rose. "Triangulated graphs and the elimination process". *Journal of Mathematical Analysis and Applications* 32.3 (1970), pp. 597–609.

[Rüh17] T. Rühmann. "A Study of Infinite Graphs of a Certain Symmetry and Their Ends". PhD thesis. University of Hamburg, 2017.

[Sco22] A. Scott. "Graphs of large chromatic number". *International Congress of Mathematicians*. 2022, pp. 6–14.

[ST93] P. D. Seymour and R. Thomas. "Excluding infinite trees". *Transactions of the American Mathematical Society* 335.2 (1993), pp. 597–630.

[Sta72] J. R. Stallings. *Group theory and three-dimensional manifolds*. Yale University Press, 1972.

[Sta83] J. R. Stallings. "Topology of finite graphs". *Inventiones mathematicae* 71.3 (1983), pp. 551–565.

[Šva55] A. S. Švarc. "A volume invariant of coverings". *Dokl. akad. nauk sssr (ns)*. Vol. 105. 1955, pp. 32–34.

[Tho88] R. Thomas. "A counter-example to 'Wagner's conjecture' for infinite graphs". *Mathematical Proceedings of the Cambridge Philosophical Society* 103.1 (1988), pp. 55–57.

[Tho89] R. Thomas. "Well-quasi-ordering infinite graphs with forbidden finite planar minor". *Transactions of the American Mathematical Society* 312.1 (1989), pp. 279–313.

[Tho90] R. Thomas. "A menger-like property of tree-width: The finite case". *Journal of Combinatorial Theory, Series B* 48 (1990), pp. 67–76.

[TW93] C. Thomassen and W. Woess. "Vertex-transitive graphs and accessibility". *Journal of Combinatorial Theory, Series B* 58 (1993), pp. 248–268.

[Tut84] W. T. Tutte. *Graph Theory*. Vol. 21. Encyclopedia of Mathematics and Its Applications. Addison-Wesley Publishing Company, 1984.

[Wag37] K. Wagner. "Über eine Eigenschaft der ebenen Komplexe". *Mathematische Annalen* 114.1 (1937), pp. 570–590.

[WJ08] M. J. Wainwright and M. I. Jordan. *Graphical Models, Exponential Families, and Variational Inference*. Foundations and Trends in Machine Learning. Now Publishers, Inc., 2008.

[Wal71] C. Wall. "Pairs of relative cohomological dimension one". *Journal of Pure and Applied Algebra* 1.2 (1971), pp. 141–154.

[Wol15] P. Wollan. "The structure of graphs not admitting a fixed immersion". *Journal of Combinatorial Theory, Series B* 110 (2015), pp. 47–66. arXiv: 1302.3867.

Appendices

CHAPTER A

Summary of the results

This thesis discusses problems in graph theory, a branch of discrete mathematics. Most of these problems ask for a specific type of decomposition of a given graph that is modelled on another graph. The thesis is divided into three parts. In the first part, we make the global structure of graphs visible. The second part deals with questions in the recently emerging field of coarse graph theory. The third part addresses several problems concerning tree-like structures and the obstructions to their existence.

Part I: Local versus global aspects of graphs

Chapter 1: Canonical graph-decompositions via coverings

We present a canonical way to decompose finite graphs into highly connected local parts. The decomposition depends only on an integer parameter whose choice sets the intended degree of locality. The global structure of the graph, as determined by the relative position of these parts, is described by a coarser *model*. This is a simpler graph determined by the decomposition, not imposed.

The model and decomposition are obtained as projections of the tangle-tree structure of a covering of the given graph that reflects its local structure while unfolding its global structure. In this way, the tangle theory from graph minors is brought to bear canonically on arbitrary graphs, which need not be globally tree-like.

Our theorem extends to locally finite quasi-transitive graphs, and in particular to locally finite Cayley graphs. It thereby yields a canonical decomposition of finitely generated groups into local parts, whose relative structure is displayed by a graph.

Chapter 2: Canonical graph-decompositions via local separations

In Chapter 1 we have shown that every finite graph G can be decomposed in a canonical way that displays its global structure. Recall that these decomposition is defined via a suitable globally more tree-like covering of G, whose tangle-tree structure is projected down to G. The covering graphs needed here are almost always infinite. The canonical decompositions they induce on G are therefore not computable following their definition.

We reconstruct these decompositions of G from finite information in G itself that is sufficiently local to be reflected in the cover. This involves the reconstruction of canonical tangle structure in terms of a new theory of local separations in finite graphs, which we develop for this purpose.

As an application, we find that the canonical graph-decompositions from Chapter 1 are computable.

Chapter 3: Locally chordal graphs

A graph is locally chordal if, locally at each vertex, it is chordal, i.e. composed of cliques glued along a tree. In addition to extending many classic characterizations of chordal graphs, we show that locally chordal graphs are precisely those graphs G which can be decomposed into cliques arranged in the shape of a high-girth graph. The decomposition in the shape of a high-girth graph simultaneously witnesses that G is, locally at each vertex, composed of cliques glued along a tree. We also show that these decompositions can be computed efficiently. Our analysis of locally chordal graphs is inspired by the framework to study local versus global structure in graphs built in Chapter 1 and also Chapter 2.

Part II: Coarse graph theory

Chapter 4: 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.

Chapter 5: A characterisation of 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.

Chapter 6: 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 $\mathrm{d}_G(P_1, P_2) \geqslant d$, or there exists $v \in V(G)$ such that the ball of radius $f(d)$ centered at v intersects every $X - Y$ path.

Part III: Tree-like structures and tangles

Chapter 7: 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.

Chapter 8: 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.

Chapter 9: Efficiently distinguishing all tangles in locally finite graphs

While finite graphs have tree-decompositions that efficiently distinguish all their tangles, locally finite graphs with thick ends need not have such tree-decompositions. We show that every locally finite graph without thick ends admits such a tree-decomposition, in fact a canonical one. Our proof exhibits a thick end at any obstruction to the existence of such tree-decompositions and builds on new methods for the analysis of the limit behaviour of strictly increasing sequences of separations.

Chapter 10: 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.

Chapter 11: A grid theorem for strong immersions of walls

We show that a graph contains a large wall as a strong immersion minor if and only if the graph does not admit a tree-cut decomposition of small 'width', which is measured in terms of its adhesion and the path-likeness of its torsos.

Chapter 12: The Lovász-Cherkassky theorem for locally finite graphs with ends

Lovász and Cherkassky discovered independently that, if G is a finite graph and $T \subseteq V(G)$ such that the degree $d_G(v)$ is even for every vertex $v \in V(G) \setminus T$, then the maximum number of edge-disjoint paths which are internally disjoint from T and connect distinct vertices of T is equal to $\frac{1}{2}\sum_{t\in T} \lambda_G(t, T \setminus \{t\})$ (where $\lambda_G(t, T \setminus \{t\})$ is the size of a smallest cut that separates t and $T \setminus \{t\}$). From another perspective, this means that for every vertex $t \in T$, in any optimal path-system there are $\lambda_G(t, T \setminus \{t\})$ many paths between t and $T \setminus \{t\}$. We extend the theorem of Lovász and Cherkassky based on this reformulation to all locally-finite infinite graphs and their ends. In our generalisation, T may contain not just vertices but ends as well, and paths are one-way (two-way) infinite when they establish a vertex-end (end-end) connection.

CHAPTER B

Zusammenfassung der Ergebnisse

Diese Arbeit befasst sich mit Problemen der Graphentheorie, einem Teilgebiet der diskreten Mathematik. Die meisten dieser Probleme verlangen nach einer bestimmten Art von Zerlegung eines gegebenen Graphen, modelliert auf einem anderen Graphen. Die Arbeit ist in drei Teile gegliedert. Im ersten Teil machen wir die globale Struktur von Graphen sichtbar. Der zweite Teil befasst sich mit Fragen aus dem jüngst entstandenen Forschungsgebiet der groben Graphentheorie. Der dritte Teil beschäftigt sich mit verschiedenen Problemen, die baumartige Strukturen und deren Hindernisse betreffen.

Teil I: Lokale versus globale Aspekte von Graphen

Kapitel 1: Kanonische Graphenzerlegungen mittels Überlagerungen

Wir präsentieren eine kanonische Weise, um endliche Graphen in hochgradig verbundene lokale Teile zu zerlegen. Die Zerlegung hängt nur von einem ganzzahligen Parameter ab, dessen Wahl den gewünschten Grad der Lokalität festlegt. Die globale Struktur des Graphen, die durch die relative Position dieser Teile bestimmt wird, wird durch ein *Modell* beschrieben. Dabei handelt es sich um einen einfacheren Graphen, der durch die Zerlegung entsteht und nicht im Vorhinein auferlegt wird.

Wir erhalten das Modell und die Zerlegung als Projektionen der Knäuel-Baum-Struktur einer Überlagerung des gegebenen Graphen G. Die betrachtete Überlagerung spiegelt die lokale Struktur von G wider und faltet gleichzeitig seine globale Struktur auf. So kann die Knäueltheorie kanonisch auf beliebige Graphen angewandt werden, die nicht global baumartig sein müssen.

Außerdem zeigen wir unser Theorem auch für lokal endliche quasi-transitive Graphen, wozu insbesondere auch lokal endliche Cayleygraphen gehören. Dadurch ergibt sich eine kanonische Zerlegung endlich erzeugter Gruppen in lokale Teile, deren relative Struktur durch einen Graphen dargestellt wird.

Kapitel 2: Kanonische Graphenzerlegungen mittels lokaler Teilungen

In Kapitel 1 haben wir jeden endlichen Graphen G auf kanonische Weise so zerlegt, dass seine globale Struktur dargestellt wird. Wir erinnern uns, dass diese Zerlegung über einen global baumähnlicheren Überlagerungsgraphen von G definiert wird, dessen Knäuel-Baum-Struktur auf G projiziert wird. Der hierfür benötigte Überlagerungsgraph ist in der Regel unendlich. Die oben beschriebene kanonische Zerlegung ist daher ihrer obigen Definition folgend nicht berechenbar.

Wir rekonstruieren diese Zerlegung von G aus endlicher Information in G selbst, die hinreichend lokal ist, um sie in der Überlagerung wiederzufinden. Dies beinhaltet die Rekonstruktion der kanonischen Knäuel-Struktur mithilfe einer neuen Theorie lokaler Teilungen in endlichen Graphen, die wir zu diesem Zweck entwickeln.

Als Anwendung ergibt sich, dass die kanonischen Graphenzerlegungen aus Kapitel 1 berechenbar sind.

Kapitel 3: Lokal chordale Graphen

Ein Graph ist lokal chordal, falls er lokal an jeder Ecke chordal ist, d.h. aus Cliquen besteht, die entlang eines Baumes geklebt sind. Neben der Verallgemeinerung vieler klassischer Charakterisierungen von chordalen Graphen zu lokal chordalen Graphen zeigen wir, dass lokal chordale Graphen genau die Graphen G sind, die in Cliquen zerlegt werden können, die in Form eines Graphen hoher Taillenweite angeordnet sind. Diese Zerlegung bezeugt wiederum, dass G lokal an jeder Ecke aus Cliquen besteht, die entlang eines Baumes geklebt sind, also lokal chordal ist. Wir zeigen zudem, dass diese Zerlegungen effizient berechnet werden können. Unsere Analyse von lokal chordalen Graphen ist inspiriert von dem Rahmen zur Untersuchung lokaler versus globaler Strukturen in Graphen aus Kapitel 1 und auch Kapitel 2.

Teil II: Grobe Graphentheorie

Kapitel 4: Eine strukturelle Dualität für Wegzerlegungen in Teile mit kleinem Radius

Es ist eine einfache Beobachtung, dass wenn ein Graph G eine Wegzerlegung zulässt, deren Teile einen kleinen Radius haben, dann enthält G keine große Unterteilung von $K_{1,3}$ oder K_3 als (quasi-)geodätischer Untergraph. Wir zeigen, dass dies tatsächlich die einzigen Hindernisse für solche Wegzerlegungen mit kleiner radialen Weite sind, und wir beweisen analoge Ergebnisse für Zerlegungen, die auf Kreisen und unterteilten Sternen anstelle von Wegen modelliert sind.

Mit unseren Ergebnissen bestätigen wir in eine starke Form einer Vermutung von Georgakopoulos und Papasoglu über Charakterisierungen von Graphen, die quasi-isometrisch zu Wegen, Kreisen und Wegen bzw. unterteilten Sternen sind, mithilfe von fetten Minoren. Dazu präsentieren wir eine neue Sichtweise auf Quasi-Isometrien zwischen Graphen durch Graphenzerlegungen mit beschränkter radialen Weite und Ausdehnung. Diese neue Sichtweise ermöglicht es uns, weitere Ergebnisse in der groben Graphentheorie zu beweisen, und ist daher möglicherweise von unabhängigem Interesse sein.

Kapitel 5: Eine Charakterisierung von Graphen, die quasi-isometrisch zu Graphen ohne K_4-Minor sind

Wir beweisen, dass es eine Funktion f gibt, sodass jeder Graph ohne K-fetten K_4-Minor $f(K)$-quasi-isometrisch zu einem Graph ohne K_4-Minor ist. Dies löst den K_4-Fall der allgemeinen

Vermutung von Georgakopoulos und Papasoglu. Mit unserer Beweistechnik lässt sich auch der entsprechende K_4^--Fall kurz beweisen, der zuvor von Fujiwara und Papasoglu gelöst wurde.

Kapitel 6: Ein Menger-artiger Satz für zwei induzierte Wege

Wir beweisen einen Menger-artigen Satz, der beschreibt, wann ein Graph G zwei X–Y Wege P_1 und P_2 enthält, sodass $P_1 \cup P_2$ ein induzierter Teilgraph von G ist. Genauer gesagt beweisen wir, dass es eine Funktion $f(d) \in O(d)$ gibt, sodass für jeden Graphen G und $X, Y \subseteq V(G)$, entweder zwei X–Y Wege P_1 und P_2 existieren, deren Abstand $\geqslant d$ ist, oder eine Ecke $v \in G$ existiert, sodass der Ball mit Radius $f(d)$ um v jeden X–Y Wege trifft.

Teil III: Baumähnliche Strukturen und Knäuel

Kapitel 7: Verbundene Baumzerlegungen in endliche Teile

Wir beweisen, dass jeder Graph, der eine Baumzerlegung in endliche Teile zulässt, sogar eine verwurzelte Baumzerlegung in endliche Teile hat, die verbunden, eng und komponententreu ist.

Als Anwendung erhalten wir, dass jeder Graph ohne Halbgitter-Minor eine schlanke Baumzerlegung in endliche Teile hat, was das entsprechende Ergebnis von Kříž und Thomas für Graphen mit endlich beschränkter Baumweite verstärkt. Insbesondere folgt daraus, dass jeder Graph ohne Halbgitter-Minor eine Baumzerlegung hat, die effizient alle Enden und kritischen Eckenmengen unterscheidet, was die Ergebnisse von Carmesin und von Elm und Kurkofka für diese Graphenklasse verstärkt.

Als weitere Anwendung unseres Hauptergebnisses folgt, dass jeder Graph, der eine Baumzerlegung in endliche Teile zulässt, eine Baumzerlegung in endliche Teile hat, die alle Enden von G und ihre kombinierten Grade darstellt, womit eine Frage von Halin aus dem Jahr 1977 gelöst wird. Die letztgenannte Baumzerlegung liefert außerdem kurze, einheitliche Beweise für die von Robertson, Seymour und Thomas aufgestellten Charakterisierungen von Graphen ohne Halbgitter-Minor und von Graphen ohne binäre Baumunterteilung.

Kapitel 8: Gegenbeispiele zu verbundenen und schlanken Baumzerlegungen von unendlichen Graphen

Kříž und Thomas zeigten, dass jeder (endliche oder unendliche) Graph von Baumweite $k \in \mathbb{N}$ eine schlanke Baumzerlegung der Weite k hat. Wir diskutieren eine Reihe von Gegenbeispielen, die die Grenzen möglicher Verallgemeinerungen ihres Resultats auf beliebige unendliche Baumweite aufzeigen.

Insbesondere konstruieren wir einen lokal endlichen, ebenen, zusammenhängenden Graphen, der keine schlanke Baumzerlegung besitzt.

Kapitel 9: Effiziente Unterscheidung aller Knäuel in lokal endlichen Graphen

Zwar haben endliche Graphen Baumzerlegungen, die alle ihre Knäuel effizient unterscheiden, doch lokal endliche Graphen mit dicken Enden müssen keine solche Baumzerlegungen zulassen. Wir zeigen, dass jeder lokal endliche Graph ohne dicke Enden eine solche Baumzerlegung hat, und zwar sogar eine kanonische. Unser Beweis zeigt, dass in jedem Limit, welches verhindert, dass eine solche Baumzerlegung existiert, ein dickes Ende lebt. Der Beweis baut auf neuen Methoden zur Analyse des Grenzverhaltens streng monoton steigender Sequenzen von Teilungen auf.

Kapitel 10: Über Knäuel induzierende Eckenmengen

Diestel, Hundertmark und Lemanczyk stellten die Frage, ob jedes k-Knäuel in einem Graphen von einer Eckenmenge durch Mehrheitsbeschluss induziert wird. Wir reduzieren deren Frage auf Graphen, deren Größe durch eine Funktion in k begrenzt ist. Zusätzlich zeigen wir, dass, wenn dieses Problem für ein festes k eine positive Antwort hat, dann wird jedes k-Knäuel durch eine Eckenmenge induziert, deren Größe in k beschränkt ist. Allgemeiner beweisen wir für alle k, dass jedes k-Knäuel in einem Graphen G durch eine Gewichtsfunktion $V(G) \to \mathbb{N}$ induziert wird, deren Gesamtgewicht in k beschränkt ist. Als Schlüsselschritt unserer Beweise zeigen wir, dass jedes gegebene k-Knäuel in einem Graphen G die Hebung eines k-Knäuels in einem topologischen Minor von G ist, dessen Größe in k beschränkt ist.

Kapitel 11: Ein Gittertheorem für starke Immersionen von Mauern

Wir zeigen, dass ein Graph eine große Mauer als starken Immersionsminor enthält, genau dann, wenn der Graph keine Baumschnittzerlegung mit kleiner Weite zulässt. Hierbei wird die Weite einer Baumschnittzerlegung anhand der Adhäsion und der Wegähnlichkeit der Torsos gemessen.

Kapitel 12: Der Satz von Lovász und Cherkassky für lokal endliche Graphen mit Enden

Lovász und Cherkassky entdeckten unabhängig voneinander, dass, wenn G ein endlicher Graph und $T \subseteq V(G)$ so ist, dass der Grad $d_G(v)$ für jede Ecke $v \in V(G) \setminus T$ gerade ist, dann ist die maximale Anzahl von kantendisjunkten Wegen, die jeweils intern disjunkt zu T sind und verschiedene Ecken von T verbinden, gleich $\frac{1}{2} \sum_{t \in T} \lambda_G(t, T \setminus \{t\})$ (wobei $\lambda_G(t, T \setminus \{t\})$ die Größe eines kleinsten Schnitts ist, der t und $T \setminus \{t\}$ trennt). Anders betrachtet, bedeutet dies, dass für jede Ecke $t \in T$ in jedem optimalen Wegesystem $\lambda_G(t, T \setminus \{t\})$ viele Pfade zwischen t und $T \setminus \{t\}$ existieren. Wir erweitern den Satz von Lovász und Cherkassky auf der Grundlage dieser Neuformulierung auf alle lokal endlichen unendlichen Graphen und ihre Enden. In unserer Verallgemeinerung kann T nicht nur Ecken, sondern auch Enden enthalten, und Pfade sind einseitig (bzw. zweiseitig) unendlich, wenn sie einen Ecke und ein Ende (bzw. zwei Enden) verbinden.

CHAPTER C

Publications related to this dissertation

Part I:

Chapter 1 is based on [DJKK22].

Chapter 2 is based on [CJKK25].

Chapter 3 is based on [AJKK25].

Part II:

Chapter 4 is based on [ADEFJKW23].

Chapter 5 is based on [AJKW24].

Chapter 6 is based on [AHJKW24].

Part III:

Chapter 7 is based on [AJKP24a]

Chapter 8 is based on [AJKP24b].

Chapter 9 is based on [JK24].

Chapter 10 is based on [ABJKW24].

Chapter 11 is based on [DJKW23].

Chapter 12 is based on [JJKKM23].

CHAPTER D

Declaration of my 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.

Chapter 1 is based on [DJKK22], which is joint work with Reinhard Diestel, Raphael W. Jacobs and Jan Kurkofka. The basic idea of the construction of the canonical graph-decomposition by folding a specific canonical tree-decompositions of the local covering is due to Reinhard Diestel and Johannes Carmesin. The whole theory in all its details was then worked out by Raphael W. Jacobs, Paul Knappe and Jan Kurkofka. In particular, Raphael W. Jacobs and Paul Knappe developed the proof presented in Section 1.5 and the following sections. The final presentation is due to all authors.

Chapter 2 is based on [CJKK25], which is joint work with Johannes Carmesin, Raphael W. Jacobs, and Jan Kurkofka. Raphael W. Jacobs, Jan Kurkofka and I developed the overarching idea for a theory of local separations, drawing inspiration from the work of Johannes Carmesin, but independently of it. The idea to use displacement as a sufficient criterion for settings in which the theory of local coverings works particularly well and also the construction of the respective classes of examples (Section 2.5), is due to Raphael W. Jacobs and me. Raphael W. Jacobs, Paul Knappe and Jan Kurkofka have developed Sections 2.3, 2.4 and 2.6 and worked out their details. Paul Knappe and Jan Kurkofka have developed Section 2.7 and worked out its details. Paul Knappe developed Section 2.8, whose details Paul Knappe and Jan Kurkofka worked out. Paul Knappe and Jan Kurkofka then worked out a first presentation, and the final presentation is due to Raphael W. Jacobs, Paul Knappe and Jan Kurkofka.

Chapter 3 is based on [AJKK25], which is joint work with Tara Abrishami, Raphael W. Jacobs and Jonas Kobler. The idea to study graphs with chordal local coverings via graph-decompositions into cliques is due to Raphael W. Jacobs and me (Section 3.14). Their early work in particular included the proof of the existence of a canonical tree-decomposition into cliques for chordal, locally finite graphs (Section 3.13). With this as input, Jonas Kobler wrote a master thesis supervised by myself, in which he mainly developed a detailed outline for the proof of Theorem 6, from which Tara Abrishami and I further developed the final proofs of Section 3.2. I developed the ideas for Sections 3.7–3.9 and worked out the details. Raphael W. Jacobs and I developed the idea for Section 3.10.2 and I worked out the details. The ideas for all other sections in this chapter have been developed by Tara Abrishami and myself, whose

details we have also worked out together. The final presentation is due to Tara Abrishami and myself.

Chapter 4 is based on [ADEFJKW23], which is joint work with Sandra Albrechtsen, Reinhard Diestel, Ann-Kathrin Elm, Eva Fluck, Raphael W. Jacobs, 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 21 and 22, where we excluded 3-geodesic subdivision of K^3. Ann-Kathrin Elm came up with and proved Lemma 4.4.4 which improved our earlier version of Theorem 21 to the one presented in this thesis. Sandra Albrechtsen proved Theorem 23. Sandra Albrechtsen, Ann-Kathrin Elm, Eva Fluck, Raphael W. Jacobs and I worked out the details of the proofs and finalised a first version [ADEFJKW23]. The final presentation of the results in Sections 4.1 and 4.3 in this thesis that takes into account several recent developments in coarse graph theory is due to Sandra Albrechtsen, Raphael W. Jacobs and myself.

Chapter 5 is based on [AJKW24], which is joint work with Sandra Albrechtsen, Raphael W. Jacobs, 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 25 and 26, but were unsuccessful in our attempts. Sandra Albrechtsen then proved Theorems 25 and 26 building on some of these ideas; in particular, the proof of Lemma 5.5.2 is completely by Sandra Albrechtsen. I have formalised the inductive proof method described in Section 5.3. Sandra Albrechtsen and I worked out the details of the proofs of Theorems 25 and 26 and developed the final presentation.

Chapter 6 is based on [AHJKW24], which is joint work with Sandra Albrechtsen, Tony Huynh, Raphael W. Jacobs, Paul Wollan. This project was started during a 3-weeks research visit of Paul Wollan in Hamburg in June and July 2022. Sandra Albrechtsen, Raphael W. Jacobs, Paul Wollan and I have in close collaboration developed Conjecture 6.1.3 and tried various proof ideas Theorem 6.1.2, but were unsuccessful in our attempts. Tony Huynh and Paul Wollan then proved Theorem 6.1.2 building on some of our ideas. Sandra Albrechtsen spotted and repaired a major error in the construction of the two paths P_1, P_2 starting on page 261. Sandra Albrechtsen, Raphael W. 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 261.

Chapter 7 is based on [AJKP24a], which is joint work with Sandra Albrechtsen, Raphael W. Jacobs, Max Pitz. Max Pitz started the project by asking Raphael W. Jacobs and me 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, we worked out the details. I 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 7.7). We presented this result to the Discrete Mathematics group in Hamburg, and asked whether one can improve the tree-decomposition to be even into finite parts. Sandra Albrechtsen suggested considering critical vertex sets. She found a tree-decomposition displaying the critical vertex sets (Section 7.5),

which she suggested to refine with the previous linked, componental, tight tree-decomposition into rayless to obtain the final tree-decomposition into finite parts (Section 7.4.3). Theorem 30 and Corollary 32 are due to Sandra Albrechtsen and the appliccation in Corollaries 7.1.2 and 7.1.3 are due to Max Pitz. Sandra Albrechtsen, Raphael W. Jacobs 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 Sandra Albrechtsen and me.

Chapter 8 is based on [AJKP24b], which is joint work with Sandra Albrechtsen, Raphael W. Jacobs and Max Pitz. Sandra Albrechtsen developed Construction 8.3.1 and proved Examples 35 and 36. For Example 37, Max Pitz suggested considering the graph from [Car19, Example 3.7], and Sandra Albrechtsen proved that it indeed works. Raphael W. Jacobs, Max Pitz and I developed Theorem 8.5.1. Sandra Albrechtsen and I worked out the details of the proofs.

Chapter 9 is based on [JK24], which is joint work with Raphael W. Jacobs. Raphael W. Jacobs and I developed all the proofs and their details together as well as the complete presentation.

Chapter 10 is based on [ABJKW24], which is joint work with Sandra Albrechtsen, Hanno von Bergen, Raphael W. Jacobs, 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. Sandra Albrechtsen and Hanno von Bergen developed a more detailed outline of the proof. Afterwards, Sandra Albrechtsen, Raphael W. Jacobs and I worked out the full details of the proofs and developed the final presentation.

Chapter 11 is based on [DJKW23], which is joint work with Reinhard Diestel, Raphael W. Jacobs and Paul Wollan. During a 3-weeks research visit of Paul Wollan in Hamburg in October 2021, we developed an outline of the proof of the main result. Raphael W. Jacobs and I then worked out all details and developed a complete presentation.

Chapter 12 is based on [JJKKM23], which is joint work with Raphael W. Jacobs, Attila Joó, Jan Kurkofka and Ruben Melcher. The authors developed the proof idea during a series of online meetings. Raphael W. Jacobs, Jan Kurkofka and I then worked out the details and developed the complete presentation.

CHAPTER E

Danksagungen

Mein erster Dank gilt meinem Betreuer Reinhard Diestel für seine geduldige Unterstützung und Förderung. Weiterhin bedanke ich mich bei Max Pitz, der mich während der Betreuung meiner Bachelorarbeit für die mathematische Forschung begeistert hat, sowie bei Jan Kurkofka, der dieses Interesse während der Betreuung meiner Masterarbeit weiter verstärkt hat. Ohne alle drei hätte ich diesen Weg nicht eingeschlagen.

Ich bedanke mich bei all meinen Koautoren für die erfolgreiche Zusammenarbeit. Außerdem möchte ich mich bei all den Mitstudierenden bedanken, durch die das gesamte Studium viel Spaß gemacht hat; explizit möchte ich davon Johann Bellmann, Kamil Bredow, Dennis Breutigam, Hugo Eberhard, Florian Gut, Kirke Joamets, Insa Killian, Kevin Sames, Tim Seifert und Tim Vieth nennen. Der gesamten Arbeitsgruppe Diskrete Mathematik der Universität Hamburg danke ich für die angenehme Atmosphäre. Besonders bedanke ich mich bei meinen Mitdoktorand:innen Sandra Albrechtsen und Raphael W. Jacobs für die Zeit im gemeinsamen Büro, die moralische Unterstützung und die intensive inhaltliche Zusammenarbeit.

Besonders bedanke ich mich bei meinen engsten Freund:innen, die immer ein offenes Ohr für mich haben und mich mental unterstützen; hiervon möchte ich explizit Finn Hansen und Anne Heinzer nennen. Zuletzt gilt mein größter Dank meiner Familie, Sonja, Ulrich und Lars Knappe, denn ohne ihre unbedingte Unterstützung wäre es mir nicht möglich, das zu tun, was ich möchte.

CHAPTER F

Eidesstattliche Versicherung

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.

_____

Paul Knappe